

University of Texas at Arlington

MavMatrix

Electrical Engineering Dissertations

Department of Electrical Engineering

2023

Hardware in the Loop Test Stand for Battery Management System Validation in Shipboard Controls

Cole Dawson Tschritter

Follow this and additional works at: https://mavmatrix.uta.edu/electricaleng_dissertations



Part of the [Electrical and Computer Engineering Commons](#)

Recommended Citation

Tschritter, Cole Dawson, "Hardware in the Loop Test Stand for Battery Management System Validation in Shipboard Controls" (2023). *Electrical Engineering Dissertations*. 385.

https://mavmatrix.uta.edu/electricaleng_dissertations/385

This Dissertation is brought to you for free and open access by the Department of Electrical Engineering at MavMatrix. It has been accepted for inclusion in Electrical Engineering Dissertations by an authorized administrator of MavMatrix. For more information, please contact leah.mccurdy@uta.edu, erica.rousseau@uta.edu, vanessa.garrett@uta.edu.

HARDWARE IN THE LOOP TEST STAND FOR BATTERY MANAGEMENT
SYSTEM VALIDATION IN SHIPBOARD CONTROLS

By Cole Tschritter

A dissertation
submitted in partial fulfillment
of the requirements for the degree of
DOCTOR OF PHILOSOPHY IN ELECTRICAL ENGINEERING

UNIVERSITY OF TEXAS AT ARLINGTON

August 2023

© 2023

Cole Tschritter

ALL RIGHTS RESERVED

University of Texas at Arlington

DEFENSE COMMITTEE AND FINAL READING APPROVALS

of the dissertation submitted by

Cole Tschritter

Dissertation Title: Hardware in the Loop Test Stand for Battery Management
System Validation in Shipboard Controls

Date of Final Oral Examination: 07 August 2023

The following individuals read and discussed the dissertation submitted by student Cole Dawson Tschritter, and they evaluated the student's presentation and response to questions during the final oral examination. They found that the student passed the final oral examination.

David Wetz, Ph.D. Chair, Supervisory Committee

Wei-Jen Lee, Ph.D. Member, Supervisory Committee

Rasool Kenarangui, Ph.D. Member, Supervisory Committee

The final reading approval of the dissertation was granted by David Wetz, Ph.D., Chair of the Supervisory Committee. The dissertation was approved by the Graduate College.

DEDICATION

I would like to dedicate this work to my parents, Gordon and Ellen Tschritter. They encouraged me consistently throughout my life to reach my best. I couldn't have accomplished any of this without you both. I love you mom and dad.

I would also like to acknowledge and give thanks to Don Bray for planting the seed of passion for electrical engineering into my life. Thank you for your continued friendship, advice, and support.

Finally, I would like to give a big thanks to Pavani Rambachan, for her unwavering support. Thank you for being a wonderful friend and partner, life with you is like clockwork.

ACKNOWLEDGMENTS

I would like to sincerely thank my dissertation committee members, Dr. David Wetz, Dr. Wei-Jen Lee, and Dr. Rasool Kenarangui for their time, support, and guidance throughout my graduate studies.

I would like to thank ONR for their financial support of this effort. Their support through grants N00014-18-1-2714, N00014-16-1-2956, and N00014-21-1-2124 enabled all of the technical work on the IDEAL testbed to be possible. Any opinions and findings are those of the author and not those of ONR or NSWC-Philadelphia.

Thank you also to General Technical Services (GTS) for their support under subcontract number GTS-S-20-156. The prime sponsor is NSWC-Philadelphia Division (NSWC-PD) and I thank them for their support. Their support allowed all the work on the battery emulator to be possible. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the GTS or NSWC-PD.

COPYRIGHT

In reference to IEEE copyrighted material which is used with permission in this dissertation, the IEEE does not endorse any of UTA's products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to http://www.ieee.org/publications_standards/publications/rights/rights_link.html to learn how to obtain a License from RightsLink. If applicable, University Microfilms and/or ProQuest Library, or the Archives of Canada may supply single copies of the dissertation.

ABSTRACT

Traditionally, large-scale power systems have relied on power generation always exceeding the total load demand. Though this is most often the case, there are rare occurrences where failures in the generation, distribution, or even planning, that have led to situations where load demand quickly surpasses the generation capacity. In such cases, uncontrolled shedding of loads occurs, leading to brownouts or blackouts that can have catastrophic impact to both people and property. Advanced architectures, such as microgrids, offer a variety of improvements that are aimed at mitigating or even preventing these failures. Improvements include distributed power generation, active source and load monitoring, manageable load control, enhanced resilience against distribution failures, and in some instances the integration of energy storage to buffer transient loads. Despite all these benefits, system complexity is massively increased and heightened system monitoring is required, which is costly and difficult to implement.

Shipboard power system architectures resemble microgrids and the US Navy has proposed zonal shipboard power systems that employ low voltage (LV) AC, medium voltage (MV) AC, and even MVDC nodes within the same architecture. To both evaluate and validate such architectures, the Intelligent Distributed Energy Analysis Laboratory (IDEAL) testbed was established at the University of Texas at Arlington (UTA). IDEAL is intended to emulate one zone of a shipboard power system architecture that introduces various generation sources, power electronic converters, loads, and LFP energy storage, all of which are designed and assembled at naval relevant voltages. The platform is set up to enable Hardware in the Loop (HIL) model emulation allowing for the development, testing, and validation of control algorithms in a flexible emulated environment. Such an

environment also allows for collaborative opportunities, thus facilitating a partnership with institutions like Florida State University (FSU), University of South Carolina (USC), and Clarkson University (CU), further enhancing research and innovation in the field.

Realizing the full potential of zonal architectures and the associated algorithms needed to reliably operate them, may require effective utilization of energy storage. However, safety concerns surrounding lithium-ion energy storage has many weighing the risks against the benefits it introduces. Lithium-ion batteries are most often managed using battery management systems (BMS) that monitor and manage the state of charge (SoC) of the multitude of cells that make up the battery. As battery configurations expand and as new BMSs are designed and introduced into the market, the need to study and validate their operation before they introduced into real batteries is critical. Furthermore, the ability to interface the BMS with the overarching system level controller is essential, and its operation must be validated across all potential use cases where intervention may be required. Power HIL (PHIL) platforms offer unique capabilities for emulating batteries comprised of multiple lithium-ion cells. Such capabilities include increased flexibility for rapidly studying the BMS's reaction to normal and abnormal operating conditions, as well as a safe controlled environment for these types of scenarios. Thus, leading to increased user confidence and hopefully lead to the wider scale deployment of energy storage in shipboard power systems.

The work performed in this dissertation comprises of a few different, but interrelated thrusts. In the first, it discusses the design, rational, assembly and results obtained from a PHIL testbed used to validate BMS performance and software integration operating on batteries with up to 264 cells in series, ~1 kVDC. In the second thrust, a

collaborative effort performed by UTA, FSU and USC is discussed in which the UTA IDEAL testbed was interfaced with remote HIL platforms being operated at FSU and USC, respectively. The remote HIL co-simulation effort demonstrated the employed physical energy storage at UTA and it was used to provide ramp rate buffering of a 12 kVDC bus emulated in the co-simulated HIL. In the third and final thrust, the IDEAL testbed is used to demonstrate the effectiveness of predictive high ramp rate (PHRR) and advanced load shed (ALS) algorithms developed by Clarkson University (CU) for maintaining operability and power quality within shipboard power systems deploying continuous and transient loads. Each of these thrusts will be discussed in detail.

TABLE OF CONTENTS

DEDICATION	iv
ACKNOWLEDGMENTS	v
COPYRIGHT	vi
ABSTRACT	vii
TABLE OF CONTENTS	x
LIST OF TABLES	xiii
LIST OF FIGURES	xiv
LIST OF ABBREVIATIONS	xix
INTRODUCTION	1
BACKGROUND	4
Battery Management Systems	4
Ideal Experimental Testbed	11
BMS TESTBED EVALUATION	14
Direct Deployment onto Battery	14
Independent Power Supplies Utilizing Analog Following	15
Commercial Battery Emulators	17
Chroma 87001	18
Hioki SS7081-50	19
Speedgoat IO991	20

Speedgoat BCE unit.....	22
Hardware Option Comparison	24
Hardware in the Loop Models	24
Lithium Cell Discharge.....	28
HIL VL30AF Deployment.....	30
IDEAL TESTBED DEPLOYMENTS.....	32
Co-Simulation with FSU and USC.....	32
Co-Simulation Framework.....	34
UTA Site Zonal HIL	36
ALS and PHRR Control Methodology.....	40
ALS Control Methodology	41
PHRR Control Methodology	45
Integration of Clarkson ALS and PHRR controls into IDEAL	52
RESULTS	58
BMS Test Stand Results	58
Emulation with IO991.....	60
Emulation with Speedgoat BCE	73
USC and FSU Cosimulation.....	78
ALS and PHRR Load Scenarios.....	81
ALS Scenario 1: Load priority change and generator overload	81
ALS Scenario 2: Generator loss of capacity	85
PHRR Scenario 3: Rep-Rate Ramped Mission Load	88
PHRR Scenario 4: Ramped generation overload.....	90

ALS/PHRR Scenario 5: Ramped generation overload	92
CONCLUSION.....	95
References.....	97

LIST OF TABLES

Table 1. Summary of the hardware capabilities offered by Chroma, Hioki, and Speedgoat, respectively.	24
Table 2. Parameters of loads and generator for case scenario 1	82
Table 3. Parameters of loads and generator for case scenario 2. (© 2023 IEEE).....	85

LIST OF FIGURES

Figure 1. Battery and BMS architecture. (© 2021 IEEE).....	6
Figure 2. Example of passive balancing topologies. (Left) Bleed resistor with externally driven MOSFET. (Middle) Pass transistor biased with zener diode. (Right) Bleed resistor with locally driven MOSFET using comparator. (© 2022 IEEE).....	7
Figure 3. Example of active balancing topologies. (Left) Inductor based boost mode active balancing. (Right) Capacitor based charge shuttle active balancing. (© 2022 IEEE)	9
Figure 4. UTA IDEAL testbed one-line diagram. (© 2023 IEEE).....	13
Figure 5. Visual one-line diagram of a testbed layout using multiple individual power supplies to emulate cells. (© 2021 IEEE).....	16
Figure 6. Visual one-line diagram of a testbed layout using multiple Chroma 87001 cell emulators. (© 2021 IEEE)	18
Figure 7. Visual one-line diagram of a testbed layout using multiple Hioki SS7081-50 cell emulators. (© 2021 IEEE).....	20
Figure 8. Visual one-line diagram of a testbed layout using multiple Speedgoat IO991 cell emulators. (© 2021 IEEE).....	21
Figure 9. Visual one-line diagram of a testbed layout using multiple Speedgoat BCE units.....	23
Figure 10. Model showing the 10 cell modules connected to a load profile and temperature profile. These modules can then be further connected in series or parallel to form a battery. (© 2021 IEEE)	27
Figure 11. Lithion ion cells with different starting SoCs (93±2%) discharging at 1C rate. This slight change in SoC can cause dramatic changes in the voltage over time. (© 2021 IEEE).....	29
Figure 12. Lithion ion battery shows the sum 264 cell string discharging at 1 C rate with cells at different starting SoCs (93±2%). (© 2021 IEEE).....	29

Figure 13. Custom MATLAB application used to change the SoC of each individual cells in real time. (© 2021 IEEE).....	31
Figure 14. Setup for co-simulation with USC, FSU and UTA(© 2022 IEEE).....	33
Figure 15. One line Diagram of connections between real time simulators, desktop computers, and the Internet for cosimulation(© 2022 IEEE).....	35
Figure 16. Connection between simulators and universities using Villas node and Tinc VPN.....	36
Figure 17. Hardware configuration used by UTA IDEAL for PHIL deployment in co-simulation. (© 2022 IEEE).....	39
Figure 18. Main VI used to command supplies and loads in the IDEAL testbed. (© 2023 IEEE).....	42
Figure 19. Integration the UTA IDEAL hardware and LabVIEW control system with CU's SLDRT controller. (© 2023 IEEE).....	43
Figure 20. Block diagram of the load shedding control. (© 2023 IEEE).....	44
Figure 21. Predictive high ramp rate controller block diagram. (© 2023 IEEE).....	47
Figure 22. Flow chart diagram of the PHRR control. (© 2023 IEEE).....	48
Figure 23. Diagram of battery protection mechanism. (© 2023 IEEE).....	52
Figure 24. NI LabVIEW system dedicated to handling inputs and outputs for the deployment of the ALS controller from Clarkson.	55
Figure 25. LabVIEW subsystem inside of the IDEAL controller that handles the integration of the PHRR system.	57
Figure 26. Predicted BMS behaviour with 2 cells in normal operation. (© 2021 IEEE) .	59
Figure 27. Predicted BMS behaviour with 2 cells. The graph demonstrates the user's ability to rapidly change individual cell(s) SoC to evaluate BMS behaviour quickly. (© 2021 IEEE).....	60
Figure 28. Simplified Circuit diagram of the battery and the current measurement to be implemented.....	61
Figure 29. Current sense PCB attached to a BMS and IO991 modules.	61

Figure 30. Voltage and current data collected while 5 emulated cells are connected and balanced by the BMS. Solid lines indicate cell current while dashed lines represent cell voltages. (© 2021 IEEE) 63

Figure 31. The graph demonstrates the ability to rapidly change a cell(s) SoC to evaluate a BMS’s performance near cell bounds quickly. Solid lines indicate cell current while dashed lines represent cell voltages. (© 2021 IEEE)..... 64

Figure 32. Image of K2 BMS board and current sense board (in black). The green card to the left is the Slave card. The master is the white PCB in the middle. The blue board is for communications. The K2 BMS demonstrates operational patterns consistent with the anticipated behavior illustrated in Figure 27. Unlike the Actia BMS in attempting to balance cells towards the lowest recorded voltage, the BMS endeavors to achieve a designated balance threshold. Evaluating the K2 BMS's balancing functionality necessitates a modified testing approach. In this case, the cells are charged until they reach a maximum threshold, subsequently prompting the BMS to balance the cells down to the predetermined balance threshold before initiating another charging cycle. Notably, the recharging sequences in the K2 BMS are emulated within software, avoiding an external power supply to avert potential external loading effects. This approach is employed to ensure a controlled environment. A representative instance of a test involving the IO991 card on the K2 BMS is depicted in Figure 33. 65

Figure 33. K2 balance test. This BMS requires cycles of recharging after balancing to bring the cells into proper balance. 66

Figure 34. Chroma unit in white. It is a 16 cell emulator capable of up to 5 amps continuous current on each channel. 68

Figure 35. Toolbox to communicate with instrument in MATLAB..... 69

Figure 36. Block diagram depicting use of TCP blocks to connect to Chroma 87001..... 70

Figure 37. Chroma TCP connect for standard Ethernet interface..... 71

Figure 38. Chroma system balancing from K2 BMS. Top graph contains the values commanded by Simulink while the bottom graph contains the voltage values recorded by the BMS. 72

Figure 39. Graphs of voltages from K2 BMS with Chroma emulator. Graph shows the voltage of cells falling until their output drops to zero and current is no longer drawn. 73

Figure 40. Graphs of currents from K2BMS with Chroma emulator. Graph shows constant current being drawn from several cells until their output drops to zero and current can no longer be maintained.	73
Figure 41. Completed circuit board to attach the cells in series and interface with the BMS.	75
Figure 42. Front and back of one of the BMS testbed cabinets	77
Figure 43. BCE software to enable multiple modules and control them efficiently at once.	78
Figure 44. The specific subsystem that forms the ramp rate for cosimulation.	79
Figure 45. Co-Simulation results with FSU and USC. Bottom series of graphs contain the unbuffered bus, The top series of graphs are the buffered version, and results with reduction of spikes on the PGM. (© 2022 IEEE).....	81
Figure 46. The simulated results of scenario 1 on the ALS controller. (© 2023 IEEE)...	83
Figure 47. Experimental load values measured by the ALS controller. (© 2023 IEEE)..	84
Figure 48. Experimental data collected by the PXI chassis during the operation of scenario 1. (© 2023 IEEE).....	84
Figure 49. The simulated results of the ALS controller for scenario 2. (© 2023 IEEE) ..	86
Figure 50. Experimental load values measured by the LabVIEW controller during scenario 2. (© 2023 IEEE).....	87
Figure 51. Experimental data collected by the PXI chassis during the operation of ALS scenario 2. (© 2023 IEEE).....	87
Figure 52. Simulation of the PHRR controller operating in a repeated mission load scenario 3. (© 2023 IEEE).....	88
Figure 53. Experimental load values measured by the LabVIEW controller while the PHRR controller operates on the IDEAL testbed during scenario 3. (© 2023 IEEE).....	89
Figure 54. Experimental data collected by the PXI chassis during the operation of PHRR scenario 3. (© 2023 IEEE).....	90
Figure 55. Experimental values of TDK, mission load, and battery current measured on IDEAL when battery is in an unsafe operating condition in scenario 4. (© 2023 IEEE).....	91

Figure 56. Experimental values of battery voltage and TDK current measured on IDEAL when battery is in an unsafe operating condition in scenario 4. (© 2023 IEEE)..... 91

Figure 57. Experimental data collected by the PXI during scenario 4. (© 2023 IEEE)... 92

Figure 58. The simulated results of running the PHRR controller and ALS controller for scenario 5. (© 2023 IEEE)..... 93

Figure 59. Experimental load values measured by the LabVIEW controller during the combined scenario 5. (© 2023 IEEE)..... 94

Figure 60. Experimental data collected by the PXI chassis during the operation of scenario 5. (© 2023 IEEE)..... 94

LIST OF ABBREVIATIONS

AC	Alternation Current
BCE	Battery Cell Emulator
BMS	Battery Management System
CDAQ	Compact Data Acquisition
DC	Direct Current
DG	Distributed Power Generators
ESM	Energy Storage Management
FSU	Florida State University
IDEAL	Intelligent Distributed Energy Analysis Laboratory
LFP	Lithium-Iron-Phosphate
LNG	Liquified Natural Gas
LV	Low Voltage
MATLAB	Matrix Laboratory
MOSFET	Metal Oxide Semiconductor Field Effect Transistor
MPC	Model Predictive Control
MV	Medium Voltage
NI-MAX	National Instruments Measurement and Automation eXplorer
NSWC-PD	The Naval Surface Warfare Center- Philadelphia Division
ONR	Office of Naval Research
PCB	Printed Circuit Board

PCI	Peripheral Component Interconnect
PCM	Power Conversion Module
PGM	Power Generation Module
PXI	PCI eXtension
SLDRT	Simulink Desktop Real Time
SoC	State of Charge
SPS	Shipboard Power System
SR	Spinning Reserve
USC	University of South Carolina
VFD	Variable Frequency Drive
VI	Virtual Instrument
VSD	Variable Speed Drive

INTRODUCTION

Diesel-electric propulsion was initially introduced in the Vandal vessel in 1903. Subsequently, during the 1920s, there was a notable surge in the construction of ships integrating electric propulsion on a larger scale. In this arrangement, steam turbine generators were harnessed to power the propeller motors. The rotational speed of the propeller motor was skillfully controlled by manipulating the generator's speed [1]. However, the prominence of this technology waned swiftly as diesel engines took precedence across the maritime industry.

It wasn't until the emergence of power electronic devices in the 1980s that a revitalization of the all-electric ship concept was witnessed. This resurgence was further amplified with the advent of viable variable speed drives (VSDs) in the 1990s [2]. Over the last few decades, an array of comprehensive investigations has been conducted concerning shipboard power system architectures. These studies have explored the potential employment of various topologies, such as zonal medium voltage direct current (MVDC), medium voltage alternating current (MVAC), or a hybrid combination thereof. Encouraging strides have been taken in recent times, suggesting a promising outlook for wider implementation soon.

Presently, there are already instances of cruise vessels, icebreakers, and liquefied natural gas (LNG) tankers employing electric propulsion, featuring a spectrum of AC and DC loads. This operational framework, encompassing both constant and pulsed power modes [3], has provided valuable case studies. These practical scenarios underscore both

the merits and drawbacks inherent to adopting a microgrid architecture. It is important to note, however, that substantial efforts are imperative to fully harness the latent potential of these architectural paradigms. The subsequent discourse represents a select few among the multifaceted endeavors required to propel these innovations forward.

The integration of energy storage within a microgrid structure brings forth a significant advantage, enabling the implementation of algorithms, like ramp rate control. These algorithms, designed to enhance power quality on the bus, demonstrate their utility. Notably, batteries play a pivotal role in realizing this potential. Despite the wide-ranging applications of batteries in diverse electronic devices and their admirable safety record, it is crucial to emphasize that improper monitoring and neglect of specific management practices can introduce risks. This underscores the importance of employing a Battery Management System (BMS) to ensure detailed oversight.

This necessity for diligent management becomes even more pronounced as the number of individual cells within battery units escalates. In this context, the monitoring and control exercised through a BMS becomes paramount. With the plethora of available BMS styles, it becomes crucial that these systems undergo thorough evaluation before their deployment into shipboard control systems. This stringent vetting process is essential to uphold the integrity and safety of the overall power system.

An additional compelling rationale for the adoption of microgrid architecture arises from the widespread utilization of transiently operated loads. These dynamic loads exert fluctuating demands on a generator's bus voltage, inducing sagging and surging, and causing frequency deviations. Such fluctuations have the potential to displace the voltage and current beyond the established thresholds dictated by pertinent power quality

standards. The operation of numerous transient loads also introduces the risk of surpassing the power generation capacity. In instances where the available power falls short of meeting the cumulative load requirements, a rapid shedding of non-essential loads becomes imperative to safeguard the continuity of critical loads [4, 5].

In the absence of real-time monitoring and control mechanisms, load shedding might transpire in an arbitrary manner, devoid of systematic considerations regarding its impact on overall operability. With the integration of sensing and control capabilities, a certain degree of latency emerges between the two elements, significantly influencing the optimization of power quality, ramp rate support, and load shedding efficiency. The undertaking detailed herein involves collaborative efforts between researchers at the University of Texas at Arlington (UTA), Florida State University (FSU) in Tallahassee, Florida, the University of South Carolina (USC) in Columbia, South Carolina, Clarkson University (CU) in Potsdam, New York, and the Naval Surface Warfare Center - Philadelphia Division (NSWC-PD) in Philadelphia, Pennsylvania. These entities have collaborated to demonstrate a multitude of control algorithms within the framework of the Intelligent Distributed Energy Analysis Laboratory (IDEAL) power system testbed, situated at the UTA campus [6]. This report provides a comprehensive discussion of the control strategies implemented on the IDEAL testbed, along with the establishment of a dedicated testbed for the evaluation of BMSs. The latter plays a pivotal role in enabling the integration of electrochemical energy storage into shipboard controls, as batteries and their BMS constitute fundamental components of various control strategies.

BACKGROUND

Battery Management Systems

Electrochemical batteries play a pivotal role in the success of controllers used to reliably and efficiently operate microgrid power systems. While batteries find widespread applications across many electronic devices and boast a commendable safety track record, it remains critical to underscore that they can pose risks if not subjected to proper monitoring and vigilant management practices. The necessity for meticulous management becomes especially pronounced with the escalation in the number of individual cells within a battery assembly.

In practice, high-voltage and high-energy battery configurations are materialized by interconnecting multiple cells in a configuration of n series by m parallel connections. Demonstrations have showcased battery voltages reaching a threshold of 1 kVDC [7, 8, 9]. To attain this voltage, a composition of nearly 270 cells with a nominal rating of 3.3V are arranged in a series configuration. Such battery architectures are presently under consideration for deployment as power sources within prospective zonal shipboard power system designs [7, 9], as well as in vehicular applications [8], among other promising use cases.

Of particular significance is the role that batteries play in buffering transient loads or prioritized loads, especially in scenarios where an alternative power source is unavailable. This buffering function holds pivotal importance of electrical loads aboard next-generation vessels. Employing batteries to buffer power sources yields a cascade of advantages, such as: the enhancement of power quality, attenuation of harmonics, potential gains in fuel efficiency, and the creation of temporal leeway for generators to respond to

abrupt shifts in power demand, thereby optimizing their output. This multifaceted utility underscores the salient role of energy storage systems, particularly batteries, in forging a firm and adaptive framework for maritime and vehicular electrification.

The oversight and control of individual battery cells are achieved through the utilization of a BMS. The responsibilities entrusted to BMSs are many with a focus on preserving charge equilibrium among cells interconnected within a series configuration. This crucial function serves to mitigate undue aging effects and avert hazardous operational circumstances [10]. An illustrative representation of the typical architecture characterizing a BMS system for high-voltage battery modules is depicted in Figure 1.

The design of this system involves the organization of individual cells into modules, structured as nS/mP configurations. Each discrete cell, or a cluster of cells interconnected in series, is subject to continuous monitoring facilitated by a dedicated balance circuit. At a more granular level, numerous monitoring channels are amalgamated onto lower-tier boards, with each module of cells being overseen by one of these channels. To scale the design as necessary, multiple lower-tier boards are serially interconnected, interfacing with an upper-tier board. This upper-level component establishes communication links with a host controller or supplementary electronic components, culminating in a cohesive hierarchical structure that underpins the functioning of the BMS [10].

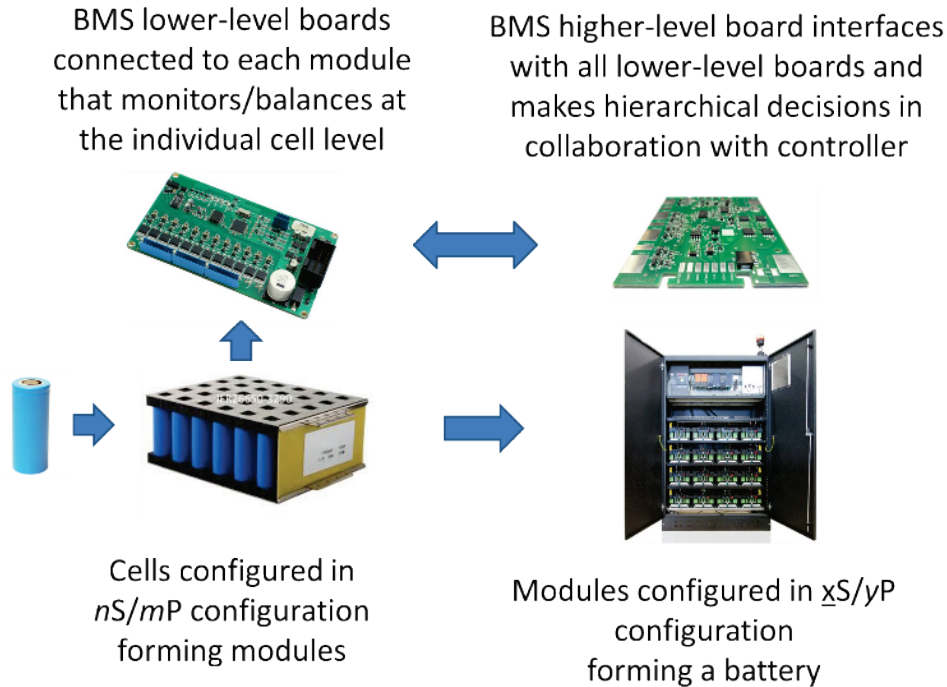


Figure 1. Battery and BMS architecture. (© 2021 IEEE)

In lithium-ion batteries, the process of balancing frequently involves employing a resistor in conjunction with cells exhibiting elevated voltage. This arrangement facilitates the dissipation of surplus energy in the form of heat. This approach is notable for its cost-effectiveness and established technical maturity. Nevertheless, there exist certain limitations to consider. Notably, this technique demonstrates inefficiency in energy utilization and primarily addresses over-voltage occurrences. Alternative strategies for battery management have been explored, yet they remain comparatively less mature and necessitate further investigation and cost reduction endeavors.

Numerous BMS configurations have been devised, all serving the common purpose of passive balancing through the utilization of a series transistor with or without an additional bleed resistor. These designs span a spectrum, ranging from rudimentary setups such as simple Zener diode-biased transistor bleed circuits, to more intricate systems

employing comparator-driven over-voltage detectors, which leverage MOSFETs and dissipation resistors for energy dissipation. A few such configurations are exemplified in Figure 2.




Figure 2. Example of passive balancing topologies. (Left) Bleed resistor with externally driven MOSFET. (Middle) Pass transistor biased with zener diode. (Right) Bleed resistor with locally driven MOSFET using comparator. (© 2022 IEEE)

While certain designs, such as the Zener-biased pass transistor, exhibit dynamic response characteristics by adjusting the balance current in proportion to the degree of over-voltage, none of these analog systems possess the sophistication to customize the balance current based on the extent of imbalance among cells. At best, an output current proportional to the voltage surpassing the threshold can be achieved, albeit in a static manner. Consequently, even when cells attain perfect balance upon reaching the threshold, all balancing circuits remain active. In turn, chargers must be designed intelligently to accommodate this ‘extra’ current during the constant current (CC) – constant voltage (CV) charging algorithm and the subsequent termination stage.

It is pertinent to note that BMSs typically engage in power dissipation at comparatively modest rates. Common balance currents generally range from a few milliamperes to a few hundred milliamperes, at most. In scenarios where the charge current exceeds the dissipation capacity of cells with elevated voltages, the BMS has limitations

in bleeding current rapidly enough to keep the cells safe. Consequently, charge cutoffs are frequently initiated before achieving uniform charge distribution among all cells. BMS variants capable of accommodating higher balance currents, on the order of a few amperes at most, are available. However, it's important to consider that these solutions can introduce bulkiness, elevated cost, and potential cooling requirements if the thermal design is insufficiently addressed. The limitations inherent in passive balancing approaches and configurations are widely recognized.

In systems characterized by high power, the discharge currents during operation are often several magnitudes greater than the recharge currents. As a result, a substantial portion of imbalance occurs during discharge phases, a period during which the balancing circuitry remains inactive. Even when the balancing circuitry is engaged, many passive designs exhibit a binary balance current nature meaning it operates in either an on or off mode. In this case, the balance current is solely contingent on cell voltage proportions. Consequently, cells with varying degrees of imbalance are treated similarly, irrespective of the extent of that imbalance.

It becomes evident that an active system capable of not only transferring energy among cells within the pack but also doing so proportionately, both during charging and discharging, would represent a markedly more efficient and effective approach to sustaining balance within the system. Extensive literature delves into active battery management strategies, encompassing a range of concepts from capacitively coupled charge transfer systems to coupled inductor hybrid DC-DC converter systems. A selection of these configurations is illustrated in

Figure 3.




Figure 3. Example of active balancing topologies. (Left) Inductor based boost mode active balancing. (Right) Capacitor based charge shuttle active balancing. (© 2022 IEEE)

One active topology seen in the left side of Figure 3, known as the inductor-based boost mode balancer, employs boost converters to achieve its intended purpose. It facilitates the charging of an inductor through the switch beneath it, while the body diode within the upper switch permits the transfer of charge to the upper cell once the lower switch is deactivated. While this topology is characterized by its simplicity in terms of component count and its potential for high energy transfer rates, it does possess limitations. Notably, energy transfer can only be facilitated for a single cell at any given time. Furthermore, when extending the topology for bidirectional energy flow, the system's complexity escalates significantly.

The capacitor-based charge shuttling topology is seen on the right side of Figure 3. Despite its notable simplicity in implementation and arguably being the most efficient among active topologies, faces considerable limitations. Balancing in this configuration involves simultaneous operation of all switches, each in either the high or low position. When switched, cells whose voltage is higher than that of the capacitor charge the capacitor, while cells with lower voltage discharge the capacitor. This process continues,

effectively redistributing charge from higher cells to lower ones until balance is attained. This balancing process inherently lacks control thus, it treats every cell uniformly, with balancing being either fully engaged or deactivated. The energy transfer rate is directly tied to capacitor size and the switching frequency. The effectiveness of this approach however diminishes as cell voltages become more aligned. Furthermore, the energy transfer mechanism exhibits a high level of elasticity – the process of moving charge from multiple cells into a single weaker cell is notably slow. These characteristics collectively underscore the constraints of the capacitor-based charge shuttling method.

Beyond their role in cell balancing, BMSs undertake a multitude of critical functions. In instances where higher costs are incurred for BMSs, these systems often possess the capability to individually monitor cell voltages and relay this information to a host controller. Generally, each lower-level board is equipped to record a limited number of temperature measurements derived from thermocouples or thermistors strategically affixed to specific points of interest within the module. This monitoring helps ascertain the presence of temperatures surpassing acceptable thresholds.

In the event that any cell experience over-discharge, over-charge, or if temperatures breach designated thresholds, the BMS can command an output relay that facilitates the isolation of the battery from its source or load until safe conditions can be reinstated. These fail-safe mechanisms underscore the pivotal safety measures that advanced BMSs are poised to enact. Every BMS manufacturer highlights their distinct enhancements in balance topology, control algorithms, communication strategies, cooling technologies, interface methods, and more, all aimed at enticing potential customers to adopt their BMS solutions. Given the absence of universal standards dictating the design and comprehensive testing

of all BMS variations, it becomes imperative for prospective users to possess the means to assess a BMS's short- and long-term performance before its integration into a battery system. The consequences of a BMS failing to perform as advertised can range from substantial financial implications to potential safety hazards, underscoring the significance of their evaluation.

In the scope of this project, a BMS testbed has been developed that employs battery emulation using PHIL technology. This assembly incorporates 264 individual cells, capable of reaching potentials as high as 1 kVDC. The unique capability to dynamically modify the state of charge (SoC) of each individual cell in real-time facilitates the safe and swift evaluation of both standard and anomalous conditions. This adaptive environment serves to bridge the knowledge gap pertaining to potential challenges associated with BMS operation in high-voltage battery setups. Moreover, this initiative contributes to the advancement and validation of controls that will eventually be employed in tandem with a real-world 1 kV battery configuration. This multi-dimensional approach represents a pivotal step toward fostering a comprehensive understanding of BMS performance characteristics, thereby fortifying the groundwork for dependable advance control algorithms.

Ideal Experimental Testbed

The UTA IDEAL testbed was purposefully crafted to replicate a distinct zone within an MV AC/DC electric ship, encompassing LVAC, MVAC, and MVDC sources and loads. Using the IDEAL testbed, researchers gain access to the requisite hardware for validating diverse control algorithms applicable to next-generation electronic shipboard systems, after the verification and comprehensive integration of BMSs on energy storage

systems. This testbed comprises a comprehensive ensemble, comprising an electric motor-generator (M-G) set, an array of AC and DC buses, power converters, transformers, variable resistive loads, and a lithium iron phosphate (LFP) battery. While the majority of IDEAL's functionalities and constituent components have been extensively detailed in prior publications [6, 11, 12, 13, 14], noteworthy enhancements and refinements have been introduced since those were published. At the core of the testbed, an electric M-G set, designated as G1 in Figure 4, assumes role of the primary AC power source. This unit, governed by a variable frequency drive (VFD), furnishes 480 VAC and is rated for 150 kW. The M-G set, operating as a four-pole synchronous generator, sustains a rotational frequency spanning 1500 to 2000 rpm. Its control is entrusted to a four-pole, 300 hp induction motor.

In reference to the one-line diagram depicted in Figure 4, the testbed's load framework encompasses four distinctive Mosebach resistive loads, identified as Load Group 1 (L1), Load Group 2 (L2), Load Group 3 (L3), and Mission Load (ML). Elaborating further, Load Group 1 (L1) is characterized as a 480 VAC, 350 kW resistive load, characterized by an incremental resolution of 1 kW. Load Group 2 (L2) is configured as a 6 kVDC, 150 kW resistive load, with an incremental resolution of 50 kW. Similarly, Load Group 3 (L3) is fashioned as a 12 kVDC, 100 kW resistive load, affording a stepwise resolution of 50 kW.

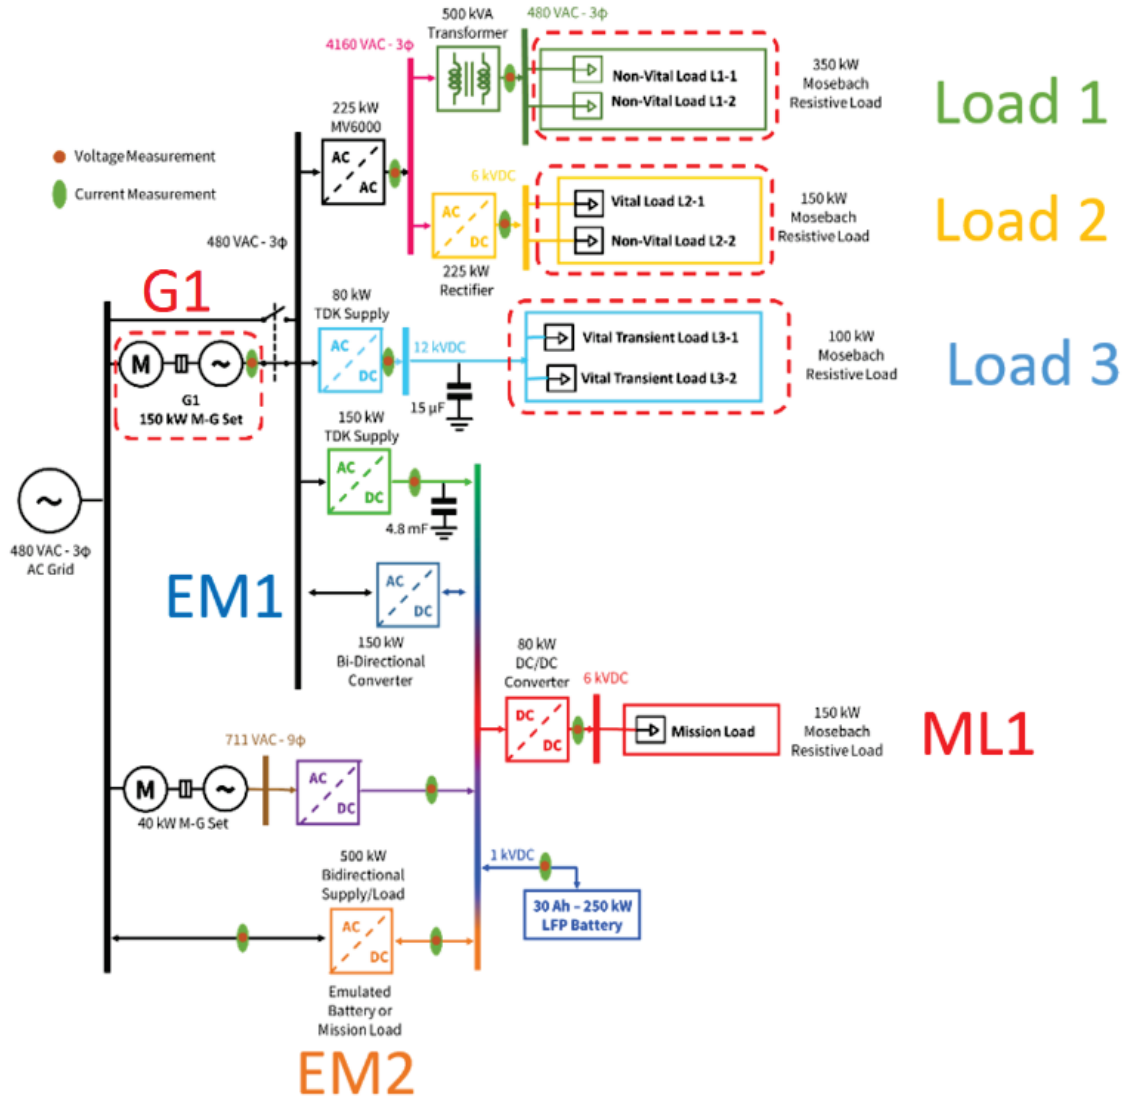


Figure 4. UTA IDEAL testbed one-line diagram. (© 2023 IEEE)

The integration of transient load profiles into Load Group 3 (L3) and the Mission Load (ML) is achieved through the modulation of the respective AC/DC converter outputs that are interfaced with the resistive load. This modulation imparts dynamic transient characteristics to Load Group 3 (L3) and the Mission Load (ML), enhancing the testbed's versatility in simulating real-world scenarios and facilitating comprehensive evaluations. The testbed can be defined or characterized differently for the testing of a variety of controls, but this will be the primary configuration for this work.

BMS TESTBED EVALUATION

At the start of this study, a comprehensive array of options were assessed in formulating the design for a testbed on which BMSs could be studied. The foundational requisites for this system were defined as follows: must be able to study a minimum of 260 ‘cells’ in series, responsiveness registering within hundreds of milliseconds or less, and a galvanic isolation of 1000 V or beyond. These requirements were decided upon for the 1 kV battery configuration based on lithium-ion cells featuring an approximate nominal value of 3.6 volts per cell. Such a battery configuration would be capable of buffering a 1 kVDC bus, as illustrated in Figure 4. A pivotal stipulation involves ensuring swift hardware and software response times measured on the order of milliseconds to ensure successful HIL deployment. This update frequency facilitates a robust modeling environment, empowering the assessment of BMS responses. Delving into the testbed's developmental trajectory, three primary options emerged and are discussed.

Direct Deployment onto Battery

The initial approach under consideration involved the direct study of a BMS installed on an actual battery. This requires the assembly of 264 cells connected in series, which are subsequently interconnected in 26x series modules to constitute the battery configuration. In this setup, the BMS is affixed and subjected to testing within the physical framework of the system. This solution is the most direct of the options and the simplest; nonetheless, this avenue is beset by a multitude of challenges.

Foremost among these challenges is the intrinsic safety risks associated with lithium-ion batteries, alongside other battery variants. Concerns span a spectrum of potential issues, such as chemical leaks, arc flash occurrences during assembly or

maintenance activities, and the risk of thermal runaway leading to the outbreak of a fire. Furthermore, this setup would inevitably impose downtime, obligating the system to undergo cycles of charging or discharging to achieve desired test conditions. Over time, the cells would undergo gradual wear and degradation, impacting the repeatability of test outcomes. While conceptually straightforward and reliant on minimal software intervention, this approach is underscored by a series of drawbacks and constraints. It falls short of the comprehensive capabilities sought in a testbed, compelling research into alternative solutions that surmount these limitations.

Independent Power Supplies Utilizing Analog Following

The second option considered was connecting multiple programmable power supplies in series, each of which emulates a single lithium-ion cell. A schematic of this setup is shown in Figure 5. This involved the interconnection of numerous programmable power supplies arranged in series, with each supply emulating an individual lithium-ion cell. A visual representation of this configuration is depicted in Figure 5.

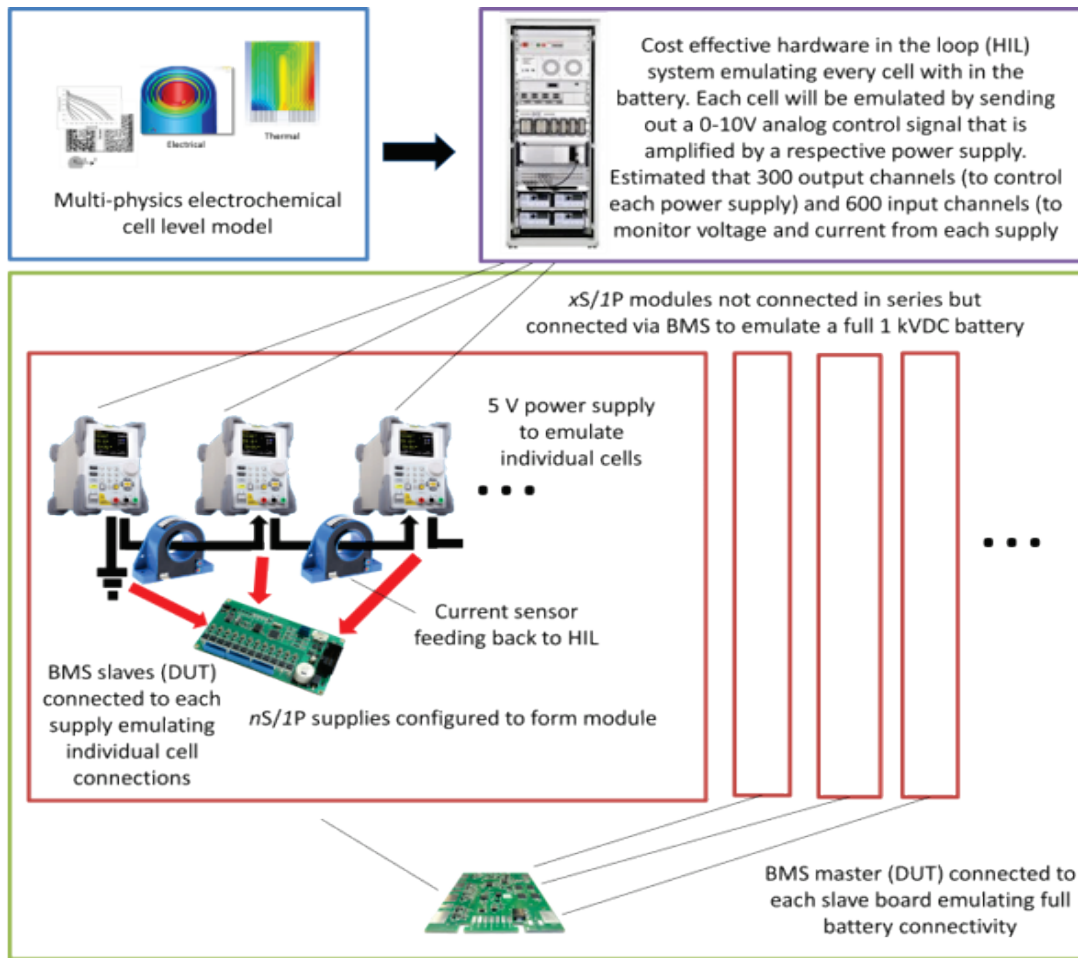


Figure 5. Visual one-line diagram of a testbed layout using multiple individual power supplies to emulate cells. (© 2021 IEEE)

Creating a 1000 V battery in this context necessitates up to 264 power supplies. Furthermore, it's important to acknowledge that power supplies are not conventionally engineered to be interconnected in series in such a configuration, rendering this approach a formidable endeavor.

For a start, each power supply would need to offer galvanic isolation exceeding 1 kVDC, while simultaneously featuring a linear output stage to circumvent switching artifacts. Moreover, the analog following slew rate of each power supply would need to be responsive enough to align with the update rate of a HIL model. The selected HIL platform

itself would need to be endowed with a substantial array of analog input and output channels, thereby creating a notable financial commitment.

To elaborate on the scale of IO, every power supply necessitates a dedicated analog output channel within the HIL platform. Simultaneously, for each cell, a minimum of two analog input channels would be required to monitor output current and voltage, alongside other parameters yet to be identified. Additionally, contingent on the communication capabilities of the power supplies, independent voltage and current sensors would be essential for the monitoring of the 'cell' voltage and current outputs. This information is necessary to facilitate real-time updates of the HIL model based on the power drawn by each channel of the BMS.

Despite its conceivable feasibility, this approach was swiftly deemed not viable due to several reasons such as: the intricate complexity, challenges in sourcing power supplies adhering to the series connection and isolation requirements, the logistical intricacies involved in the setup, and the substantial financial investment it would demand. While potentially workable, this option was swiftly discarded in favor of more viable alternatives.

Commercial Battery Emulators

A third alternative that underwent scrutiny entails the utilization of commercial battery cell emulators, exemplified by offerings from prominent entities such as Chroma Inc., Hioki USA, and Speedgoat Inc. The requirement for batteries in domains spanning electric vehicles to grid energy storage has catalyzed the emergence of these specialized cell emulators, manifesting functionality tailored to this specific application. An overview of these emulators follows.

Chroma 87001

The Chroma 87001 battery emulator is distinguished by its channel number and current capability, boasting a total of sixteen channels within each unit. Each channel encompasses a utilizable voltage span spanning 0 V to 5 V, complemented by a source/sink current capacity extending up to a continuous 5 A. Notably, the datasheet enumerates a substantial galvanic isolation threshold of up to 2000 VDC. The units therefore have an emulation capacity encompassing as many as 480 cells with its isolation. A graphical representation of this arrangement is encapsulated within Figure 6, illustrating the hardware setup in a one-line diagram.

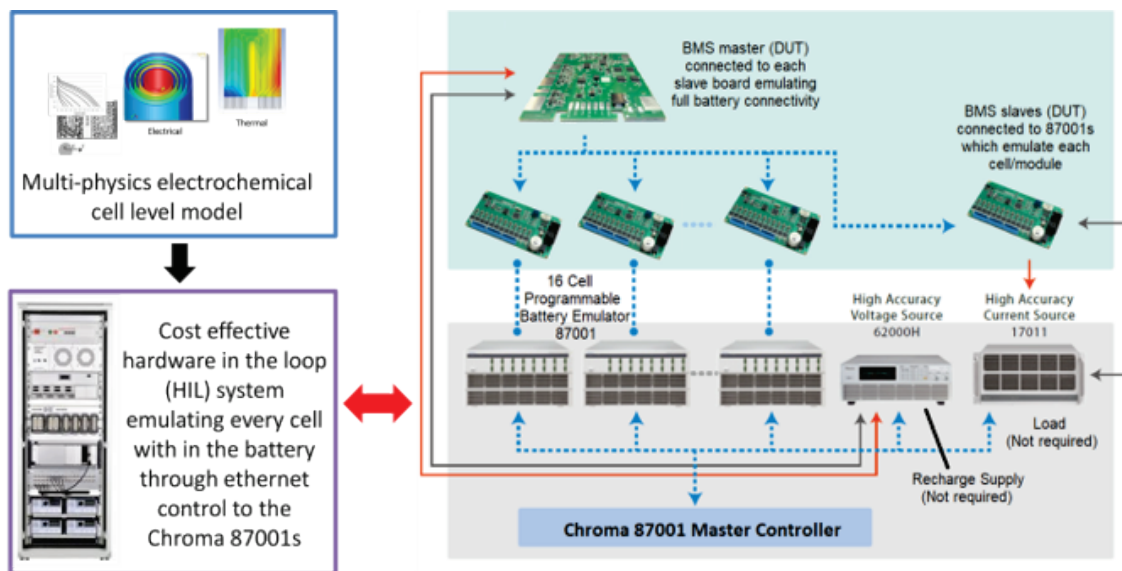


Figure 6. Visual one-line diagram of a testbed layout using multiple Chroma 87001 cell emulators. (© 2021 IEEE)

Chroma introduces a proprietary battery software emulation solution that facilitates the configuration of each channel, with an impressively swift response time of 10 ms per a command. This software empowers users to input parameters that faithfully encapsulate

the attributes of the designated cell. However, it's imperative to discern that this representation doesn't attain the same level of validation as an HIL model would. The 87001, while it doesn't inherently offer instantaneous integration capabilities with an HIL system, does accept remote directives from an HIL unit, leveraging their TCP communication port.

Notably, the interface for remote control can be established through either an Ethernet or Controller Area Network (CAN) communication standard. This communication framework facilitates remote command and control of each individual channel, resulting in response time of approximately 10 ms. The efficacy of this option is underscored by the feasibility of seamless coordination between the proprietary battery emulation unit and the HIL system. Due to the advantages of its practicality, a single unit has been procured for preliminary testing and evaluation.

Hioki SS7081-50

The Hioki SS7081-50 stands as a counterpart to the Chroma 87001, bearing striking similarities in its functionality. Marketed in sets comprising twelve cells, this emulator exhibits the capability to generate outputs ranging from 0 to 5.1 V, accommodating a continuous source/sink current of up to 210 mA. The schematic depiction of a testbed integrating Hioki cell emulators is encapsulated within Figure 7, portraying the hardware configuration in a one-line diagram.

Its galvanic isolation reaches 1400 V, exceeding the necessary minimum 1000 VDC. This attribute culminates in the potential emulation of 260 to 280 cells, provided the output voltage remains conservatively below the 5 V threshold. Akin to its Chroma counterpart, it boasts a proprietary battery emulation software geared for optimal response

time. Notably, the unit also boasts the versatility to accept remote directives through the Ethernet TCP protocol. Hioki asserts that this remote interface offers control over each channel, yielding a commendable response time of approximately 50 ms.

This alternative, currently under contemplation for future investigation, introduces another option to consider. However, it's important to acknowledge that while it holds promise, it was not obtained for demonstration in this work.

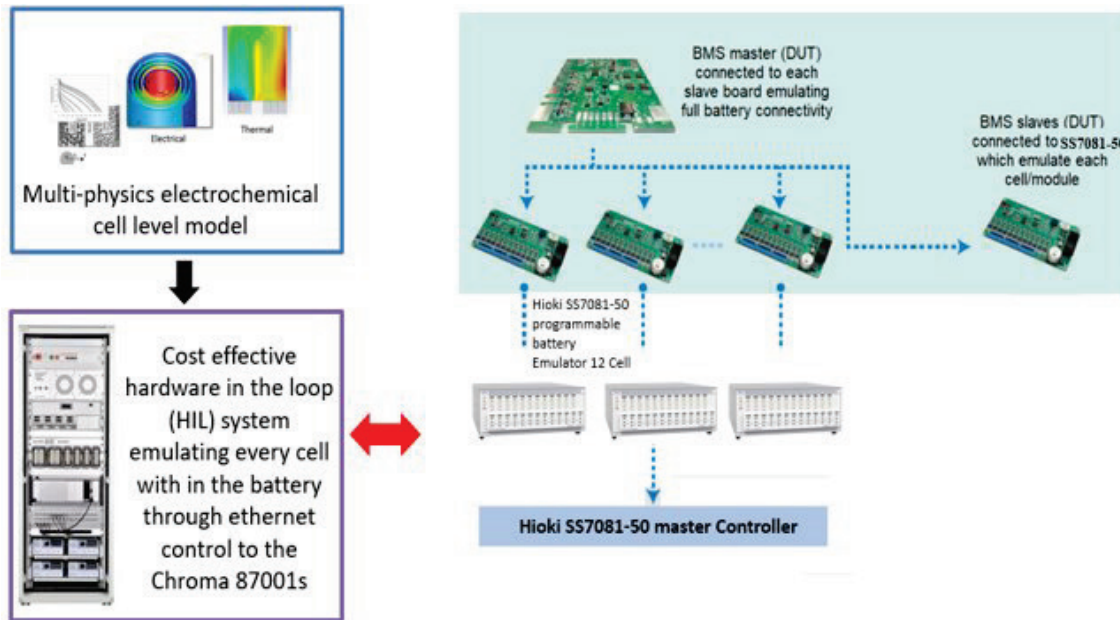


Figure 7. Visual one-line diagram of a testbed layout using multiple Hioki SS7081-50 cell emulators. (© 2021 IEEE)

Speedgoat IO991

Speedgoat introduces an integrated battery emulator solution, distinguishing itself from the offerings of both Chroma and Hioki. The IO991 module from Speedgoat presents an array of six channels, each poised to deliver outputs spanning 0 to 7 V, with the capacity to source 300 mA and sink 100 mA. However, its galvanic isolation threshold of up to 750 V falls slightly short of the target voltage of 1000 VDC. To implement this configuration, IO991 modules are assembled within one or more multi-slot PXI chassis, effectively

stacked in series to achieve the desired configuration. The operational edge of Speedgoat resides in its innate capacity to instantaneously emulate a HIL battery model within the realm of MATLAB/Simulink.

Speedgoat's cell emulation hardware is purposefully engineered to harmonize with its HIL target platform. While Speedgoat's hardware is optimally tuned to interface with its proprietary HIL machine with their block set, it's notable that the IO991 card may not interface with alternative systems. This stands in contrast to the operational modality of Chroma and Hioki systems, which interfaces through the standard Ethernet protocol. As part of the ongoing exploration, three units of the IO991 module were procured for testing during this initiative. Figure 8 illustrates a one-line diagram of the IO991 cards.

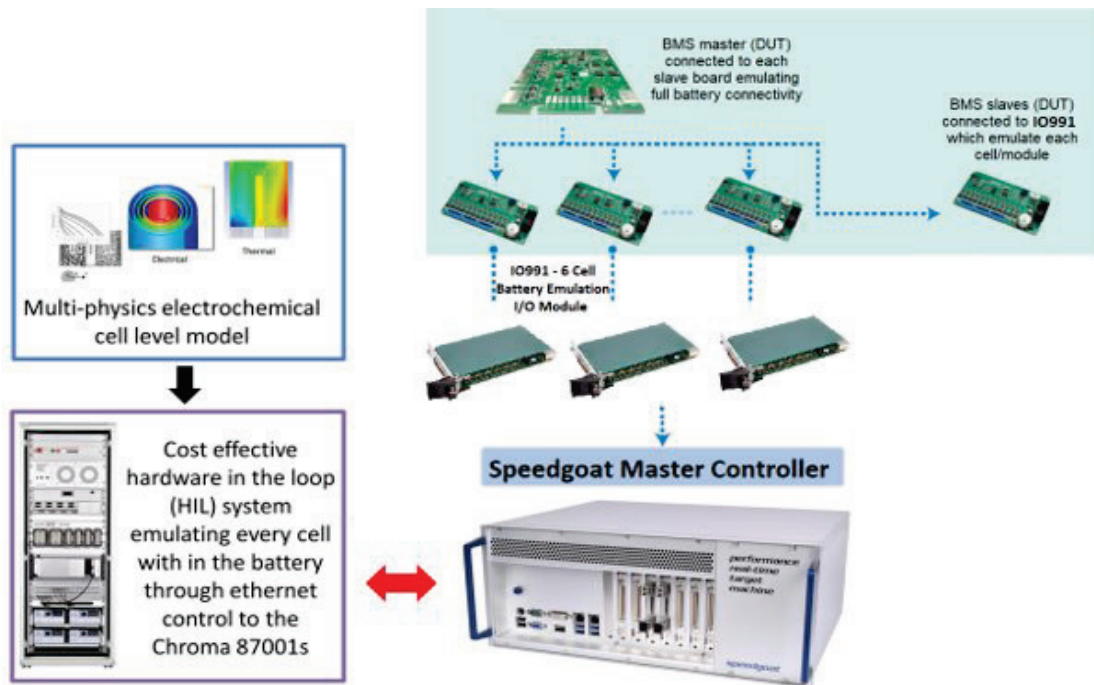


Figure 8. Visual one-line diagram of a testbed layout using multiple Speedgoat IO991 cell emulators. (© 2021 IEEE)

Speedgoat BCE unit

At the beginning of this endeavor, Speedgoat was actively engaged in the development of an enhanced iteration of their battery emulator solution. This iteration features a notable elevation in galvanic isolation, scaling up to 1600 V. Channel specifications for this new model encompass a voltage range of 0 to 5 V, complemented by the capacity to source or sink 5 A per channel.

This emulator is anticipated to communicate via a standard Ethernet connection, akin to the operational modality of the Hioki and Chroma emulators. However, a distinguishing facet lies in the approach to communication and settings management. Unlike manual TCP commands, the forthcoming iteration is expected to encompass a user-friendly, ready-to-use block set. This block set is designed to handle communication protocols and configuration settings, eliminating the need for manual programming of a TCP communication setup within the HIL model.

The design is projected to offer a response rate of around 10 ms for command execution, alongside the incorporation of an integrated current sensing mechanism. The features of this innovative design caused it to emerge as the preferred choice for the emulation of a 264-cell battery configuration. While in development, initial testing was concurrently conducted on both the Chroma and IO991 platforms, culminating in the eventual acquisition of the preferred Speedgoat solution in the second quarter of 2023. Thus, an in-depth exploration of the Speedgoat unit is conducted, while also completing comprehensive testing and assessment of the available alternatives. A one-line diagram of the Speedgoat units is represented in Figure 9

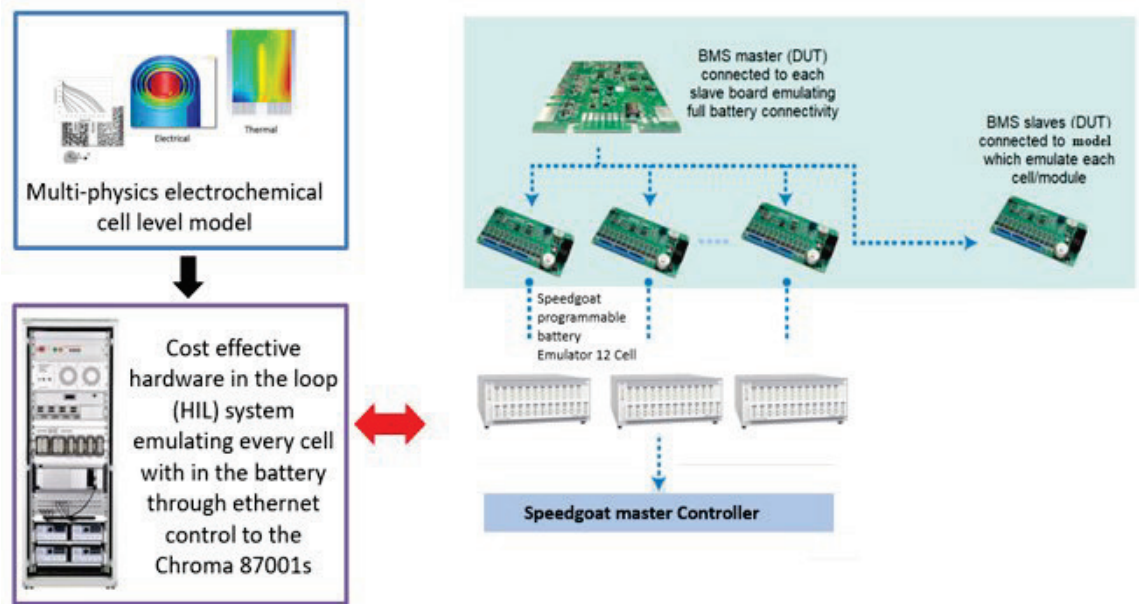


Figure 9. Visual one-line diagram of a testbed layout using multiple Speedgoat BCE units.

Hardware Option Comparison

A comparison between the capabilities of all four commercial solutions is presented in Table 1.

Table 1. Summary of the hardware capabilities offered by Chroma, Hioki, and Speedgoat, respectively.				
Battery Emulator	Chroma 87001	Hioki SS7081-50	Speedgoat IO991	Speedgoat BCE Unit
Channels per Lower-Level Device	16	12	6	12
Voltage range	0 V - 5 V	0 V - 5.1 V	0 V - 7 V	0-5V
Sink Current	5 A	1 A	100 mA	5A
Source Current	5 A	1 A	300 mA	5A
Voltage Accuracy	(0.02%FS) ±1 mV	±0.0150% ±500 μV	±0.2% ±20 mV	±0.01% ±0.5 mV
Voltage Isolation	1000 V ch-ch, 2000 V to gnd	1000 V	750 V	96 V ch-ch 1600 V to gnd

Hardware in the Loop Models

As delineated earlier, the overarching aim is to facilitate the development of multi-physics battery models within the MATLAB/Simulink framework, subsequently enabling their emulation through the PHIL platform. Notably, several models have been crafted and are currently operational within the testbed.

Within the realm of Simulink, two fundamental methodologies exist for formulating battery models. The initial and relatively uncomplicated approach entails the utilization of pre-installed Simulink models. These models present a user-friendly interface, wherein users can conveniently input customizable settings that are contingent upon the specific battery chemistry under scrutiny. Moreover, certain models within this category even provide provisions for calculating cell temperature and State of Health (SoH). This approach, while simplified, effectively streamlines the process of demonstrating how the PHIL model interfaces with the emulator and the responsiveness of the BMS to rapid variations in State of Charge (SoC), orchestrated by the emulator.

Conversely, the second approach empowers users to forge their own mathematical or empirically derived battery models. This technique affords a higher degree of autonomy to users, enabling them to tailor the battery model to encapsulate the desired level of granularity and fidelity. This more intricate approach, the focal point of our current discussion, underscores the significance of PHIL capability within the testbed's framework.

While the initial selection has favored Speedgoat's hardware for battery emulation, it's important to acknowledge the continued viability of the Chroma and Hioki alternatives. Both Chroma and Hioki hardware can interface with any HIL target machine with a TCP interface. Consequently, battery models are being concurrently developed and investigated utilizing both Speedgoat and OPAL-RT HIL systems, leveraging the respective Chroma hardware option.

Remarkably, the Speedgoat Battery Cell Emulator (BCE) unit, the improved revision of their battery emulator, exhibits the potential of interfacing with other target machines. However, it's notable that the IO991 module within the Speedgoat framework

lacks this interfacing capability. Notably, the Speedgoat target machine establishes a direct connection with its IO991 hardware through the Simulink Real-Time environment, exclusively employing a PXI chassis for this purpose.

To achieve model deployment from Simulink to Speedgoat's hardware, a MATLAB add-on, Simulink Real-Time, must be acquired and installed. This supplementary toolset empowers Simulink to transition the model into Real-Time operational mode. In addition, Speedgoat's blockset must be downloaded from their website and installed into MATLAB libraries. Upon completion of this, the IO991 modules from Speedgoat can be configured to interface with individual cell models, effecting a successful emulation process.

OPAL-RT employs a similar approach to Speedgoat's Simulink Real-Time environment, operating through its proprietary platform known as RT-Lab. This platform serves as the conduit for deploying Simulink models onto Opal's Real-Time (RT) computer, thereby establishing an interaction with various emulation hardware components. Much akin to the Simulink Real-Time framework, RT-Lab transitions Simulink models into a Real-Time operational environment, effectively priming them for integration with a diverse array of emulation hardware elements.

However, it's imperative to underscore that the operational nuances of RT-Lab and Simulink Real-Time diverge significantly. Within the RT-Lab paradigm, the Simulink model requires partitioning into distinct subsystems, prominently encompassing the Subsystem Master (SM), Subsystem Slave (SS), and Subsystem Console (SC). This modular architecture facilitates the distribution of computational workloads. SM and SS systems allocate individual CPU cores of the target machine to effectively process

calculations in their subsystems. Concurrently, the SC facilitates data collection and user control operations, fostering a link to the desktop user interface.

It's worth noting that while RT-Lab empowers users with an enhanced degree of configurability for model compilation and presents a spectrum of options for emulator deployment, this heightened flexibility is accompanied by an escalation in the intricacies of the model development process. Figure 10 encapsulates a battery model deployed to the computer using Simulink.

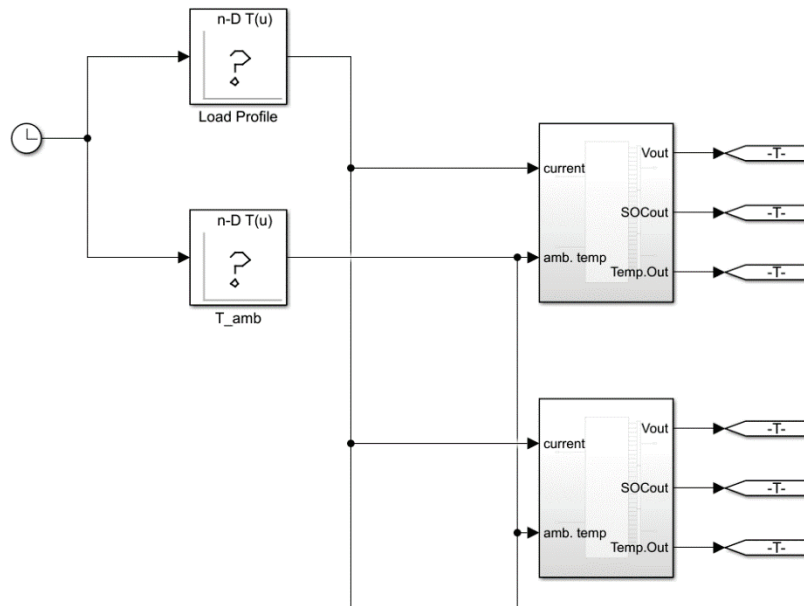


Figure 10. Model showing the 10 cell modules connected to a load profile and temperature profile. These modules can then be further connected in series or parallel to form a battery. (© 2021 IEEE)

The successful deployment of this model has been realized through both the Speedgoat and OPAL-RT HIL target machines, marking a significant achievement. However, the forthcoming report focuses its work towards the model's implementation on the Speedgoat HIL machine, affording it prominence with its deployment on the Speedgoat

BCE units. The Opal deployment will have a focus with the cosimulation controls later in the work.

The underlying architecture of this battery model revolves around the utilization of multiple instances of a Simulink cell model that is interconnected in a series configuration. This systematic arrangement materializes as 10S modules and an aggregation of 26 such modules create the complete 1000 V battery. In a strategic augmentation, a 27th module, encompassing 4S, is strategically incorporated to achieve the desired total of 264 cells.

The forthcoming report will delve into the model's emulation processes, traversing a variety of emulators and BMSs. This includes a subset of IO991 modules, a solitary Chroma 87001 unit, and the comprehensive testbed featuring up to 22 Speedgoat BCE units. The results chapter will chronicle and present these emulation scenarios, culminating in an analysis of the BMSs.

Lithium Cell Discharge

Figure 11 and Figure 12 serve as depictions of a generic battery simulation portraying the interplay of the 264 cells in series, constituting a 1000 V lithium battery model. In the initial plot, the discharge process unfolds as each of the 264 individual cells undergoes discharge at a rate of 1C. The subsequent plot visually encapsulates the effect of SoC on these cells.

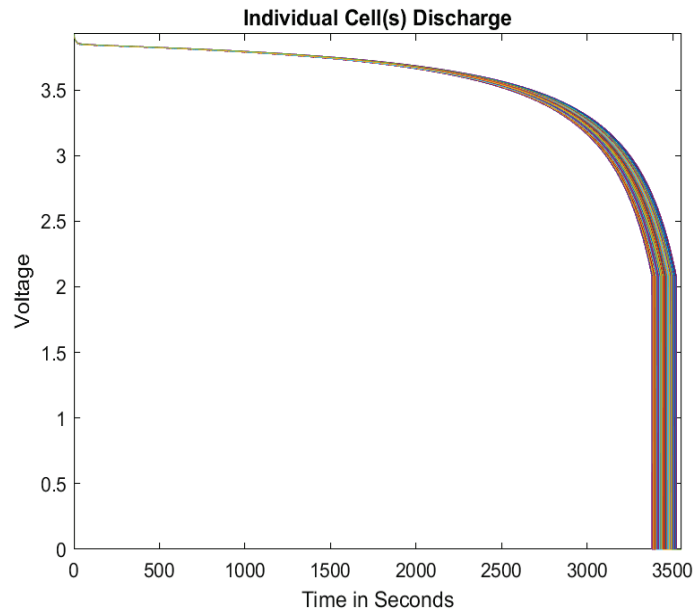


Figure 11. Lithion ion cells with different starting SoCs ($93\pm 2\%$) discharging at 1C rate. This slight change in SoC can cause dramatic changes in the voltage over time. (© 2021 IEEE)

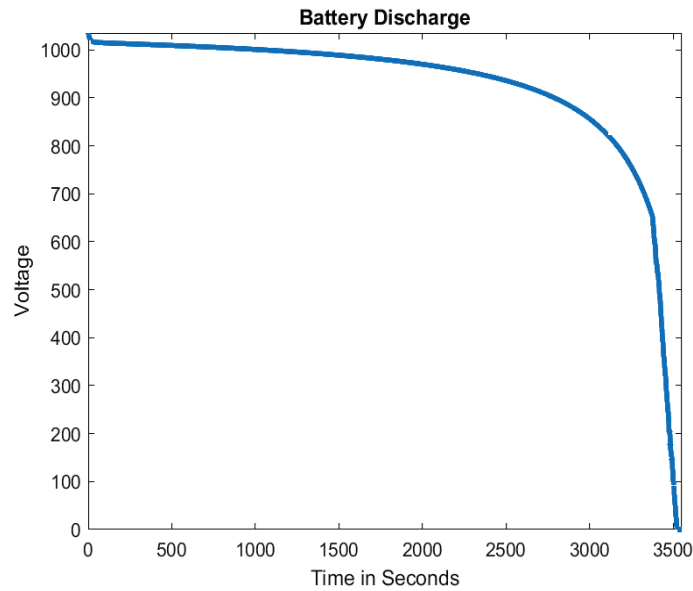


Figure 12. Lithion ion battery shows the sum 264 cell string discharging at 1 C rate with cells at different starting SoCs ($93\pm 2\%$). (© 2021 IEEE)

It is paramount to acknowledge that the model, in its current instantiation, does not account for all details, such as potential parasitic losses due to series interconnections

within the battery. Rather, it serves the purpose of illustrating a pertinent facet: the difference of a $\pm 2\%$ state of charge (SoC) alteration does not unequivocally correlate with a corresponding $\pm 2\%$ variation in voltage. This perceptible deviation inherently poses a challenge for the BMS, elevating the intricacy of its goal to accurately estimate cell charge during routine operational cycles. This highlights the multifaceted nature of this endeavor and the critical role of the BMS in navigating this intricate landscape.

HIL VL30AF Deployment

The battery model employed in this testing endeavor is a Multiphysics representation derived from the Saft VL30AF battery cell. This validated and verified model has been deployed within the Simulink Real-Time framework for Speedgoat deployment, as well as within the Rt-Lab environment for Opal deployment.

It is important to note that due to the proprietary nature of the model, the extent of information that can be shared regarding it is limited. Nevertheless, Simulink offers a range of simplified models, exemplified in Figure 10, designed for convenient deployment. Additionally, users have the option to pursue the development of more intricate, cell-specific models, thereby affording a higher degree of detail and accuracy as used here.

Within an actual battery, effecting a rapid and safe alteration in a cell's State of Charge (SoC) remains an unattainable feat. However, with the utilization of cell emulators, the capability to swiftly modify a cell's SoC for the purpose of assessing the response of a BMS at voltage thresholds and equilibrium levels has become a tangible reality. This approach accelerates the evaluation process by obviating the need for the time-consuming charging or discharging procedures requisite in a physical battery setup.

This achievement is facilitated through the integration of a MATLAB application enabling real-time adjustment of each individual cell's SoC. This application empowers users to manipulate the SoC of individual cells within the battery during runtime, eliminating the necessity for model recompilation and upholding the model's optimal performance. A visual representation of this application, presenting the SoC status of each cell, is depicted in Figure 13. This application expedites assessment protocols and enhances the efficiency of BMS evaluation in simulated scenarios.

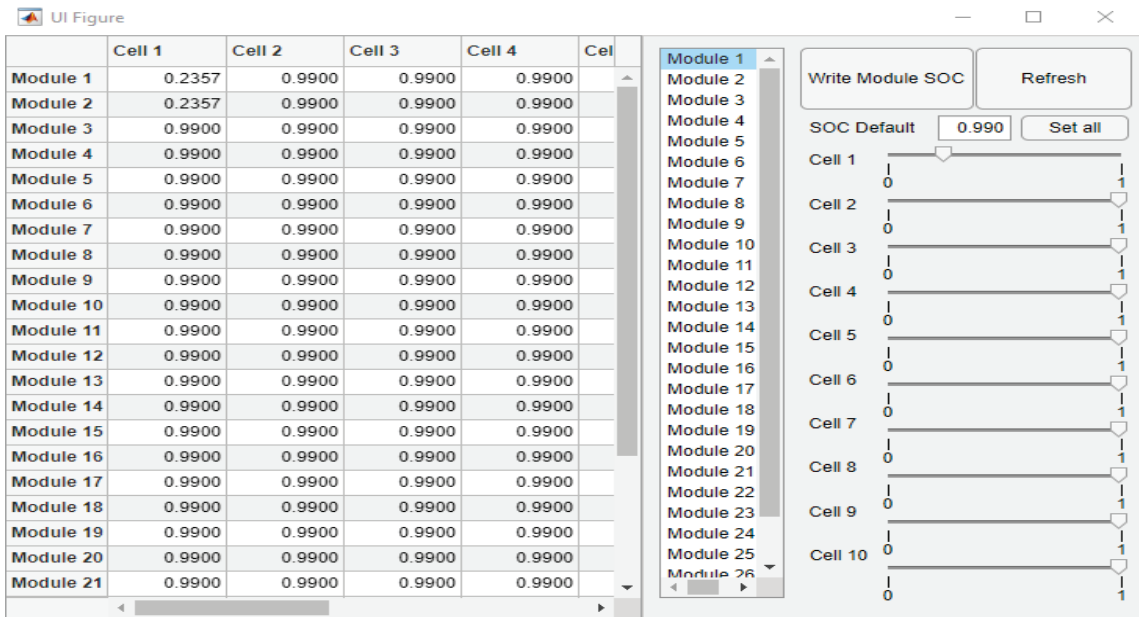


Figure 13. Custom MATLAB application used to change the SoC of each individual cells in real time. (© 2021 IEEE)

IDEAL TESTBED DEPLOYMENTS

This section encompasses the deployment of various control strategies onto the testbed. The initial endeavor involves a cosimulation initiative in collaboration with FSU and USC, where HIL deployment is integrated with the IDEAL testbed. Demonstrating PHIL deployment on IDEAL after verifying the BMS's operational effectiveness will prepare the groundwork for a PHIL deployment using the real 1 kV battery. The subsequent phases of this section encompass the implementation of Advanced Load Shedding (ALS) and Predictive High Ramp Rate (PHRR) controllers, both developed by CU. These controllers are incorporated onto the IDEAL testbed to evaluate their performance and efficacy. Specifically, these experiments are conducted using the 1kV battery in conjunction with the K2 BMS after its testing with the battery emulators.

Co-Simulation with FSU and USC

As previously explained in the introduction, the complexity of shipboard power systems, akin to the utility grid, have been steadily increasing alongside the rising demand for enhanced monitoring and control capabilities [15]. Employing a model-based approach for the prototyping and engineering of such power systems offers distinct advantages, as it mitigates the logistical and financial challenges often associated with hardware experimentation. Successful implementation of a model-based engineering strategy hinges on the utilization of verified and validated (V&V) models [16]. Numerous research groups are actively developing model sets tailored for engineering next-generation shipboard power systems. The advent of modern HIL simulators enables real-time simulation of these models, with time steps fine enough to attain the intended fidelity crucial for effective prototyping [17]. However, the cost and computational demands associated with HIL

equipment can pose significant barriers. These obstacles can be surmounted by leveraging distributed HIL resources across different laboratories, even when they are geographically separated. This can be achieved through cosimulation, employing open-source software tools for simulator communication. This will be the focus of this effort. Notably, the laboratories involved in this study are geographically distant, with FSU approximately 503 km from USC, and UTA spanning a distance of 1237 km from FSU. This work entails conducting a collaborative co-simulation across these laboratories, focusing on a microgrid system depicted in Figure 14, which serves as a segment of a DC shipboard power system (SPS).

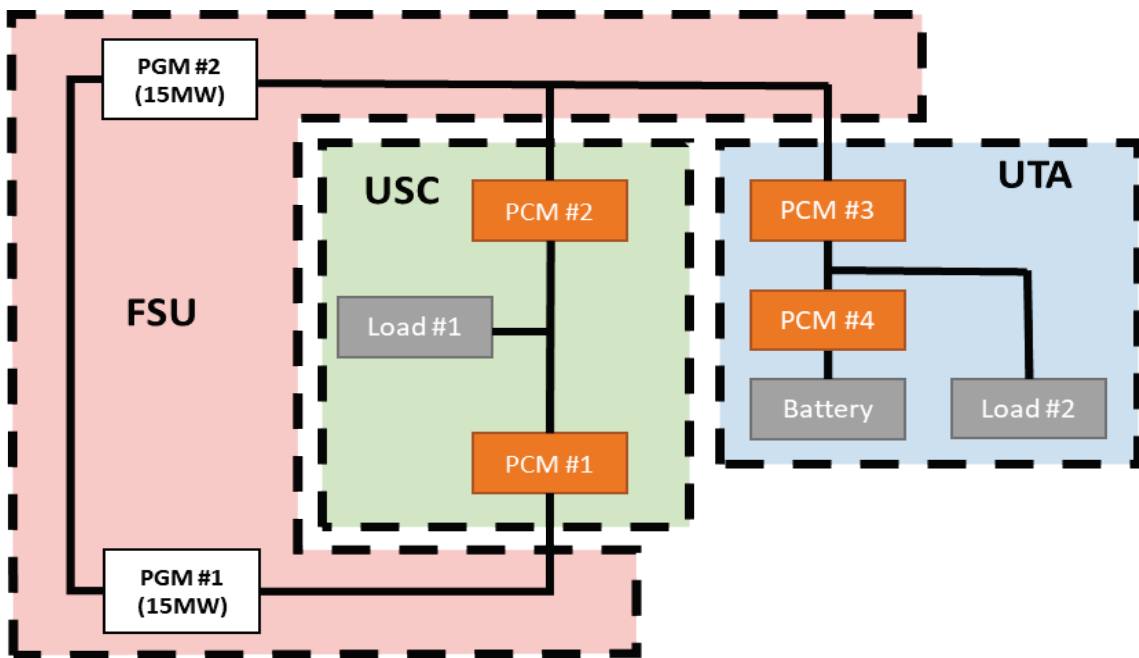


Figure 14. Setup for co-simulation with USC, FSU and UTA(© 2022 IEEE)

The focal point of this study is a self-contained DC microgrid, characterized by the presence of distributed power generators that supply power to designated load centers via a DC distribution network. The microgrid configuration incorporates a collection of power converter modules, facilitating interconnections between the generators, loads, energy

storage systems, and the overall energy flow. The entire microgrid system is compartmentalized into three distinct subsystems, each subject to real-time simulation and control within dedicated HIL platforms across different laboratories. These HIL simulators are interlinked over the Internet using an open source cosimulation framework and networking tools. This collaborative setup enables the unified simulation of the microgrid system with relative accuracy and time fidelity. The subsequent sections provide an exploration of the networked cosimulation framework, and analysis of results from the conducted experiments. The work presented here will focus on UTA's contribution; however, more information over the work can be found in in an OSMSES paper [18]. Results of the cosimulation will be discussed in the Results chapter.

Co-Simulation Framework

The co-simulation framework adopted for this research is built upon VILLAS node [19], an open-source network gateway developed in C/C++ that serves as a conduit for real-time co-simulation interface data. This versatile platform facilitates communication between disparate simulators or sources, utilizing a wide range of communication protocols. Moreover, it offers valuable features such as data communication statistics and network emulation, crucial for effectively managing co-simulations spanning geographically distant computers [20].

Within each laboratory, a dedicated VILLAS node server was deployed on Internet-connected x86-based hardware, forming an interconnected local network with the respective HIL simulator. This architecture enabled communication between the HIL simulators of different laboratories through their respective VILLAS node servers. Configuration files were employed to tailor the behavior of the VILLAS node servers,

enabling them to efficiently route communication data between the HIL simulators. A diagram showing the local connection between the systems and computers is in Figure 15

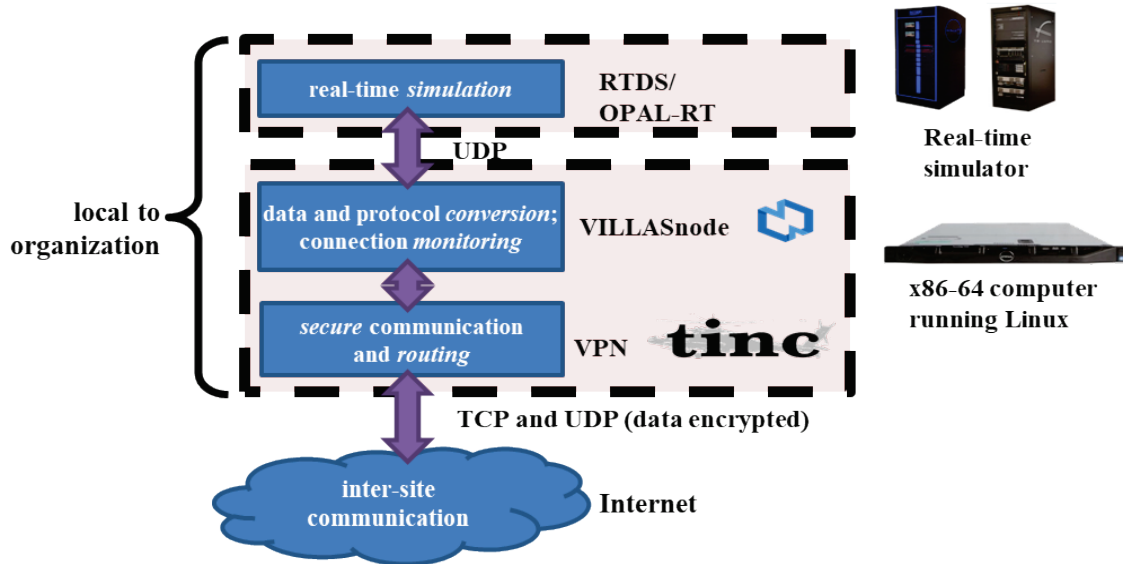


Figure 15. One line Diagram of connections between real time simulators, desktop computers, and the Internet for cosimulation(© 2022 IEEE)

To facilitate communication, the VILLAS node servers and the HIL simulators utilized the User Datagram Protocol (UDP). The robustness and security of communication across laboratories were fortified by leveraging the Virtual Private Network (VPN) capabilities of the open-source TINC-VPN daemon [21] This VPN connection established a secure and reliable communication channel over the Internet, ensuring the confidentiality and integrity of data exchange between the laboratories. A visual diagram of this connection is shown in Figure 16

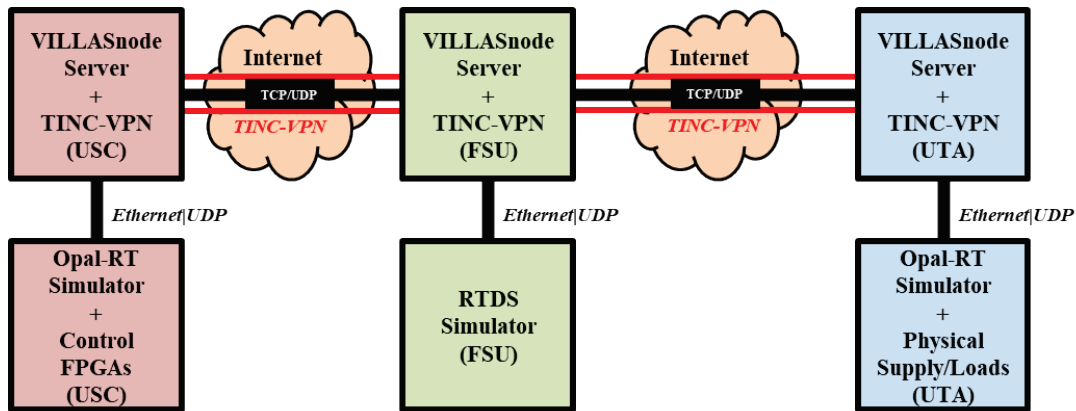


Figure 16. Connection between simulators and universities using Villas node and Tinc VPN.

Latency in network communication between FSU and USC was assessed through network pinging, yielding an estimated latency of approximately 16 ms. Comparatively, the latency between UTA and FSU was measured at approximately 30 ms. These latency figures provide valuable insights into the responsiveness and synchronicity of the distributed co-simulation framework.

UTA Site Zonal HIL

Illustrated in Figure 14, UTA's designated zone undertakes the simulation of a transient load, specifically denoted as Load #2, which finds its buffering support from an energy storage unit – in this instance, a lithium iron phosphate (LFP) battery boasting a capacity of 30 Ah and an operational voltage of 1000 V. PCM #3, functioning as a pivotal power control module, is entrusted with the task of governing power dynamics derived from the shared 12 kVDC bus, which serves as a common link for FSU, USC, and UTA. This bus's energy supply emanates from PGM #2.

The strategic operation of PCM #3 is executed in the role of a controlled current source, adroitly managed through reference set points originating from UTA. Meanwhile, the regulation of the LFP battery is governed through PCM #4. Notably, PCM #3 and PCM

#4 collaborate to maintain integrity of the 12 kVDC bus against potential disruptive transients, such as the incursion of a square pulse. This collaborative approach is executed via the implementation of ramp-rate control mechanisms, ensuring controlled power contributions to the 12 kVDC bus.

Within the realm of UTA's simulation framework, the operational domain centers around the OPAL-RT platform. The present model configuration capitalizes on three real-time cores, executing at a fixed time interval of 2 milliseconds. This allocates the first core to the execution of the system in XPS mode, thereby managing the hardware synchronization inside of Opal. This mode, in turn, enables the analog inputs and outputs OP7000 for use.

The remaining two cores serve model execution purposes within this network. The second core is dedicated to the host subsystem, encompassing critical components such as PGMs (Power Generation Modules), simulated loads, analog inputs, analog outputs, and the UDP communication linkage to the Villas Node computer. The third core engages in the administration of the client subsystem, housing the Opal version of the Multiphysics VL30AF battery model discussed earlier in the BMS Testbed Evaluation

Lastly, in this framework is the console subsystem, facilitating local manual input control and runtime data visualization. It is important to underscore that this subsystem operates asynchronously, enabling interaction with the real-time OPAL-RT machine from the user desktop.

To create the cosimulation setup, data is exchanged between FSU and UTA. Specifically, FSU contributes load enable signals and generator voltage to the UTA model from their PGM model, while UTA reciprocates by providing feedback to FSU with data

pertaining to the generator's current draw and the State of Charge (SoC) values of the battery.

A segment of the UTA simulation is being replicated in hardware through the utilization of a power hardware in the loop (PHIL) arrangement, visually depicted in Figure 17. Instead of the existent battery at UTA, a programmable power supply engineered by TDK Lambda, boasting specifications of 1200 V, 125 A, and 150 kW, will take on the role of emulating the series-connected battery and PCM #4.

It is worth noting that while the battery can deliver up to 325 kW, the power supply has limitations that render it unsuitable for emulating the battery's complete usable range. However, this discrepancy does not hinder the successful execution of the experiments undertaken in this context. The emulation of the load is achieved through a bi-directional programmable load/supply crafted by Chroma, designed with capacities of 1200 V, 700 A, and 500 kW.

To ensure precise control, the reference points for both the TDK supply and the Chroma load are dynamically managed via 0 – 10 V analog reference signals delivered directly by the OPAL-RT platform. In tandem with this setup, a hall effect current sensor, denoted as CT, featuring a 50 kHz bandwidth, is positioned at the output of the TDK power supply. The sensor's measurements are harnessed and fed back into an analog input interface on the OPAL-RT system, thus facilitating real-time updates to the battery's state of charge (SOC).

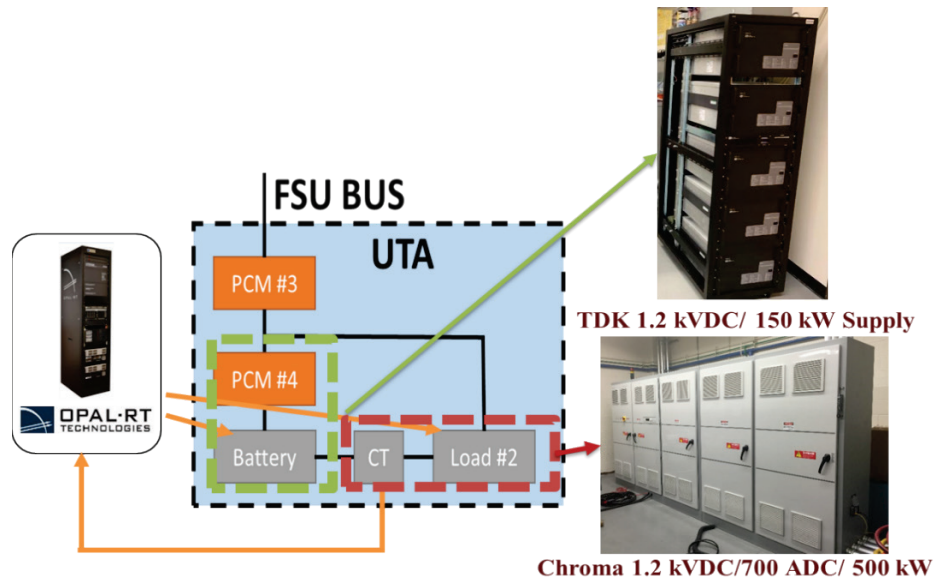


Figure 17. Hardware configuration used by UTA IDEAL for PHIL deployment in co-simulation. (© 2022 IEEE)

Since all components in the UTA setup are being emulated, this is not truly demonstrating opportunities that a PHIL execution affords. For example, since UTA has a battery, PHIL emulation of the PCM and load using the physical battery would enable several electrical and thermal properties of the battery to be studied. Emulation of the battery into a physical PCM and emulated load would allow the PCM to be studied. Finally, emulation of the battery and PCM into a physical load would allow the load to be studied. For the sake of this work, all three devices are being emulated to show viability of a remote co-simulation deploying hardware emulated components as well. The PHRR controller discussed later in the work will make use of the real 1000V battery.

In the UTA configuration, where all elements are being replicated through emulation, the full potential of power hardware in the loop (PHIL) execution may not be fully demonstrated. To illustrate, in the presence of UTA's physical battery, employing PHIL emulation to replicate the PCM and load using the tangible battery could enable in-

depth exploration of various electrical and thermal attributes of the battery. Similarly, introducing the emulation of the battery alongside a physical PCM and emulated load would facilitate an in-depth analysis of the PCM's behavior.

Furthermore, the integration of battery and PCM emulation with a physical load could provide a comprehensive understanding of load characteristics. However, within the scope of this work, the decision has been made to emulate all three devices. This choice is driven by the objective of showcasing the feasibility of a remote co-simulation utilizing hardware-emulated components. It is important to note that the PHRR controller, which will be discussed in subsequent sections of this work, will actively involve the authentic 1000V battery, thereby ensuring a well-rounded and comprehensive assessment of controls utilizing a BMS and real 1 kV battery.

ALS and PHRR Control Methodology

The control of the IDEAL testbed is effectively managed by employing real-time sensors that establish a connection with National Instruments PXI and CDAQ hardware components. Operating in tandem, the PXI controller digitizes the data captured by the sensors, transmitting it to the virtual instrument (VI) responsible for supervising the CDAQs associated with the hardware. Within this framework, the hardware system can be engaged manually via the VI or in an autonomous manner through the integration of control algorithms embedded within it.

For this particular project, Clarkson University (CU) has undertaken the development of Advanced Load Shedding (ALS) and Predictive High Ramp Rate (PHRR) controllers within the Simulink Desktop Real-Time (SLDRT) environment. These controllers establish an interface with the VI, enabling them to impart instructions on how

the hardware should be directed and operated. In the upcoming sections, a comprehensive exploration of these algorithms will be provided, delving into their intricacies and functionalities.

ALS Control Methodology

For implementing the ALS controls, the M-G set G1 operates at its maximum capacity of 150 kW under a 480V rating, and the load groups (1, 2, and 3) consist of six distinct loads that can be independently operated. Load Group 2 (L2) and Load Group 3 (L3) each include two 50 kW internal steps, designated as L2-1, L2-2, L3-1, and L3-2. Each load however achieves these steps by different means. For Load Group 1 (L1), it is represented as two separate loads in software and it has many internal contactors which give it a 1kW step resolution, while L2 has two distinct steps within the unit thus making it similar to L1, but more simplistic. L3 although has two distinct steps, is regulated by the current provided by the 12 kV TDK to emulate two steps. The Mission Load (ML) is a single step transient load buffered by the lithium-ion battery. The Advance Load Shedding (ALS) controller is designed to activate when G1's rated power is surpassed, with the primary goal of sustaining critical loads while optimizing overall system functionality. The controller assigns the highest priority to L2-1 and the mission load, ensuring that they are never shed. Manual control of load states is facilitated through the VI, which communicates load states and power consumption to the ALS controller via UDP messages.



Figure 18. Main VI used to command supplies and loads in the IDEAL testbed. (© 2023 IEEE)

CU's controller has been effectively tested and integrated into the IDEAL testbed, as illustrated in Figure 19. This controller is structured with three distinct layers: external control, signal communication, and power system hardware. The external control layer is designed using SLDRT, while the signal communication layer employs the UDP protocol to facilitate data exchange between the controller and the LabVIEW hardware control system. The IDEAL testbed is connected to the LabVIEW controller, which receives load shed commands from the Advance Load Shedding (ALS) controller and transmits power supply information from IDEAL. The latency between the two control systems has been measured at approximately 100-200 ms.

Figure 20 provides an insight into the Advance Load Shedding (ALS) controller. The supervisory system is responsible for regularly refreshing IDEAL's load requirements and system structure, encompassing the operational state of the generation and energy storage elements. Embedded within the mission database is the load prioritization scheme,

driven by the criticality of individual loads. Within the ALS controller, an optimization algorithm is executed to formulate load shedding directives for each load when confronted with an over-power scenario. This iterative procedure occurs at a frequency of T_{con} during real-time operation, where T_{con} signifies the control sampling interval.

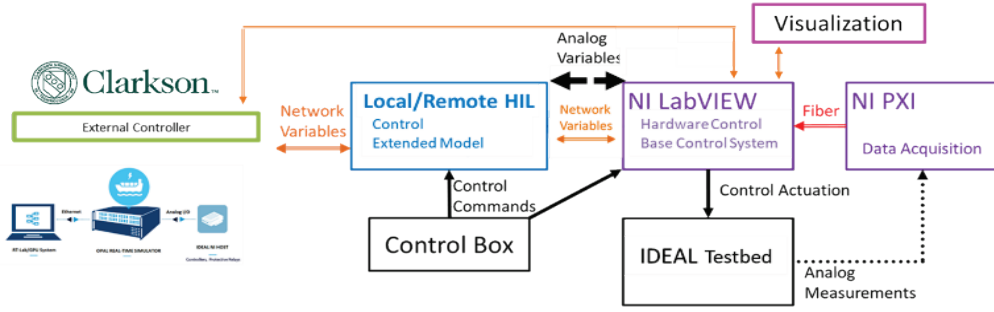


Figure 19. Integration the UTA IDEAL hardware and LabVIEW control system with CU's SLDRT controller. (© 2023 IEEE)

The weighting factors (ω_i) serve as indicators of load prioritization as dictated by the supervisory system. These factors reflect the proportion of load power to be allocated for each specific load. To accurately account for the significance of loads and their respective rated power within the context of ALS challenges, the concept of load operability (O) is introduced. This factor quantifies the extent to which a load should be powered, as outlined in equation (1).

$$O = \frac{\int_{t_0}^{t_f} \sum_{i=1}^{n_L} \hat{\omega}_i o_i^t dt}{\int_{t_0}^{t_f} \sum_{i=1}^{n_L} \hat{\omega}_i o_i^{*t} dt} \quad (1)$$

In the given equation, where n_L denotes the total number of loads, $P_i^{L, rated}$ signifies the rated power of load i , and $P_i^{L, t}$ represents the power demand of load i at time t . Additionally, ω_i denotes the normalized weight value of load i as defined in equation (2). The variables O_i^t and O_i^{*t} correspond to the operational status of load i at time t and the

commanded operational status. Furthermore, t_0 signifies the initiation time of the operation, while t_f indicates the conclusion time of the operation.

$$\hat{\omega}_i = \frac{\omega_i P_i^{L,t}}{P_i^{L,rated}}, \forall i \in n_L \quad (2)$$

The ALS problem's formulation hinges on load weights (ω_i) and the load demand, serving as determinants for shedding load. The operational status (O_i) signifies the operability of load i and is the optimal variable within the scope of the optimization problem. Specifically, when $O_i=0$, it signifies that load i will be fully shed; conversely, $O_i=1$ denotes that load i will be completely served.

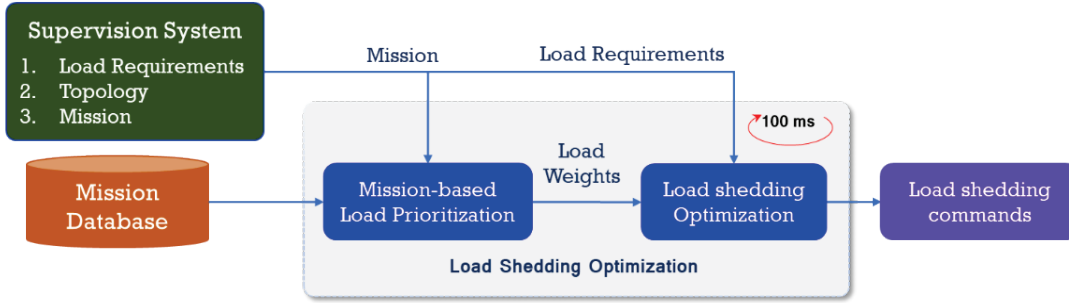


Figure 20. Block diagram of the load shedding control. (© 2023 IEEE)

The core aim of the ALS is to optimize overall load operability, as exemplified by the ALS objective function outlined in equation (3.1). Drawing inspiration from Clarkson's prior research [22], the controller's formulation has been adapted and enhanced with specific modifications. These adaptations account for ramp-rate dynamics and the discrete attributes of load devices, resulting in a more refined and efficient load shedding mechanism.

$$\max_{o_i} \frac{\sum_{i=1}^{n_L} \hat{\omega}_i o_i}{\sum_{i=1}^{n_L} \hat{\omega}_i o_i^*} \quad (3.1)$$

Subject to

$$\sum_{i=1}^{n_L} P_i^{L*} o_i \leq (1 - \beta) \sum_{g=1}^{n_G} P_g^G \quad (3.2)$$

$$r_L^{\min} \leq \Delta P_i^L \leq r_L^{\max} \quad \forall i \in n_L, \quad (3.3)$$

$$\Delta P_i^L = P_i^{L*} o_i - P_i^{Lp} o_i^p \quad \forall i \in n_L, \quad (3.4)$$

$$0 \leq o_i \leq 1 \quad \forall i \in n_L. \quad (3.5)$$

$$o_i \in \begin{cases} Z & \text{if load } i \text{ is the step load} \\ [0,1] & \text{otherwise.} \end{cases} \quad (3.6)$$

In equations (3), where n_G denotes the count of generators, β represents the reserved power percentage, P_i^{L*} signifies the load reference for load i , P_i^{Lp} corresponds to load data from the previous interval state, and r_L^{\min} as well as r_L^{\max} stand for the minimum and maximum ramp rates of loads. Here, Z is a set consisting of $\{0: \Delta O_i: 1\}$ where $\Delta O = 1/n$; and n belongs to the positive integers, Z^+ . The inclusion of constraint (3.2) enforces the generation-demand balance, ensuring that the cumulative load command remains lower than the combined available generation. Additionally, constraint (3.3) reflects the imposition of ramp-rate limitations on loads. The operational status is subject to constraint (3.5), adhering to operational boundaries. Lastly, for loads characterized by discrete changes, their operational status is governed by the constraint inherent to the discrete function (3.6).

PHRR Control Methodology

The core aim of the Predictive High Ramp Rate (PHRR) controller is to guarantee the uninterrupted operation of the mission load in scenarios where a deficit in generation

power arises due to generator ramp-rate restrictions. Much like the ALS controller, the PHRR controller is constructed using the Simulink Desktop Real-Time environment. Both the ALS and PHRR controllers are hosted on distinct personal computers, linked through identical connections to the IDEAL LabVIEW controller. The long-term goal is to consolidate these two controllers into a unified implementation, thereby offering both functionalities within a single instance of SLDRT.

The mission-based PHRR control system is engineered to synergize with the IDEAL hardware configuration, as depicted in Figure 21. The supervisory module stands as an overseer, consistently monitoring and revising critical system parameters including high ramp-rate loads, the status of generators and energy storage, and the load requisites stemming from IDEAL. Vital constants such as energy, power, and ramp-rate limitations, encapsulated within the database block, provide the foundational input that outlines the operational attributes of the generator and energy storage components. These are relayed to the ramp-rate support and Energy Storage Management (ESM) components, both integral to the PHRR controller. Applying system-wide objectives as guiding principles, the PHRR controller proceeds to calculate the ongoing reference signal for the AC/DC power converter. This converter, in turn, regulates the power distribution to the ML as well as the LFP battery.

The PHRR controller establishes communication with the LabVIEW control panel shown in Figure 18 through the transmission of UDP packets, mirroring the communication approach adopted by the ALS controller. This communication protocol is reiterated at a 10ms interval during real-time operation, ensuring the capture of high ramp-rate load dynamics with precision and fidelity.

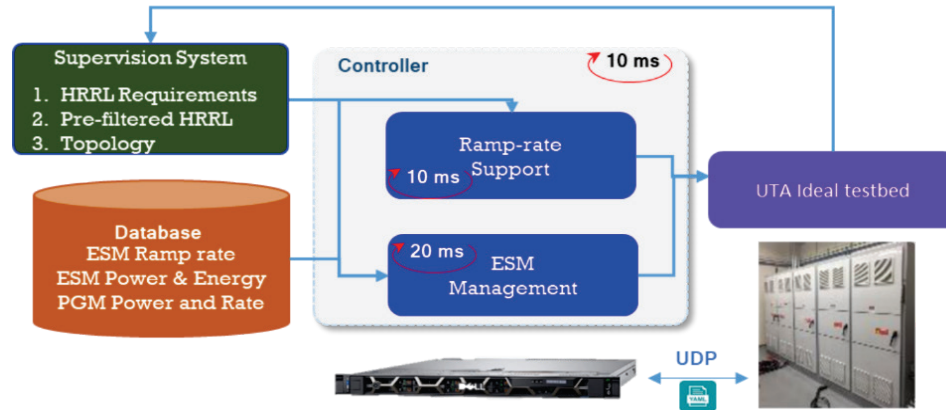


Figure 21. Predictive high ramp rate controller block diagram. (© 2023 IEEE)

The PHRR controller employs a hybrid predictive technique derived from Tuyen Vu's previous work [23], integrating both heuristic and predictive approaches. It operates based on two primary priorities:

1. Spinning Reserve (SR) for Generator Ramp Rate Support: The foremost priority involves ensuring that the generator's ramp rate requirements are met when the mission load is active. This aspect is termed "spinning reserve" (SR), and its focus is on maintaining the generator's ramp rate to accommodate the operational needs effectively.
2. Energy Storage Recharge using Model Predictive Control: The secondary priority pertains to recharging the energy storage system to achieve a predetermined State of Charge (SoC) level when generator ramp rate support is not immediately necessary. This is achieved by utilizing model predictive control (MPC) techniques to optimize the recharging process.

By integrating these two strategies, the PHRR controller optimally manages the generator's ramp rate support and energy storage recharging, ensuring the smooth operation of the mission load while maintaining overall system stability and efficiency.

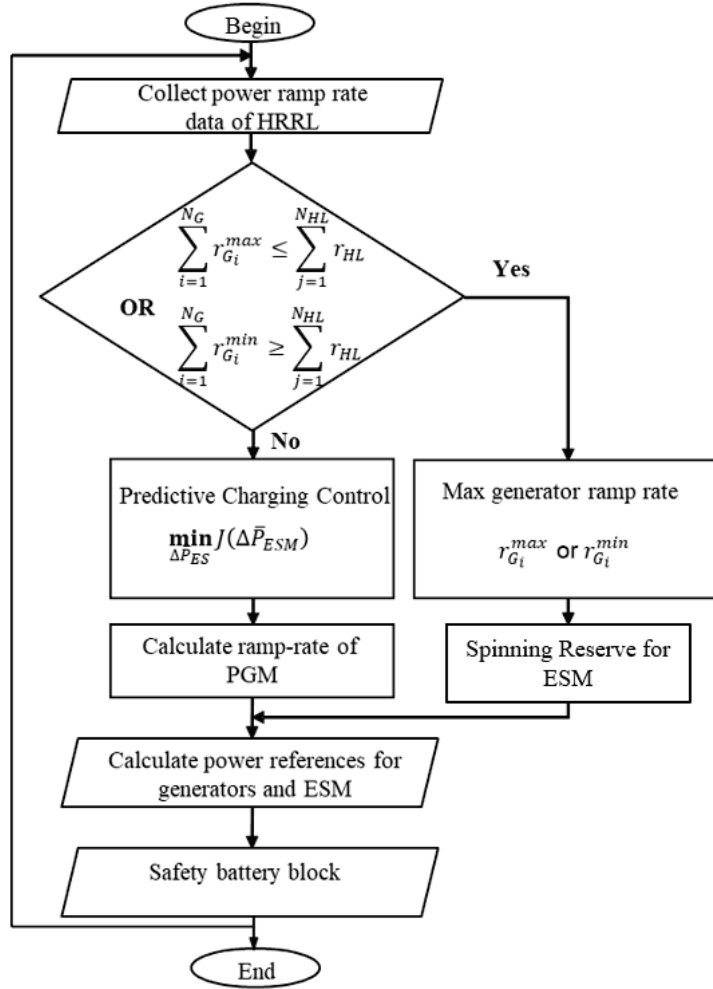


Figure 22. Flow chart diagram of the PHRR control. (© 2023 IEEE)

Figure 22 illustrates the flowchart diagram of the PHRR control process. The PHRR control operation operates by receiving input information from IDEAL through UDP packets at discrete time intervals represented by T_s . This input data serves as the basis for the controller's decision-making process.

Using the received data, the controller initiates the determination of whether a high ramp-rate load needs to be activated. This determination is carried out by comparing the aggregated power ramp rate of the generator with the requested ramp rate of the total load demand. Notably, for enhanced responsiveness, IDEAL can be configured to send a binary

enable signal one second prior to the load's intended activation. Upon assessing the activation status, the PHRR control exhibits a flexible behaviour. It dynamically transitions between the spinning reserve mode and predictive charging control mode. By seamlessly adapting between these two modes based on the activation status of HRRL, the PHRR controller ensures dynamic and optimized management of energy storage recharging and generator ramp rate support, aligning with the mission-based objectives and promoting overall system performance and reliability.

The primary goal of the predictive charging mode is to effectively regulate the SoC of the energy storage to a predefined user-defined value, all while adhering to the various physical constraints inherent to the system. To achieve this goal, a quadratic cost function J is formulated within the MPC framework, as expressed in equation (4). The aim is to minimize this cost function in order to optimize the charging process of the energy storage system:

$$\mathbf{min}_{\Delta \bar{P}_{ES}} J = \mathbf{min}_{\Delta \bar{P}_{ES}} \left((\bar{E}_{ES}^* - \bar{E}_{ES})^T (\bar{E}_{ES}^* - \bar{E}_{ES}) + \Delta \bar{P}_{ES}^T I_{N_c \times N_c} \Delta \bar{P}_{ES} \right) \quad (4.1)$$

subject to:

$$r_{ES}^{min} = r_{GEN}^{min} - r_L \leq r_{ES} \leq r_{ES}^{max} = r_{GEN}^{max} - r_L \quad (4.2)$$

$$P_{ES}^{min} = P_{GEN}^{min} - P_L \leq P_{ES} \leq P_{ES}^{max} = P_{GEN}^{max} - P_L \quad (4.3)$$

$$0 \leq E_{ES} \leq E_{ES}^{max} \quad (4.4)$$

$$P_{GEN} + P_{ES} = P_L \quad (4.5)$$

$$r_{GEN} + r_{ES} = r_L \quad (4.6)$$

where

$$\bar{E}_{ES}^* = [\bar{E}_{ES}^*]_{N_p \times 1} \quad (4.7)$$

The formulation of the MPC approach involves creating a predictive model with a prediction horizon of N_p and control steps of N_c . This predictive model is constructed based on various factors, including the characteristics of the energy storage model, forecasted load demands, and predicted power limitations of the generator. The MPC aims to find the solution for the quadratic cost function from equation (4.1), using $\Delta \bar{P}_{ES}$ which represents the power charge or discharge from the energy storage at each discrete time step T_s .

The quadratic cost function consists of two main terms: the first term minimizes the discrepancy between the current energy level and the desired energy level, while the second term focuses on keeping the rate of charge or discharge power from the energy storage as small as possible.

To ensure that the MPC solution satisfies the system's operational constraints, a series of equality and inequality constraints are introduced. These constraints are detailed in equations (4.1) to (4.7), with specific objectives:

- Equation (4.1) Is the cost function that is to be minimized to find the optimal ramp-rate
- Equation (4.2) ensures that the rate of charge or discharge of the energy storage remains within its physical capability, preventing abrupt changes that could stress the system.
- Equation (4.3) imposes a constraint to prevent the charge or discharge power of the energy storage from exceeding its operational limits.
- Equation (4.4) restricts the State of Charge of the energy storage to be below its maximum capacity, accounting for variations in energy levels.

- Equations (4.5) and (4.6) represent the balance of power and ramp rate constraints, respectively, ensuring smooth and continuous power flow while adhering to specified limitations.

The formulation of these constraints, along with the objective function, constitutes the foundation of the MPC approach. By solving this optimization problem at each time step T_s , the controller computes the optimal power adjustment from the energy storage, enabling the system to achieve the desired State of Charge while considering constraints and operational requirements.

The spinning reserve mode utilizes a low pass filter to capture the energy storage's reference power (P_{ES}^*) with the filtered power demand of the mission load (P_{ML}^{LPF}). This approach ensures smoother and more responsive energy storage output, addressing high ramp rate dynamics and enhancing system stability.

$$P_{ES}^* = P_{ML}^{LPF} \quad (5)$$

The reference current for the AC/DC power converter is calculated from the energy storage's reference power using Kirchhoff's law, as shown in equation (6).

$$I_{AC/DC}^* = P_{ML1} - \frac{P_{ES}^*}{U_{Dclink}} \quad (6)$$

The PHRR controller is designed to address a range of scenarios and critical parameters. These encompass factors such as the ongoing power consumption of the mission load, the generator's power supply to the mission load, and the status of the energy storage in relation to its minimum and maximum voltage and SoC limits. In standard operational circumstances, the generator's power output maintains a consistent and gradual

transition. However, in cases where a single battery cell's voltage falls below a specified minimum threshold, the ramping process is deactivated. This prompts the generator to assume the load's power supply, concurrently recharging the battery at its maximum charge rate until all cells achieve the user-defined nominal voltage level.

Likewise, if the battery recharged to reach the minimum threshold remains incomplete while a single cell surpasses the maximum voltage threshold, the recharging operation is ceased. This allows the battery management system to engage in rebalancing the system. The protective mechanism for the battery is confined within a dedicated Simulink subsystem. This component serves as the decision point, determining whether to utilize the reference current from the AC/DC power converter provided by the PHRR controller or to autonomously optimize the energy storage for safe and efficient operation. The graphical representation of this safeguarding process is presented in Figure 23

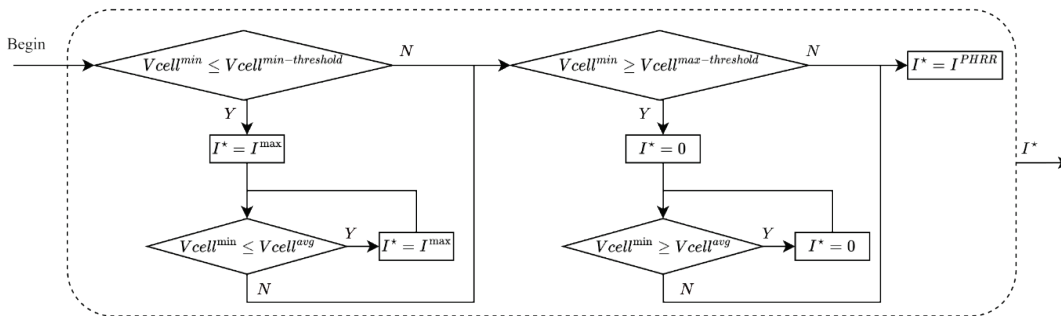


Figure 23. Diagram of battery protection mechanism. (© 2023 IEEE)

Integration of Clarkson ALS and PHRR controls into IDEAL

The Clarkson ALS controls and PHRR controls, though laden with potential, required refinements before their incorporation into an HIL simulation into IDEAL. This necessitated the creation of dedicated subsystems within LabVIEW, facilitating the

communication between the IDEAL controls and the SLDRT controls from CU. The adaptations made to the ALS controls encompassed several key aspects:

1. Precise power calculations are derived from the amalgamation of voltage and current measurements across various load steps.
2. Rounding of power values to whole integers to align with practical system capabilities.
3. Incorporation of the generator's rated power capacity to ensure coherent load shedding strategies.
4. Establishment of load state memory to avert inadvertent over-shedding of loads due to potential hardware latency.
5. Integration of a UDP (User Datagram Protocol) communication signal check to filter out invalid UDP packets that could potentially disrupt the system's integrity.
6. Safeguarding against the introduction of erroneous data regarding load ratings when the system initially executes a load shed.

Emphasizing points 4 and 6 from the aforementioned list, the introduction of a dedicated subsystem was imperative to prevent over-shedding in specific scenarios. Notably, the risk of over-shedding was particularly relevant to Load L3, controlled by a PCM with modifiable current outputs, and Load L1, characterized by load steps with a 1 kW step resolution. The underlying reason for this risk lay in the latency inherent between the two control systems and the inherent time delays associated with hardware measurements.

To address this challenge, an additional variable was introduced, monitoring the last received load shedding value and the last transmitted power value. A load shed

algorithm was implemented. After an initial loadshed was commanded, a calculation to determine the commanded power, the difference between the measured value and the value to shed, is conducted. If the commanded power value changes, then the system would respond to continue or stop the load shed, henceforth preventing the system from over shedding due to delays in the hardware to reach the desired goal. Additionally, to mitigate potential discrepancies arising from latency, particularly during load shedding events, a decision was made to momentarily saturate the max value of load steps. This was a proactive approach to counteract inaccuracies resulting from transient spikes in measurements during load shedding or enabling/disabling events. As the measurements converged back within acceptable ranges, the saturation was swiftly discontinued to maintain measurement accuracy and system fidelity. Figure 24 offers a visual representation of the dedicated LabVIEW system engineered to manage the intricate signal conditioning processes intrinsic to the ALS controls.

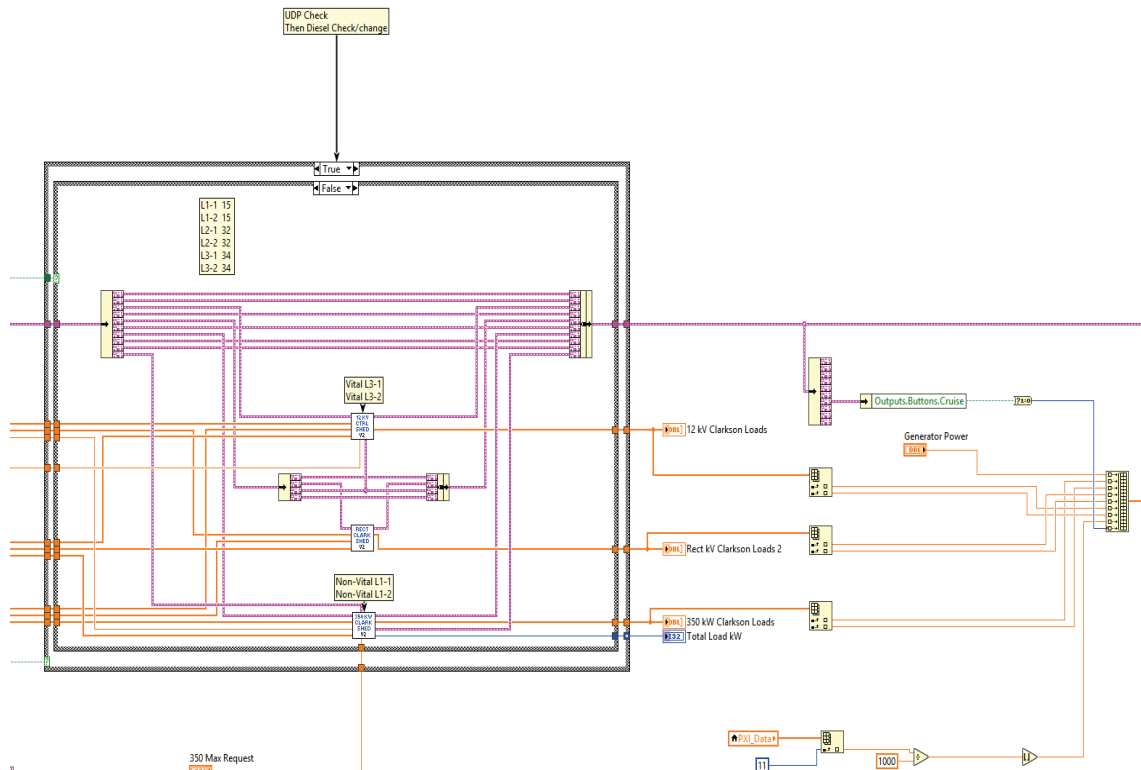


Figure 24. NI LabVIEW system dedicated to handling inputs and outputs for the deployment of the ALS controller from Clarkson.

The integration of the PHRR control into the IDEAL system necessitated a series of deliberate modifications, each contributing to its seamless operation within the broader framework. The key adaptations made to the PHRR control are as follows:

1. Development of a sophisticated lowpass filter mechanism to regulate the ramped power dynamics on the bus, ensuring a controlled and gradual transition.
2. Implementation of an accurate battery state of charge assessment through a combination of initial state determination and coulomb counting techniques.
3. Calculation of average battery voltages to enhance the precision of the control algorithm's decision-making process.
4. Creation of an intelligently informed mission load signal, designed to activate immediately upon the initiation of a user command.

5. Introduction of an enabled mission load signal, timed to activate one second subsequent to the initiation of the informed mission load signal.
6. Integration of a UDP check mechanism, validating the legitimacy of incoming UDP packets to safeguard against erroneous modifications to the controller values.

Notably, point 1 underscores the importance of the lowpass filter in orchestrating the dynamics of the ramp rate. To elaborate on its development, a conceptual model consisting of four discrete transfer blocks was presented to UTA from CU, outlining the desired characteristics of the ramp rate output. This ramp rate however suffered from its calculation being tied directly to the timestep of the model and thus limited its viability. UTA refined this concept by constructing its own lowpass filter using a combination of a moving mean and two low pass filter blocks within the Simulink environment that approximately matched what CU presented. This approach was chosen for its superior fidelity, streamlined processing, and consistent performance, even when subjected to variations in the user-defined timestep.

By addressing these crucial aspects, the PHRR control was set up to align with the dynamics of the IDEAL systems, ensuring accurate power modulation and efficient energy management. These technical refinements reflect a pursuit of precision and reliability in the control system's operation, ultimately contributing to the enhanced performance and efficacy of the PHRR control within the IDEAL environment. Figure 25 contains the LabVIEW VI dedicated with interfacing with the hardware and commands from the PHRR controller.

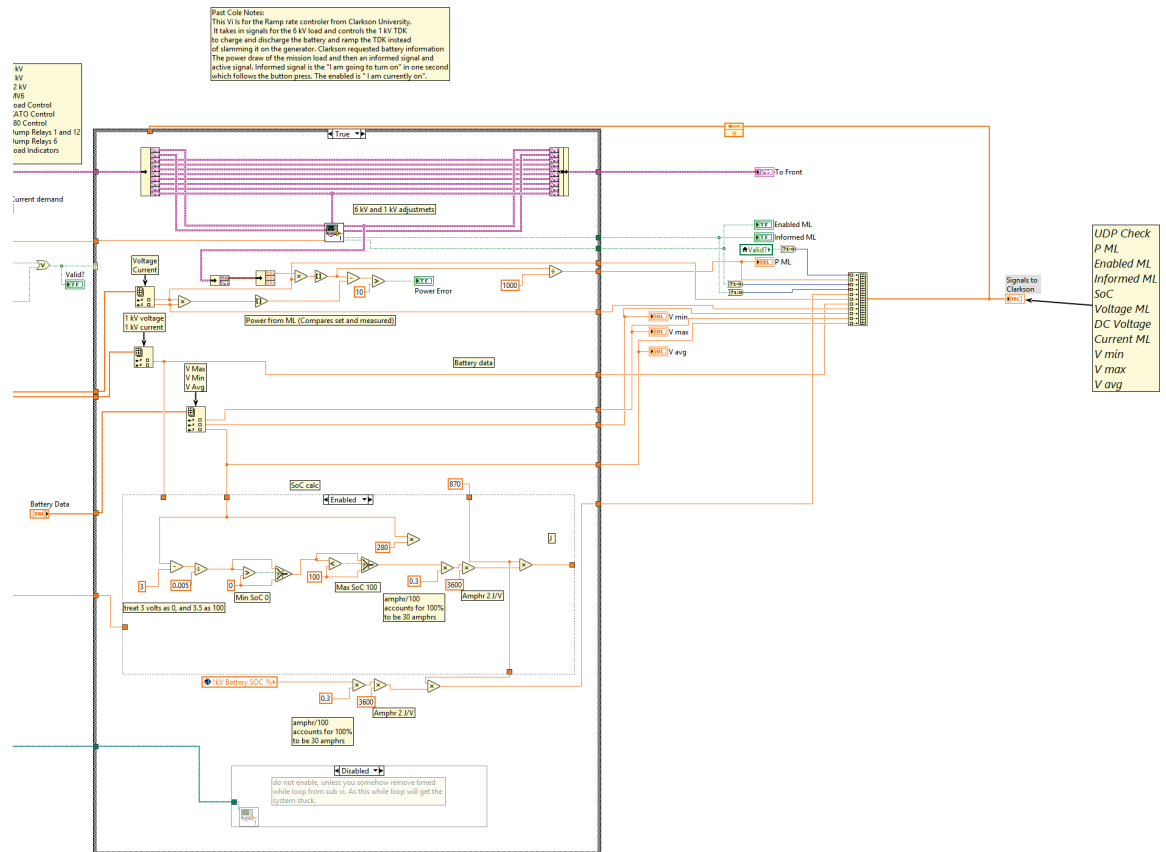


Figure 25. LabVIEW subsystem inside of the IDEAL controller that handles the integration of the PHRR system.

Upon the successful incorporation of the LabVIEW subsystems, the controllers were poised for integration into the IDEAL testbed. The subsequent results chapter will delve into comprehensive Simulink simulations, demonstrating the controllers within an optimal system framework with no communication delay or hardware delay. Following this, the ensuing sections will unveil the tangible outcomes of these controllers on real hardware, effectively portraying their real-world performance on the UTA hardware platform.

RESULTS

BMS Test Stand Results

This section presents empirical evidence highlighting the successful cell balancing across various platforms and the capability to swiftly adjust cells' SoC. To begin, Figure 26 and Figure 27 portrays two digitally generated graphs, illustrating projected operations of a BMS in distinct scenarios. In Figure 26, the initial green segment depicts a conventional charging sequence, revealing a clear imbalance between two cells. The cell with higher SoC attains its maximum charging voltage of 3.65 V, leading to the stop of the charge cycle, which subsequently ensues a balance process. Notably, in this graphical representation, it is assumed that the recharge current significantly surpasses the balance current, resulting in an imperceptible shift in the charge current beyond the balance voltage threshold. However, it is important to acknowledge that the balance circuit would activate at this juncture to dissipate excess energy. Upon reaching the balance voltage threshold, the charge process recommences until balance between the cells is achieved in the yellow segment.

Moving forward, the first pink segment showcases a discharge procedure where two cells discharge at varying rates. As one cell reaches its minimum allowable voltage, the discharge process is promptly halted. The subsequent green segment depicts a comparable recharge process to that previously outlined, followed by a merged discharge and recharge cycle presented in the ensuing blue segment.

In Figure 27, the identical charge and balance process is portrayed within the initial green segment. Comparatively, the pink segment underscores the agility in swiftly altering the SoC of one or more cells. Notably, both cells' SoC experiences a rapid reduction,

exemplifying how the BMS or host controller should gauge the minimum allowable voltage and promptly disconnect the battery from a source or load. In the subsequent blue section, a swift elevation in the SoC of a single cell up to the maximum allowable voltage is showcased. Upon reaching this threshold, the battery is once again disconnected from a source or load, initiating the balance process on the high SoC cell. Following this, several routine charge and balance sequences are displayed.

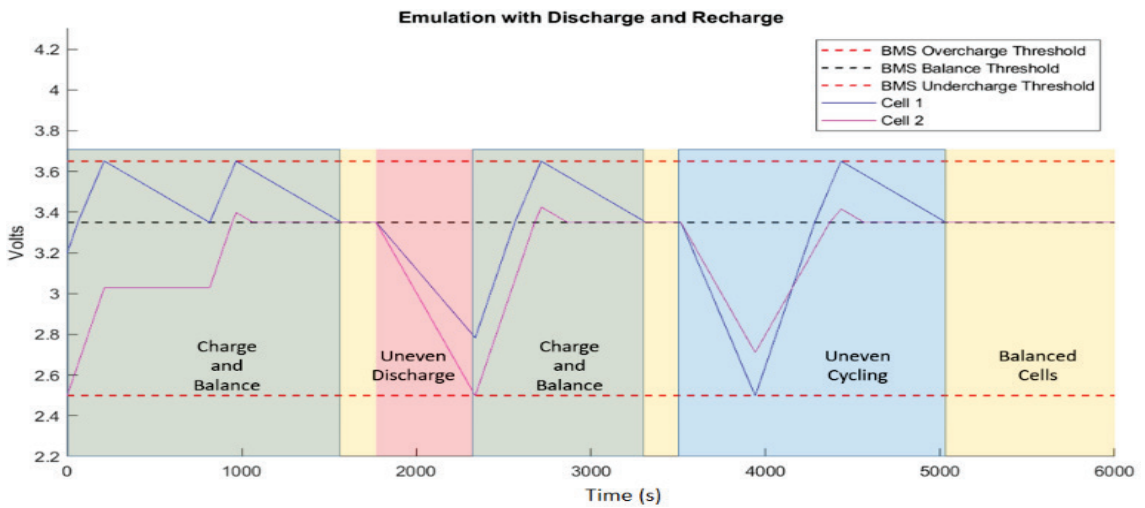


Figure 26. Predicted BMS behaviour with 2 cells in normal operation. (© 2021 IEEE)

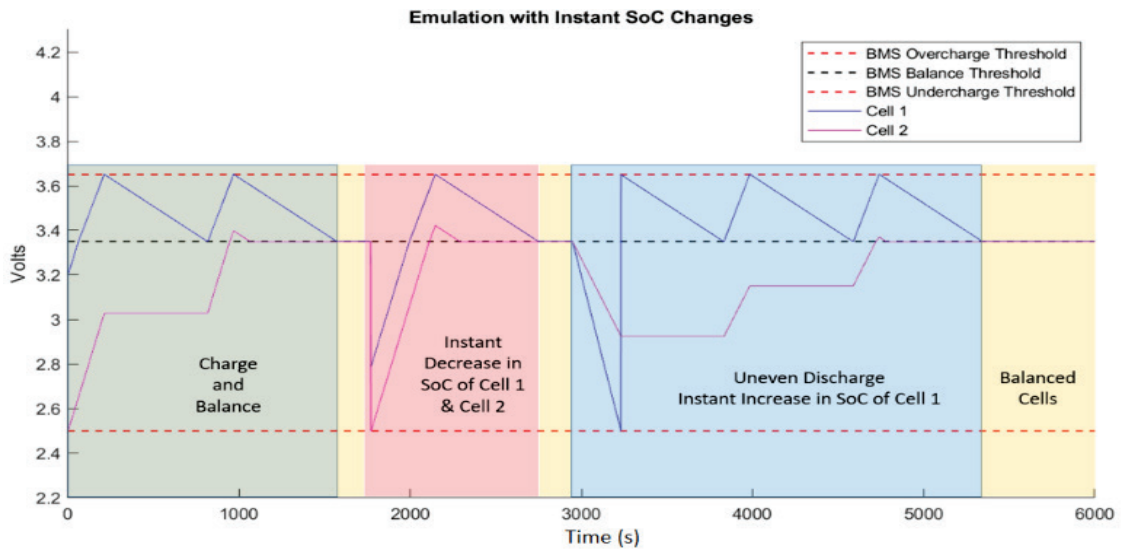


Figure 27. Predicted BMS behaviour with 2 cells. The graph demonstrates the user's ability to rapidly change individual cell(s) SoC to evaluate BMS behaviour quickly. (© 2021 IEEE)

Emulation with IO991

The IO991 modules, which constitute Speedgoat's legacy hardware, govern the cell outputs based on SoC values for the VL30AF model discussed previously, enabling the BMS to execute cell balancing. However, the legacy IO991 cards lack the capacity to measure the current supplied by each channel or cell. To address this limitation, Speedgoat provided us with an IO323 analog input card. Utilizing this card in conjunction with a specially designed in-house PCB, we successfully integrated current feedback into the system. By effectively measuring the current output from each channel, the model becomes capable of dynamically updating the SoC of each battery, enabling real-time adjustments to the emulated cell voltage. A simplified overview schematic of the current sense PCB is illustrated in Figure 28, while Figure 29 showcases images of the populated board. Notably, this circuit board is only essential for the Speedgoat IO991 modules, and not the other emulators presented.

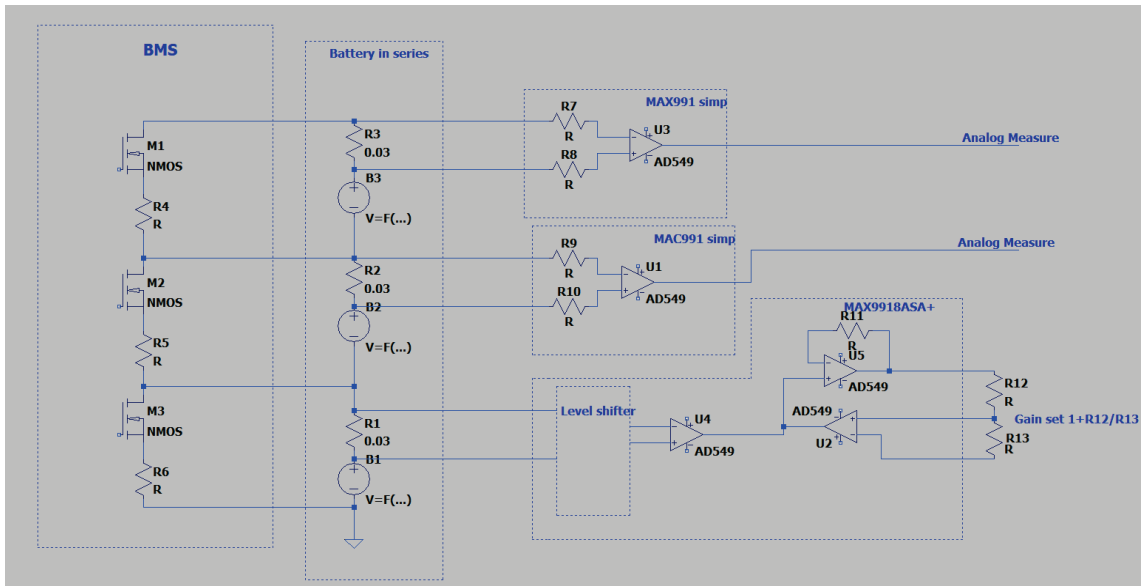


Figure 28. Simplified Circuit diagram of the battery and the current measurement to be implemented.




Figure 29. Current sense PCB attached to a BMS and IO991 modules.

After establishing the initial expectations, outlined in the beginning of this chapter, we will now present a selection of hardware results with the IO991. Figure 30 offers a depiction of the measured voltages and balance currents obtained from five out of the ten cells emulated through the IO991 interface, which is integrated with a BMS. The BMS assessment begins with Actia's proprietary BMS. While only five cells are visually represented to enhance clarity in the graphs, it's important to note that the BMS governs all ten cells. In this particular scenario, the model is in a quiescent state, unconnected to any external source or load, thus solely engaging in cell balancing activities.

Specifically, Cells 1, 2, and 3 share an identical SoC of 55%, while Cell 5 maintains a substantially higher SoC, nearing 100%. Meanwhile, Cell 4 assumes an intermediate position, hovering around 85%. A distinct observation arises from Cells 4 and 5 due to their notably elevated SoCs compared to the other three. Consequently, their balance circuits remain persistently active, drawing approximately 45 mA each from the BMS to facilitate their equilibrium. As the BMS draws current, the corresponding sense circuit measures it and promptly feeds the information back into the model, thereby facilitating adjustments to the output voltage in accordance with the updated SoC.

Cell behavior for Cells 3, 4, and 5 present intriguing dynamics. A distinctive departure emerges between the behavior of the BMS by Actia and the projected outcomes. Unlike the fixed balance point indicated in Figure 27, the BMS demonstrates a different operational pattern. Instead of adhering to a predetermined threshold, the balance point aligns with the lowest cell voltage. This design characteristic leads to a noticeable hysteresis effect. Fluctuations within the voltage sense circuit of the BMS lead to slight oscillations in voltage. Consequently, the processor alternates between which voltage is

marginally higher, thus causing an intermittent activation and deactivation of the balance mechanism.




Figure 30. Voltage and current data collected while 5 emulated cells are connected and balanced by the BMS. Solid lines indicate cell current while dashed lines represent cell voltages. (© 2021 IEEE)

Figure 31 underscores the capability of swiftly altering a cell's SoC. At the outset of data collection, all cells are uniformly set at a SoC of 55%. Subsequently, the SoCs of Cell 2, Cell 4, and Cell 5 are deliberately escalated to 60%, 85%, and 57% respectively, at sporadic intervals between 50,000 s and 70,000 s during the simulation. Whenever a cell's SoC surpasses the minimum cell voltage, the balance circuit of that cell activates, drawing current. The distinctive behaviors exhibited by balanced cells such as Cell 1 and Cell 3, as previously observed, persist due to their unchanged SoCs, resulting in their balance circuits being disengaged, or acting in the hysteresis manner previously discussed. This then extends to cells brought into balance as well after they are finished balancing.

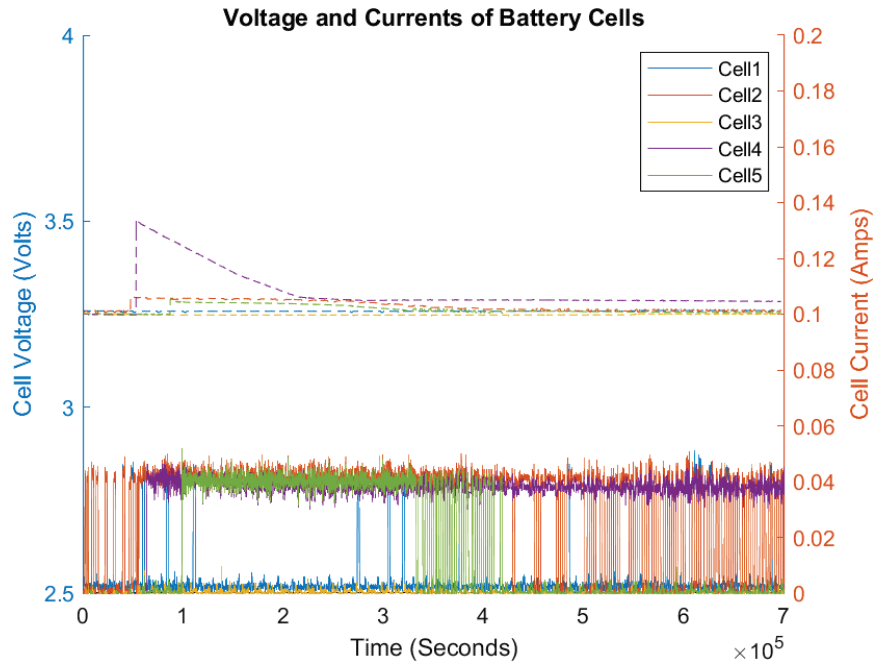


Figure 31. The graph demonstrates the ability to rapidly change a cell(s) SoC to evaluate a BMS’s performance near cell bounds quickly. Solid lines indicate cell current while dashed lines represent cell voltages. (© 2021 IEEE)

It is imperative to note that this ability to swiftly manipulate a cell's SoC, as demonstrated here, is not feasible in real-world scenarios. However, this feature significantly expedites the comprehensive assessment of BMS performance when cells approach their operational limits. Furthermore, it facilitates controlled exploration of the BMS's overvoltage characteristics, which would be challenging to achieve directly with a functioning battery due to safety considerations surrounding overcharging.

Subsequently, the IO991 modules underwent testing with a BMS manufactured by K2 Energy. This particular BMS consists of three distinct PCBs, each serving specific functions encompassing balancing, monitoring, and communication. The foremost PCB operates as the master, observing the activities of secondary cards linked to the modules. It boasts a CAN communication output that interfaces with a communication board. This communication board, in conjunction with the master, employs a fiber optic connection to

emit a signal, effectively mitigating electrical interference during battery or PCM operations. Completing the configuration are ten cell secondary cards, which are individually responsible for balancing the cells within a module. These modules exhibit a higher balance current in comparison to the Actia BMS, employing a balance resistor of approximately 39 ohms to yield an approximate balance current of 100 milliamps.

The data collected from experiments conducted on the K2 BMS is directly extracted from LabVIEW, leveraging the UDP packets transmitted by the BMS communication board. Thus results for this BMS will typically include the commanded model values as well as the measured values from the BMS. In contrast, the Actia values pertain to the model-generated data. An illustrative depiction of this setup is furnished in Figure 32.




Figure 32. Image of K2 BMS board and current sense board (in black). The green card to the left is the Slave card. The master is the white PCB in the middle. The blue board is for communications. The K2 BMS demonstrates operational patterns consistent with the anticipated behavior illustrated in Figure 27. Unlike the Actia BMS in attempting to balance cells towards the lowest recorded voltage, the BMS endeavors to achieve a

designated balance threshold. Evaluating the K2 BMS's balancing functionality necessitates a modified testing approach. In this case, the cells are charged until they reach a maximum threshold, subsequently prompting the BMS to balance the cells down to the predetermined balance threshold before initiating another charging cycle. Notably, the recharging sequences in the K2 BMS are emulated within software, avoiding an external power supply to avert potential external loading effects. This approach is employed to ensure a controlled environment. A representative instance of a test involving the IO991 card on the K2 BMS is depicted in Figure 33.




Figure 33. K2 balance test. This BMS requires cycles of recharging after balancing to bring the cells into proper balance.

The assessment performed on the K2 BMS, along with the recorded external measurements, brings to light a noteworthy observation regarding potential noise on the output of the IO991 module. Subsequent comparisons with tests conducted on the Chroma equipment reveal that the measurements exhibit considerably reduced noise or variance when employed with the K2 BMS. This discrepancy prompts an exploration into potential factors influencing the noise levels, such as the current sensing methodology or other variables not directly linked to the IO991 hardware. It is imperative to account for this

aspect when interpreting test outcomes and considering the potential impact on hysteresis effects.

Emulation with Chroma 87001

Similarly, the Chroma emulator was interfaced with the K2 BMS in a manner analogous to that of the IO991 without the current sense PCB. This 16-cell emulator, functioning with a response rate of 10 ms, utilizes a distinctive approach where multiple commands are dispatched and subsequently updated simultaneously at the conclusion of the cycle, resulting in an effective update rate of approximately 180 ms. Notably, this emulator is equipped with an integrated current sensing capability, enabling direct measurement from the emulator itself. A notable divergence between the Chroma emulator and the Speedgoat legacy emulator lies in the more consistent voltage output observed in the case of the Chroma system as recorded by the K2 BMS. The arrangement of the Chroma unit, securely mounted within a rack, is depicted in Figure 34. Due to internal incorporation of current sensing in the Chroma unit, wiring of cells in series is simply done using ring terminals and barrier blocks.



Figure 34. Chroma unit in white. It is a 16 cell emulator capable of up to 5 amps continuous current on each channel.

Communication with the Chroma unit is facilitated through a TCP connection, involving the configuration of the unit's IP address located on its rear panel. Once connectivity is established, control commands can be issued from MATLAB or other compatible software like NI MAX. Employing the instrument commands first in NI-MAX, allows validation of them before the configurations were incorporated into the MATLAB environment. Within MATLAB, a suite of commands *is* deployed, encompassing conventional VISA instrument commands such as *IDN, along with custom directives tailored to voltage adjustment and current retrieval on the Chroma 87001 unit. A dedicated MATLAB toolbox, showcased in Figure 35, streamlines the communication process with the Chroma emulator.

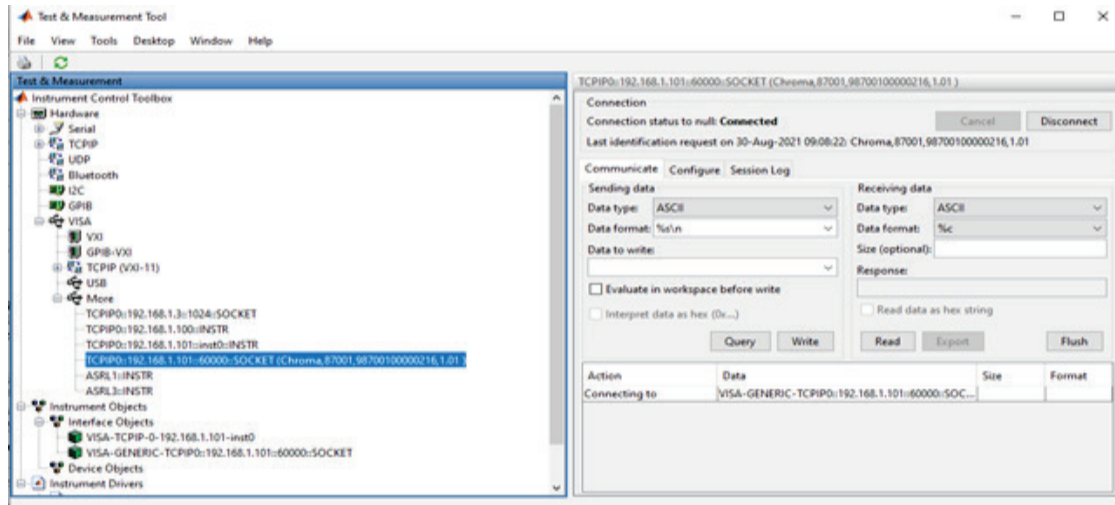


Figure 35. Toolbox to communicate with instrument in MATLAB.

Significant endeavors were dedicated to integrating the Chroma unit within the Simulink and SLRT framework. While an approach involving the invocation of code by a timer function in MATLAB, with updates from Simulink's battery values, is feasible, its response rate is inherently confined to around 500 ms. This temporal limit hinges on the timer's mechanics and the performance of the host PC, thereby introducing potential variability and inconsistent updates. Moreover, this method inadvertently entangles the development PC in the communication loop, thereby compounding intricacies.

To circumvent these limitations, Simulink's TCP blocks offer a more optimal solution. This avenue enables direct communication between the target computer and the Chroma system, fostering swifter and more consistent response times without encumbering the resources of the development computer. A purpose-built subsystem, architected for TCP communication reads and writes is visually encapsulated in Figure 36.

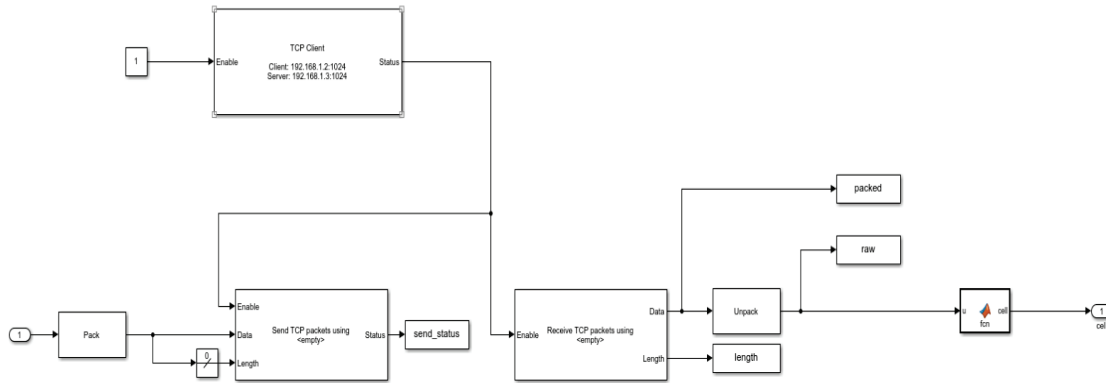


Figure 36. Block diagram depicting use of TCP blocks to connect to Chroma 87001

Leveraging the pre-established TCP connection block diagram necessitates the formulation of predefined ASCII commands. These commands configure the desired voltages for the Chroma 87001 and solicit the current measurements. Once these ASCII commands are shaped into byte representations, they are channeled into the instrument. The TCP send and receive blocks can be cyclically employed, taking cues from the client's status connection, while the readout data is systematically structured to adhere to the desired data format. Notably, this identical system framework could extend to other emulators like Hioki, necessitating only adaptations to the underlying ASCII commands and address. The functional subsystem depicted in Figure 37 is currently operational and actively employed to regulate the Chroma 87001.




Figure 37. Chroma TCP connect for standard Ethernet interface.

The encountered delay times stemming from the use of the Chroma emulator have posed a challenge if more units were to be used. It is crucial to note that the documented response time of 10 ms per instruction is a significant factor. Guiding the system to update the output channels constitutes a single command, whilst retrieving each cells' current values requires an additional command. Finally, a command is needed to update the voltage value of a single cell. Consequently, for the update of 16 individual cells, a total of 18 commands are required, translating to an approximate response rate of 180 ms. While diverse command structures and communication methods may exist, this timing consideration would apply to each cycle. In scenarios involving multiple daisy-chained Chroma units, these times would cumulatively accumulate. To mitigate excessive delay times when multiple Chroma units are utilized, an alternative solution could involve parallelizing the units using a switch, rather than daisy-chaining their TCP connections.

An illustrative instance of the Chroma system integrated onto the K2 BMS, achieving balance with a substantial SoC difference exceeding 50% between cells, is visually depicted in Figure 38. In this specific test, the balance initiation point is established

at 3.4 volts, while the termination threshold for balance is set at 3.35 volts. The 3.8-volt upper limit is set within Simulink through code governing software-based recharge. The battery attains balance at approximately 375,000 seconds, commencing the simulation with the highest cell's SoC at 98% and the lowest cell's SoC at 45%.

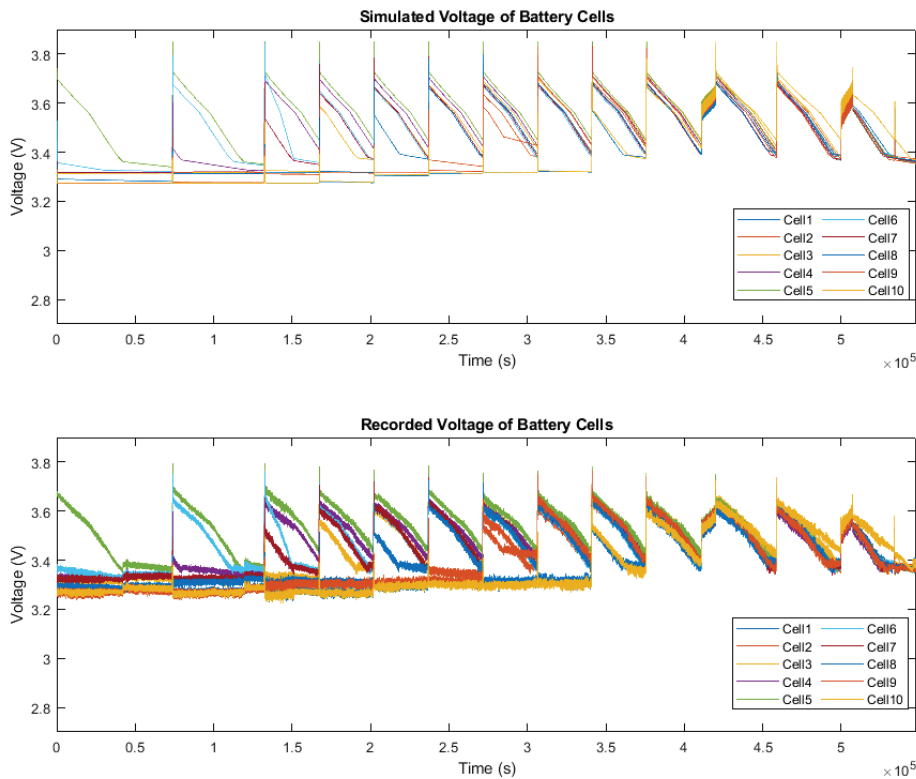


Figure 38. Chroma system balancing from K2 BMS. Top graph contains the values commanded by Simulink while the bottom graph contains the voltage values recorded by the BMS.

During the process of collecting data on the K2 BMS cell current, an unexpected and distinct failure event transpired. This failure led to several cells being balanced down to a voltage of 0. This appears to result from locked balance currents, causing them to drop below the designated balance threshold. Notably, while cell 3 demonstrated anticipated behavior with the BMS, undergoing cycles of discharge, cell 7, among others, experienced a discharge down to zero volts. This occurrence underscores the significance of the testbed,

substantiating its capability to identify potential hazards associated with the BMS. Furthermore, it establishes the invaluable role of the testbed in risk mitigation and damage prevention, offering an avenue to assess and address such failure scenarios without jeopardizing real battery systems. Visual depictions of voltages and currents are illustrated in Figure 39 and Figure 40 respectively.




Figure 39. Graphs of voltages from K2 BMS with Chroma emulator. Graph shows the voltage of cells falling until their output drops to zero and current is no longer drawn.

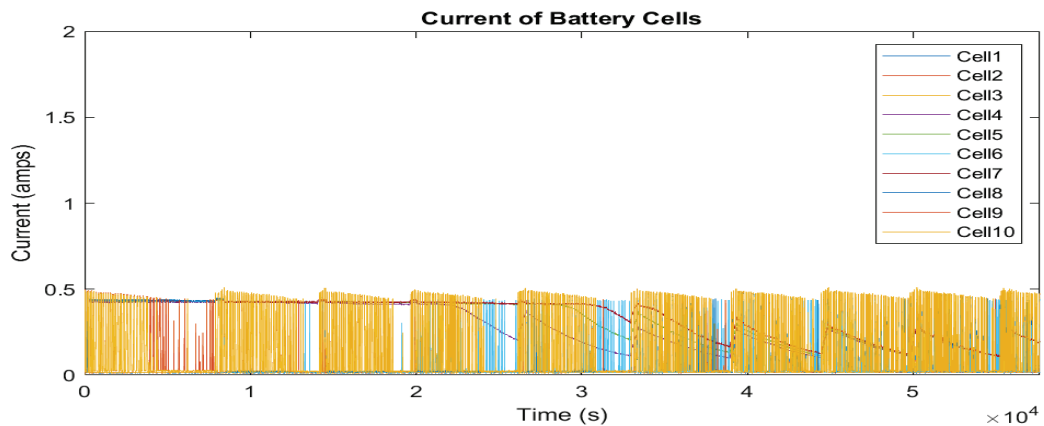


Figure 40. Graphs of currents from K2BMS with Chroma emulator. Graph shows constant current being drawn from several cells until their output drops to zero and current can no longer be maintained.

Emulation with Speedgoat BCE

In October 2020, Speedgoat conveyed a delay in the production and supply of the cell emulator hardware initially scheduled for September 2021 however, the finalized product was ultimately delivered in April 2023. An initial prototype unit was received in September 2021, but it necessitated substantial debugging efforts and did not pass the EMI test. Subsequently, five more prototype units were provided in March 2022, enabling further testing to take place. By August 2022, an update to the communication was implemented, streamlining usability. Despite progress, ongoing EMI challenges prompted additional delays, as communicated to us in June 2022. Encouragingly, by August 2022, Speedgoat indicated confidence in having resolved the EMI concern. The definitive units were received in April 2023, and subsequent final assembly and testing operations were successfully executed on the testbed. Results with the Chroma and IO991 discussed previously were conducted in the meantime before receiving the final products.

To establish the final configuration, A PCB design was devised, with the purpose of enabling the connection of emulator units to the BMS while concurrently interfacing with each channel from the emulator to form modules and the battery. This circuit board connects the serial wiring of individual cells sourced from the Battery Cell Emulator units (BCE), affording a streamlined connector that seamlessly interfaces with the K2 BMS. A total of 26 module-level cards were manufactured and integrated within the racks, aligning with the 26 individual module-level cards.

These circuit boards were subsequently mounted onto an insulated acrylic panel, ensuring an effective 1000 V isolation between the channels and the metal cabinet. Each board contains both positive and negative output lines to facilitate series or parallel interconnections of multiple boards. Additionally, the circuit board offers two distinct

options for BMS or other load attachments. The first output connection manifests as a barrier block, accommodating ring terminal connections between nodes within the series string. The second connection, a Molex 53426 series connector, similarly allows connection to all 10 cells/ 11 nodes on the circuit board. While the Molex connection exhibits a 3-amp limitation, it provides a swift and efficient connection avenue for BMS units, while simultaneously extending a higher-current alternative for loading individual cells. Figure 41 contains a single assembled PCB module card.




Figure 41. Completed circuit board to attach the cells in series and interface with the BMS.

Upon the receipt of the Speedgoat cell emulators, dedicated efforts were directed towards their integration within their designated cabinets. The initial cabinet is equipped to accommodate ten BCE units, a target computer, and a 48V power supply. Meanwhile, the second cabinet is tailored to house 12 BCE emulators, an identical power supply, and an additional target machine. Owing to spatial constraints encountered in the first cabinet, along with considerations for the PCBs, the decision was made to exclude the Chroma unit

from the final configuration of the rack setup. The connectivity arrangement includes an Ethernet cable linking the desktop to the target computer within the first cabinet, as well as an Ethernet cable and a voltage-string cable interlinking the two cabinets. The culmination of these endeavors yields a cabinet featuring 12 fully operational 10S modules, while the other cabinet boasts 14 complete modules, with an additional 4 cells remaining unused.

To provide power to each BCE, the 48-volt supply originating from the top of the cabinet was routed to the middle section behind the emulators. Within this designated space, two copper plates were securely installed, featuring insulated feet, serving as efficient power buses for each BCE unit. Subsequently, the emulators were connected to these copper buses to access the power supply. For power and ground connections, each unit was outfitted with 16AWG cables equipped with wire ferrules to prevent fraying. The cables linking the BCE channels to the circuit boards were also 16AWG, terminated with Molex flex fit connectors and then channeled through a protective wire sheath. Visual representations of the final configuration are captured in Figure 42.



Figure 42. Front and back of one of the BMS testbed cabinets

Analogous to the approach taken for the Chroma system, a dedicated subsystem was developed to oversee the EtherCAT framework responsible for managing and communicating with the BCE units. This subsystem encompasses 22 modules, each serving as a crucial interface. To speed up altering values within each subsystem, a MATLAB script was developed. This script empowers the automated configuration of individual subsystem block settings, a strategic solution that proves especially advantageous considering the near hundred EtherCAT transmits and receives inherent to each subsystem. This approach significantly streamlines the task of adjusting communications across multiple tests, ultimately saving users valuable time and effort.

For seamless execution of this subsystem, the presence of an ENI (EtherCAT Network Information) file becomes necessary. This vital file, generated via TwinCAT software, must be tailored to the specific count of units in operation. A visual depiction of this model is presented in Figure 43. Notably, the model incorporates switches, offering

the user the flexibility to enable or disable BCE units at will. This functionality proves particularly useful, enabling a gradual system activation process.




Figure 43. BCE software to enable multiple modules and control them efficiently at once.

USC and FSU Cosimulation

This section delves into the outcomes arising from the PHIL simulation conducted in collaboration with USC and FSU. As a recapitulation, this experimental setup involves FSU's PGM transmitting data to UTA, which subsequently provides power to the emulated 1kV battery and PCM which command a TDK power supply. The load emulation, in turn, is facilitated by a bi-directional Chroma supply. The overarching objective is to showcase the practical implementation of Power PHIL onto the IDEAL platform, along with the deployment of UTA's control systems, thereby spotlighting ramp rate control and its effect on a PGM.

To control the ramp rate, a hyperbolic tangent function was pursued. This starts and ends the ramp at a low rate of change, all the while having the center of the ramp rate at the steepest, thus giving the generator time to correct and account for a steady increasing rate or decreasing rate of load. The model that handles the ramp rate is shown in Figure 44

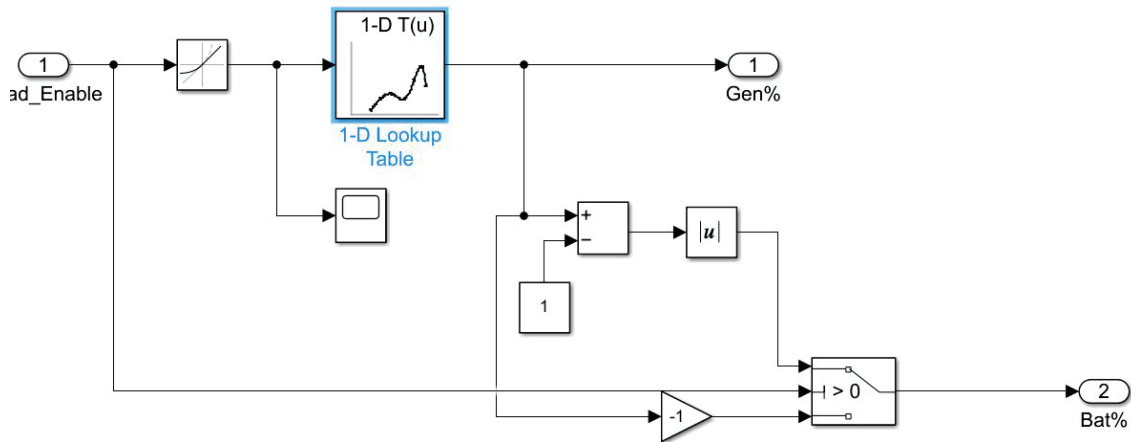


Figure 44. The specific subsystem that forms the ramp rate for cosimulation.

Consider an instance of a square pulse emanating from Load #2. As this transient load springs to life, it commands a square-shaped pulse over a span of 5 seconds. In the initial moments of activation, PCM #4 collaborates with the battery to entirely meet the load's requirements. Subsequently, PCM #3 enters the scene, providing a gradual increase in the 12 kVDC Bus's contribution while concurrently diminishing the battery's input. This continues until the 12 kVDC bus progressively shoulders 80% of Load #2's power requisites.

Upon the end of this transient load event, the battery absorbs the surplus power provided by the 12 kVDC bus from FSU's PGM. PCM #3 begins a taper of its power, culminating in a partial recharge of the battery. This choreographed sequence of actions

exemplifies the interplay between different power sources and the control mechanisms at play.

The purpose of the ramp rate support function is to serve as a buffer for the shared node, effectively mitigating most adverse effects on bus power quality caused by the transient load. The adjustable ramp rate can be fine-tuned and optimized based on the current state of the power system. In scenarios where there is no buffering, the load's transition to an on or off state could lead to an abrupt loading or unloading of the shared bus, potentially introducing poor power quality. However, with buffering enabled, the energy storage system ensures that the shared bus retains, or closely approximates, the voltage output it was providing during and after the load's transition, thereby optimizing power quality on the shared bus. Figure 45 shows the results of this cosimulation, with the ramp rate implemented on the top row of graphs, and the bottom row showing the results without ramp rate. The effects of the ramp rate and the improvement power quality are clear as the voltage spikes are reduced from approximately 400V to under 100V on the 12kV bus during initial loading and unloading.

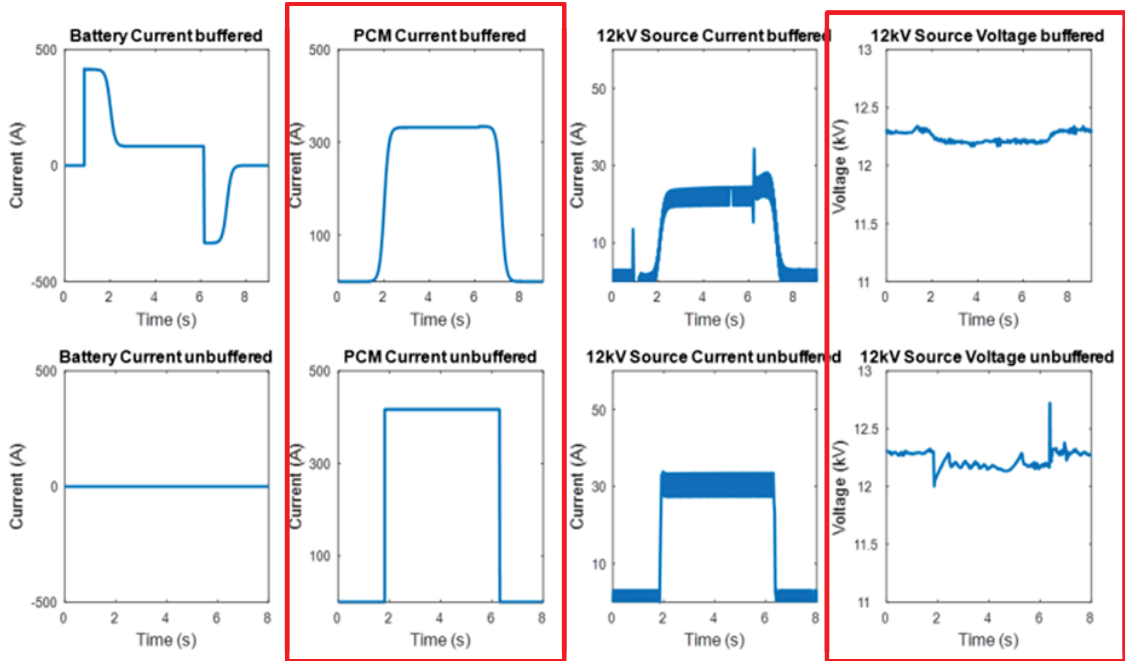


Figure 45. Co-Simulation results with FSU and USC. Bottom series of graphs contain the unbuffered bus, The top series of graphs are the buffered version, and results with reduction of spikes on the PGM. (© 2022 IEEE)

ALS and PHRR Load Scenarios

A series of experiments was conducted to showcase the successful integration and functioning of the CU ALS and PHRR algorithms on the UTA IDEAL testbed. This section presents five distinct test scenarios: the initial two focus on demonstrating the standalone operation of the ALS controller, the subsequent two highlight the standalone operation of the PHRR controller, and the final scenario illustrates the integration of both controllers. Throughout these scenarios, the defined loads and their corresponding parameters remain consistent, with the control system sampling occurring at approximately 100 ms for the ALS controller and 10 ms for the PHRR controller.

ALS Scenario 1: Load priority change and generator overload

In this scenario, we demonstrate an ALS load shed based on a predetermined set of ALS priorities. Following the initial load shedding, the user modifies the priorities, revealing how the ALS responds to another overload event. A simulation of this experiment is illustrated in Figure 46. The power generation capability of G1 is 150 kW, with 95% available for supply and 5% reserved. Table 2 outlines the seven loads and two priority sets involved.

Table 2. Parameters of loads and generator for case scenario 1			
Description	Rating (kW)	Original Weights	Adjusted Weights
Non-Vital Load 1-1	15	0.2	0.2
Non-Vital Load 1-2	15	0.2	0.2
Vital Load 2-1	32	1	1
Non-Vital Load 2-2	32	<u>0.2</u>	<u>0.5</u>
Vital Load 3-1	34	0.5	0.5
Vital Load 3-2	35	0.5	0.5
Mission Load (ML)	70	1	1
G1 M-G Set	150*		

*Over limit is set by the controller at 143 kW

Initially, the operator activates Non-Vital L1-1, Non-Vital L1-2, Vital L2-1, Non-Vital L2-2, and Vital L3-1, totalling around 128 kW. At t = 25 s, when Vital L3-2 is enabled by the operator, 143 kW is exceeded, triggering the ALS to shed Non-Vital L2-2. This choice aligns with the ALS strategy of minimizing load shedding while remaining below the 143 kW threshold. An alternative approach could involve shedding Non-Vital L1-1 and

Non-Vital L1-2, but that would shed more loads, deviating from the ALS objective of optimal load operation. At $t = 32$ s, the operator adjusts load priorities, and at $t = 40$ s, Non-Vital L2-2 is reactivated, once again surpassing 143 kW. This time, due to the modified priorities, Non-Vital L1-1 and Non-Vital L1-2 are partially shed to manage the overload.




Figure 46. The simulated results of scenario 1 on the ALS controller. (© 2023 IEEE)

Following the successful simulation of the scenario, the ALS controller was implemented on the IDEAL hardware to execute the experiment. The experimental data obtained from the PXI chassis and interpreted by the VI and sent to SLDRT is illustrated in Figure 47. Additionally, Figure 48 presents the directly measured PXI data during the operation of the ALS controller. At the scenario's outset, the operator activates the loads.

Mirroring the events depicted in Figure 46, ALS controller load sheds at $t = 27$ s and $t = 40$ s, affirming the seamless integration of the hardware with the controller.




Figure 47. Experimental load values measured by the ALS controller. (© 2023 IEEE)

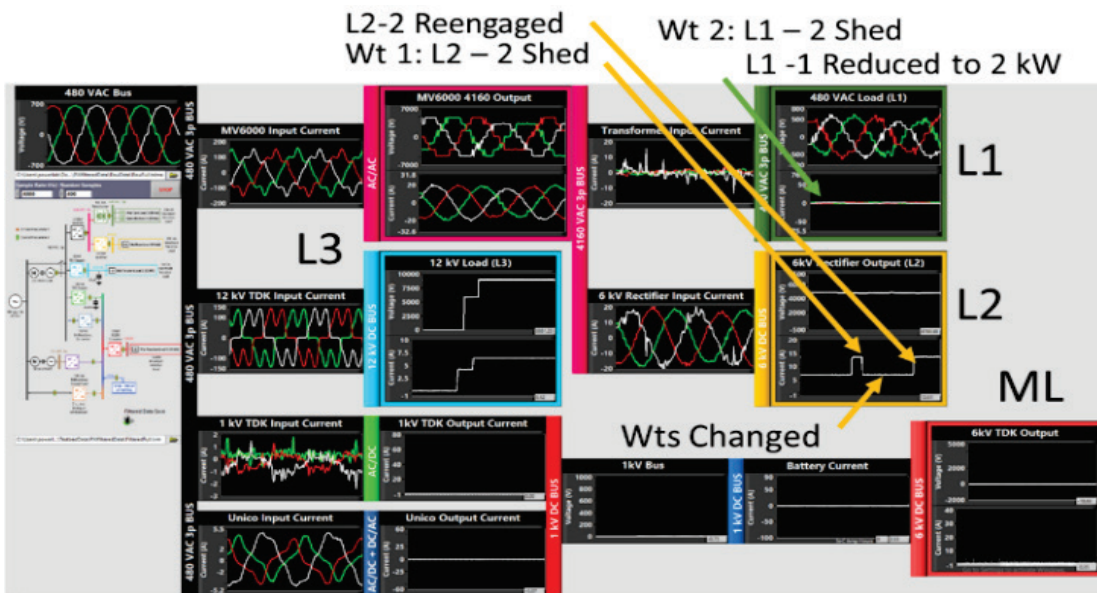


Figure 48. Experimental data collected by the PXI chassis during the operation of scenario 1. (© 2023 IEEE)

ALS Scenario 2: Generator loss of capacity

Load shedding becomes imperative in scenarios involving generator derating, a situation where the generator's maximum power output diminishes. This results in an inability to fulfill the current load demand due to the reduced generator power output. A scenario akin to scenario 1 is used to exemplify this condition. Initially, Non-Vital L1-1, Non-Vital L1-2, Vital L2-1, Vital L3-1, and Vital L3-2 are activated at T=10 s, yielding a total power consumption of approximately 131 kW. At T=25 s, the generator's rating is modified to furnish 80% of its maximum capacity (120 kW), creating a circumstance where the generator's reduced capacity falls short of the load demand. Consequently, the ALS controller must shed load to align with the diminished generator capabilities.

Given that L2-2 is inactive, the ALS controller opts to shed L1-2, followed by a partial shedding of L1-1, taking into consideration their lower weights. This strategy minimizes the number of deactivated loads while ensuring the generator remains below 95% of its newly diminished capacity of 120 kW. This scenario underscores the ALS controller's, as well as IDEAL's ability to execute a selective and partial load shedding process. Refer to Table the load weights and ratings, while Figure 49 presents the simulated outcomes achieved through the ALS controller.

Description	Rating (kW)	Original Weights	Adjusted Weights
Non-Vital Load 1-1	15	0.2	0.2
Non-Vital Load 1-2	15	0.2	0.2

Vital Load 2-1	32	1	1
Non-Vital Load 2-2	32	0.2	0.2
Vital Load 3-1	34	0.5	0.5
Vital Load 3-2	35	0.5	0.5
Mission Load (ML)	70	1	1
G1 M-G Set	150*		

*Overlimit is initially set by the controller to 142.5 kW, but it is arbitrarily adjusted by the user during the experiment to 115 kW to demonstrate a reduction of capacity.




Figure 49. The simulated results of the ALS controller for scenario 2. (© 2023 IEEE)

Following the validation of the intended functionality through simulation, the scenario was enacted on the IDEAL testbed. Between T=5 and T=15, L1-1, L1-2, L2-1, L3-1, and L3-2 were activated. At T=25, generator derating was implemented. Figure 50 presents the observed outcomes on the IDEAL testbed, while Figure 51 provides a record

of data captured by the PXI chassis throughout the scenario. It is noteworthy that the results attained from the IDEAL testbed align closely with the projected outcomes derived from simulation.




Figure 50. Experimental load values measured by the LabVIEW controller during scenario 2. (© 2023 IEEE)




Figure 51. Experimental data collected by the PXI chassis during the operation of ALS scenario 2. (© 2023 IEEE)

PHRR Scenario 3: Rep-Rate Ramped Mission Load

This scenario offers a demonstration of repeated utilization of the PHRR during a 10x repetitive transient load. Upon the initiation of the mission load by the user, a 70 kW pulse is commanded encompassing a 5-second interval. Initially, the battery serves as a buffer for the load, with a peak around 90 A, and the 1kV TDK power converter promptly escalates the power supplied from G1. This increment in TDK's power continues until it takes over the load or the mission load terminates, whichever transpires first. While the current setting in the PHRR is currently arbitrarily designated, the aim is to optimize rates and durations for diverse scenarios. Upon the disabling of the load, the TDK persists in sustaining the generator's output power for a brief period by recharging the battery.




Figure 52. Simulation of the PHRR controller operating in a repeated mission load scenario 3. (© 2023 IEEE)

Throughout this operational scenario, after roughly four transient load events, the battery-supplied current stabilizes around 25 A directed to the load, concurrently with the TDK 1 kV power supply reaching approximately 65 A. After the conclusion of the

repetitive engagements, the generator's power is upheld through battery recharge, progressively diminishing its output until the battery's SoC is recharged to 60%. This scenario underscores the PHRR's capacity to uphold a load and maintain steady power on the generator amid a repetitive transient load circumstance. This aids in minimizing total bus ripple on other sources and loads. **Error! Reference source not found.** elucidates the system simulation, while Figure 53 highlights the experimental data acquired from the hardware, including labels that highlight instances when the battery buffers the mission load, the generator supplies ramped support, and the gradual descent until the battery SoC is reinstated. While the user holds the ability to configure this value, a safety-oriented operational range of 60% was selected for the battery. Figure 54 provides visibility into the PXI chassis data.




Figure 53. Experimental load values measured by the LabVIEW controller while the PHRR controller operates on the IDEAL testbed during scenario 3. (© 2023 IEEE)

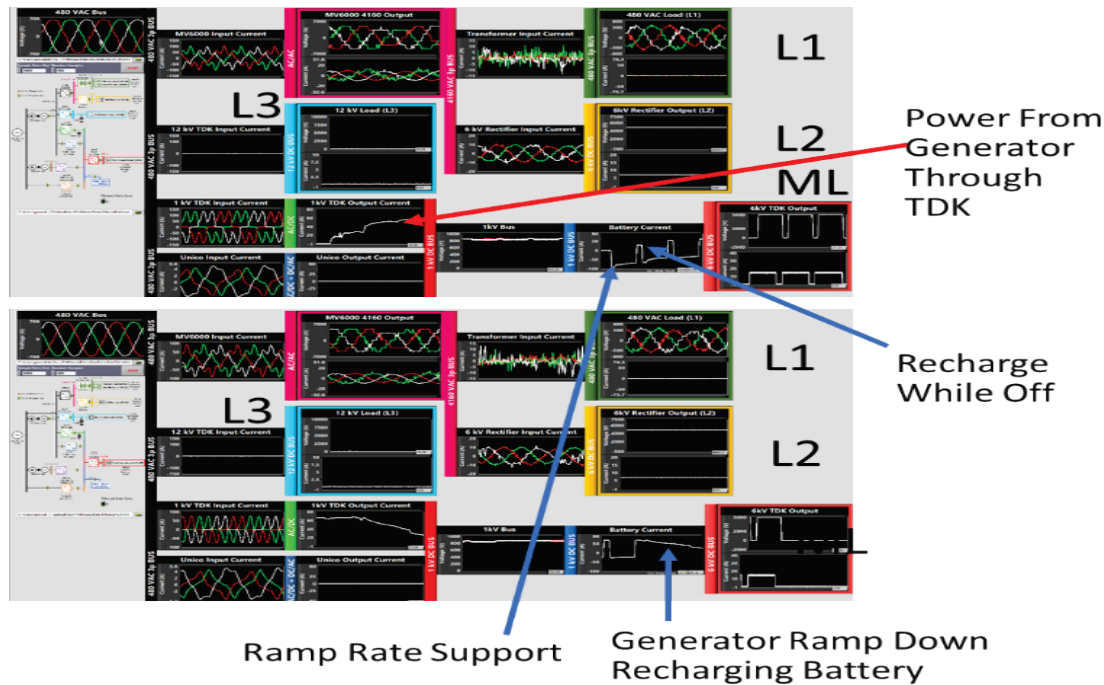


Figure 54. Experimental data collected by the PXI chassis during the operation of PHRR scenario 3. (© 2023 IEEE)

PHRR Scenario 4: Ramped generation overload

This scenario delves into the utilization of the PHRR whereby the battery functions as a buffer for the generator to accommodate the mission load. However, due to the battery's declining state, drawing power results in a drop in cell voltage below the safety threshold. Consequently, the PHRR controller promptly sources power from the generator/TDK, as the previously slower ramp rate is no longer feasible. Upon completion of the load, the PHRR undertakes the task of charging the battery at its maximum rated current, continuing until the average cell voltage of the battery ascends above the predetermined threshold or until the maximum cell voltage is attained. Subsequently, the system reverts to its normal operation. Figure 55 graphically illustrates the TDK and battery currents, alongside the power drawn by the mission load, while Figure 56 furnishes insight into the TDK current, as well as the minimum and average cell voltages. Both

figures provide a visual representation of the fluctuations in TDK and battery currents. To conclude, Figure 57 presents the PXI data garnered during this scenario's operation.




Figure 55. Experimental values of TDK, mission load, and battery current measured on IDEAL when battery is in an unsafe operating condition in scenario 4. (© 2023 IEEE)




Figure 56. Experimental values of battery voltage and TDK current measured on IDEAL when battery is in an unsafe operating condition in scenario 4. (© 2023 IEEE)

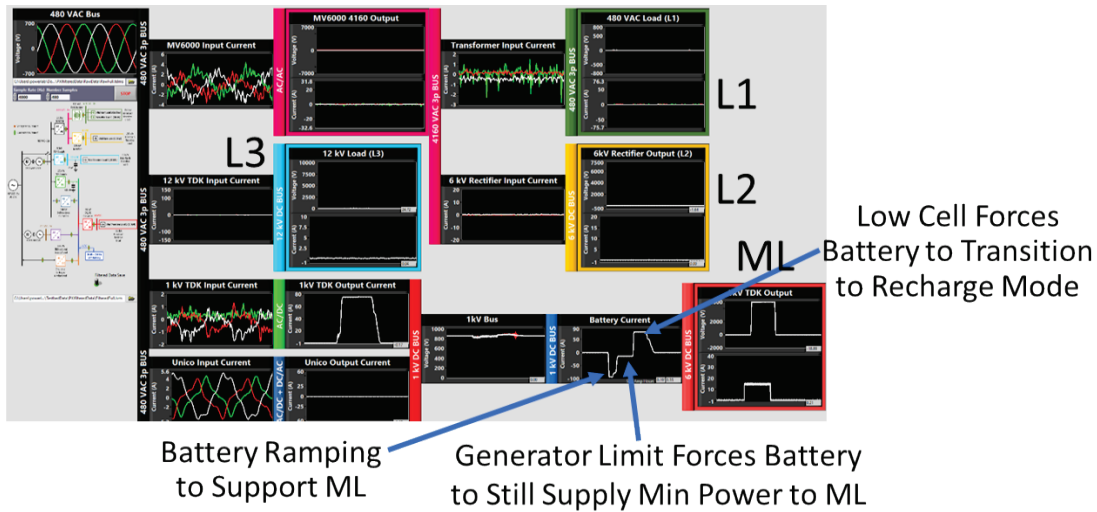


Figure 57. Experimental data collected by the PXI during scenario 4. (© 2023 IEEE)

ALS/PHRR Scenario 5: Ramped generation overload

This scenario encompasses the concurrent application of both the ALS and PHRR algorithms. A simulation depicting the unfolding of this scenario is encapsulated in Figure 58. Subsequently, Figure 59 demonstrates the empirical data derived from the executed experiment, while Figure 60 shows the PXI data interpretation gleaned from the same experiment.

Initially, a combination of vital and non-vital loads is activated, encompassing L1-1 (15 kW), L1-2 (15 kW), L2-1 (32 kW), L2-2 (32 kW), and L3-1 (34 kW), aggregating to an approximate total of 128 kW drawn from the generator. Upon initiating the Mission Load (ML), the battery initially supplies the load in its entirety, promptly followed by the intervention of the PHRR which gradually introduces generator support. As the generator's output escalates, it ultimately surpasses the threshold set by the ALS, triggering the necessity for load shedding. In this specific instance, the load with the lowest priority, L2-2, is shed. The load priorities mirror those detailed in the "Original Weight" column of Table 1. All the previously enumerated loads are simultaneously activated at the outset of

the scenario. At approximately $t = 30$ s, the initiation of the ML incites a swift influx of battery current, evident by the negative power value indicating current sourced from the battery. As the generator output surpasses the 143 kW threshold, the ALS intervenes to shed load L2-2, while preserving the supply to all other vital loads. It is pertinent to acknowledge that the battery could have exclusively powered the ML without necessitating load shedding; however, this course of action would fail to effectively demonstrate the ALS's capability. Subsequently, the battery current experiences a gradual decline while the generator power escalates, as evidenced by the gradual rise in power sourced from the TDK 1kV supply. At $t = 33$ s, the ML is deactivated, prompting a reversal in battery current direction, denoting the commencement of the recharging process. This downward trend in power persists until approximately $t = 80$ s.

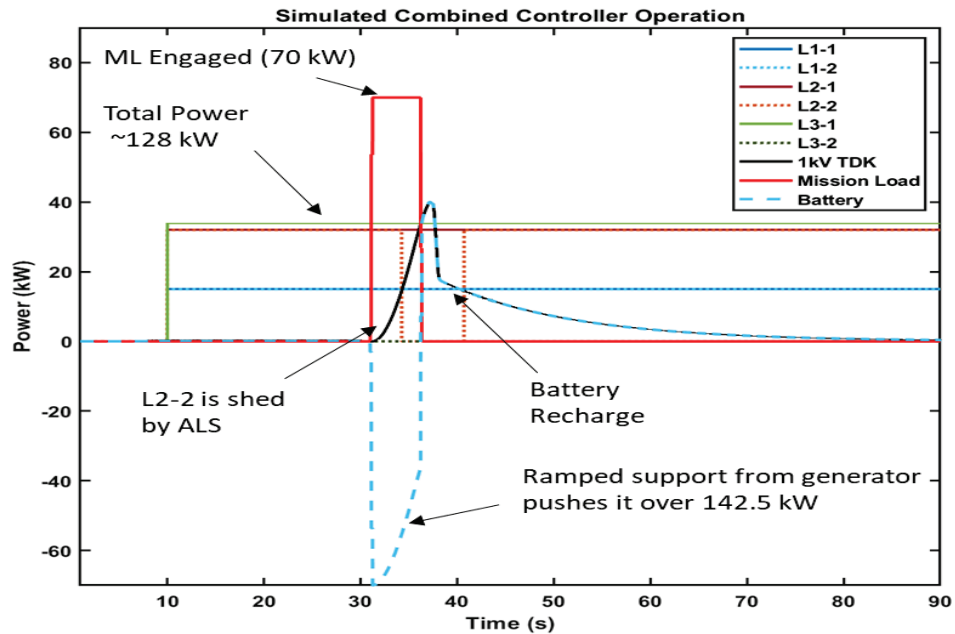


Figure 58. The simulated results of running the PHRR controller and ALS controller for scenario 5. (© 2023 IEEE)




Figure 59. Experimental load values measured by the LabVIEW controller during the combined scenario 5. (© 2023 IEEE)

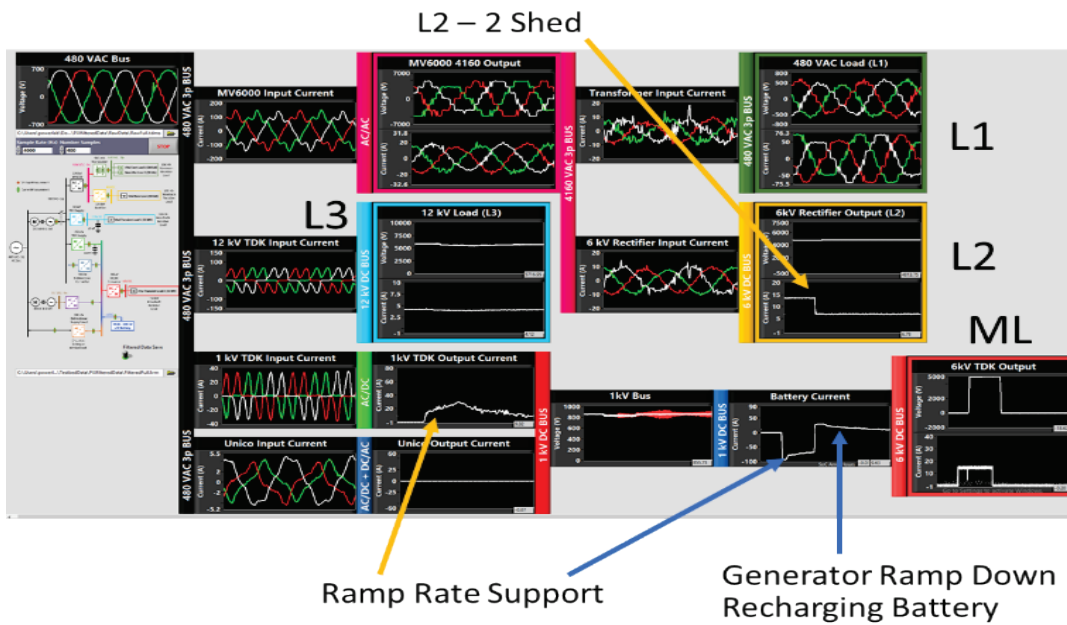


Figure 60. Experimental data collected by the PXI chassis during the operation of scenario 5. (© 2023 IEEE)

CONCLUSSION

In this paper, a meticulously designed and executed BMS testbed has been established, serving as a pivotal platform for evaluating the performance and efficacy of various BMS schemes. The testbed integrates advanced HIL models onto cell emulators, enabling real-world emulation of complex battery systems under a multitude of scenarios. A significant aspect of this endeavor involved addressing the intricate challenges posed by cell emulation, communication interfaces, and system dynamics. In addition to this, the variety of controls and frameworks each BMS has is a point concern for validation. The testbed is pivotal in validating BMS performance before integration into shipboard control systems where data and performance of the battery and BMS system is crucial.

Furthermore, a co-simulation effort with FSU and the USC further solidified the versatility and robustness of the IDEAL testbed. By combining the expertise and resources of different institutions, the study successfully demonstrated the collaboration of distinct algorithms into a cohesive and efficient system. This co-simulation not only emphasized the potential of the IDEAL testbed but also showcased its adaptability in accommodating various control methodologies. A primary example of this was the ramp rate implemented at UTA using PHIL and with it, the noticeable improvement of power quality, thus demonstrating the importance and advantages of these control algorithms.

After verification of the BMS using the BMS testbed, as well as deployment of control algorithms onto the IDEAL testbed that demonstrate the power quality improvement available by implementing battery storage, controls developed by Clarkson were integrated onto IDEAL utilizing the real 1kV battery. The scenarios executed on the IDEAL testbed underscore the capabilities of CU's ALS and PHRR algorithms. Through

a series of designed experiments, the study adeptly demonstrated the adaptability of these algorithms in load shedding and responding to transient load events, and complex scenarios where both algorithms operated simultaneously. The results not only validated the theoretical simulated underpinnings, but also provided empirical evidence of their performance onto emulated hardware.

In conclusion, the culmination of extensive efforts in designing and implementing the BMS testbed, co-simulating with external partners, and advancing the IDEAL testbed has yielded a comprehensive platform that serves as a pioneering milestone in shipboard control development. This integrated framework not only enables the exploration of innovative control algorithms but also empowers further development of advanced energy storage solutions, contributing significantly to the testing of BMSs and their integration into modern shipboard power systems.

References

- [1] J. F. Hansen and F. Wendt, "History and State of the Art in Commercial Electric Ship Propulsion, Integrated Power Systems, and Future Trends," *Proceedings of the IEEE*, vol. 103, no. 12, pp. 2229-2242, 2015.
- [2] G. Ortiz, J. Biela, D. Bortis and J. W. Kolar, "1 Megawatt, 20 kHz, isolated, bidirectional 12kV to 1.2kV DC-DC converter for renewable energy applications," in *The 2010 International Power Electronics Conference - ECCE ASIA -*, 2010.
- [3] ONR, Electric Ship Research Development Consortium , September 2021. [Online]. Available: <https://www.esrdc.com/>.
- [4] Z. Ding, S. K. Srivastava, D. A. Cartes and S. Suryanarayanan, "Dynamic Simulation Based Analysis of a New Load Shedding Scheme for a Notional Destroyer Class Shipboard Power System," *IEEE Electric Ship Technologies Symposium*, pp. 95-102, 2007.
- [5] Z. Ding, S. Srivastava and D. Cartes, "Expert System based Dynamic Load Shedding Scheme for Shipboard Power Systems," in *IEEE Industry Applications Conference Forty-First IAS Annual Meeting*, Tampa, 2006.
- [6] C. Tschritter, A. Johnston, L. Vu, T. Nguyen, D. Wetz, T. Vu, J. Langston, H. Ravindra, M. Stanovich, K. Schoder, M. Steurer, C. Schegan and J. Heinzl, "Robust Control of a Medium Voltage AC/DC Testbed," November 2022.
- [7] D. A. Dodson, D. A. Wetz, J. L. Sanchez, C. Gnegy-Davidson, M. J. Martin, B. Adams, A. Johnston, J. Heinzl, S. Cummings, N. Frank, N. A. A. Rahim and M. Davis, "Design and Evaluation of a 1000 V Lithium-Ion Battery," *Naval Engineers Journal*, vol. Volume 131, pp. 107-119, September 2019.
- [8] T. Moloughney, "2021 Lucid Air Dream Edition Revealed: Range, Pricing, Specs, And More," *Inside EVs*, September 2020.

- [9] D. A. Wetz., J. Martin, C. L. Williams, D. A. Dodson and J. M. Heinzl, "Design of Two 1000 V Batteries for Evaluation as Pulsed Power Prime Power Supplies," 2016.
- [10] L. Buccolini, S. Orcioni, S. Longhi and M. Conti, "Cell Battery Emulator for Hardware-in-the-Loop BMS Test," in *2018 IEEE International Conference on Environment and Electrical Engineering and 2018 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe)*, 2018.
- [11] A. N. Johnston, Z. R. Bailey, D. A. Wetz, G. K. Turner and J. M. Heinzl, "Design and Commissioning of a Medium Voltage Testbed with Emulated Pulsed Loads," *Naval Engineers Journal*, vol. 134, no. 3, pp. 129-144, 2022.
- [12] A. N. Johnston, G. K. Turner, D. A. Wetz, Z. R. Bailey, C. Schegan, J. M. Heinzl and M. Giuliano, "Emulation of a Single Zone of an Integrated Power System," in *American Society of Naval Engineers Intelligent Ships Symposium*, Virtual, 2021.
- [13] A. N. Johnston, D. A. Wetz, Z. Bailey, B. J. McRee, D. A. Dodson and J. M. Heinzl, "A Medium Voltage, Distributed Power Generation Testbed Deploying Transient Loads," in *International Ship Control Systems Symposium*, Delft, 2020.
- [14] A. Johnston, D. Wetz, Ramtin, Madani, A. Davoudi, G. Turner, D. Dodson, B. McRee, D. Pullaguram, Z. Bailey, J. Heinzl, M. Giuliano and C. Schegan, "Mitigating Transient Loads in Medium-Voltage Direct Current Microgrids," in *Advanced Machinery Technology Symposium*, 2020.
- [15] J. Meng, Y. Wang, C. Wang and H. Wang, "Design and implementation of hardware-in-the-loop simulation system for testing control and operation of DC microgrid with multiple distributed generation units," *IET Generation, Transmission & Distribution*, vol. 11, pp. 3065-3072, 2017.

- [16] M. M. Islam, "Front Matter," in *Shipboard Power Systems Design and Verification Fundamentals*, 2018, pp. i-xxiii.
- [17] S. G. Jayasinghe, L. Meegahapola, N. Fernando, Z. Jin and J. M. Guerrero, "Review of Ship Microgrids: System Architectures, Storage Technologies and Power Quality Aspects," *Inventions*, vol. 2, 2017.
- [18] C. A. De La O, M. Difronzo, M. Milton, H. L. Ginn, M. Stanovich, J. Park, K. Schoder, M. Steurer, D. Wetz, C. Tschritter, G. Turner and A. Johnston, "Remote Site Real-Time Co-Simulation of a Shipboard Microgrid," in *2022 Open Source Modelling and Simulation of Energy Systems (OSMSES)*, 2022.
- [19] Villas-Node, "<https://fein-aachen.org/en/projects/villas-node/>," 2021.
- [20] A. Monti, M. Stevic, S. Vogel, R. W. De Doncker, E. Bompard, A. Estebarsari, F. Profumo, R. Hovsapian, M. Mohanpurkar, J. D. Flicker, V. Gevorgian, S. Suryanarayanan, A. K. Srivastava and A. Benigni, "A Global Real-Time Superlab: Enabling High Penetration of Power Electronics in the Electric Grid," *IEEE Power Electronics Magazine*, vol. 5, pp. 35-44, 2018.
- [21] V. P. N. Tinc, "<https://www.tinc-vpn.org/documentation/Introduction.html>," 2021.
- [22] B. L. H. N. e. al., "Advanced Load Shedding for Integrated Power and Energy Systems," in *IEEE Electric Ship Technologies Symposium*, Virtual, 2021.
- [23] T. V. Vu, D. Gonsoulin, F. Diaz, C. S. Edrington and T. El-Mezyani, "Predictive Control for Energy Management in Ship Power Systems Under High-Power Ramp Rate Loads," *IEEE Transactions on Energy Conversion*, vol. 32, no. 2, pp. 788-797, 2017.
- [24] C. Tschritter, A. Johnson, David Wetz, L. Vu, T. Vu, T. Nguyen, J. Langston, H. Ravindra, K. Schoder, M. Stanovich, M. Steurer, Christian Schegan and J. Heinzl, "Advanced Load Shed and Predictive Ramp Rate Control of a Medium Voltage AC/DC Testbed," *Journal of Marine Engineering & Technology*, no. Publication Pending, 2024.

- [25] C. D. Tschritter, D. A. Wetz, G. K. Turner and J. M. Heinzl, "Battery Management System (BMS) Test Stand Utilizing a Hardware-in-the-Loop (HIL) Emulated Battery," in *2021 IEEE Electric Ship Technologies Symposium (ESTS)*, 2021.