

High-Performance Architectures for Vehicle Propulsion: An Unconventional Approach to Design, Fabrication and Analysis using Scalable Flying Capacitor Multilevel Converter Modules

Nathan Pallo



Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2021-189

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2021/EECS-2021-189.html>

August 13, 2021

Copyright © 2021, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

Much of this work was both motivated and funded by the Advanced Air Transport Technology (AATT) Project at the National Aeronautics and Space Administration (NASA) through cooperative agreement NNX14AL79A. Additional support has been provided by the Vehicle Technologies Office in the United States Department of Energy through the Electric Drive Technologies Consortium (EDTC) under Award Number DEEE0008712.

High-Performance Architectures for Vehicle Propulsion:
An Unconventional Approach to Design, Fabrication and Analysis using
Scalable Flying Capacitor Multilevel Converter Modules

by

Nathan Pallo

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Robert Pilawa-Podgurski, Chair

Professor Seth Sanders

Professor Duncan Callaway

Professor Alejandro Dominguez-Garcia

Summer 2021

High-Performance Architectures for Vehicle Propulsion:
An Unconventional Approach to Design, Fabrication and Analysis using
Scalable Flying Capacitor Multilevel Converter Modules

Copyright 2021
by
Nathan Pallo

Abstract

High-Performance Architectures for Vehicle Propulsion:
An Unconventional Approach to Design, Fabrication and Analysis using
Scalable Flying Capacitor Multilevel Converter Modules

by

Nathan Pallo

Doctor of Philosophy in Engineering – Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Robert Pilawa-Podgurski, Chair

The consequences of climate change grow more pronounced each year as carbon emissions continue unabated. Decarbonizing all aspects of the transportation sector is a major requirement for a sustainable pathway to mitigate these consequences. However, this requires overcoming substantial engineering challenges, not the least of which are the power density and efficiency requirements of the inverter system in future electric drivetrains. This work examines a high-performance and scalable approach for modeling, designing, fabricating and testing an architecture that meets aggressive industry targets through innovative techniques and the use of an unconventional topology. Several prominent and promising topologies are reviewed, and an unconventional approach using the flying-capacitor multilevel topology is introduced to drive the order-of-magnitude improvements required by the industry targets. A 10-level, 1 kV, dual-interleaved converter module serves as the platform for an in-depth study of the electrical and thermal design and performance attained using this approach, where a peak power of 18.9 kW and a peak efficiency of 98.95% is experimentally demonstrated. Scalability and reliability of the proposed architecture is discussed and tested using individual modules as well as an array paralleling nine of these 38.4 kW/kg, 24.4 kW/L modules. Finally, potential extensions to the work in future studies are reviewed.

To my husband, Xin. He will likely never read this document, nor particularly care about any of the implications of this work. However, he was instrumental in the support, nourishment and encouragement required to complete the efforts within.

Also, to my family and friends. My family has been an anchor in all of the years of my adulthood, while my found family has provided additional enrichment and encouragement beyond the borders of my hometown.

Contents

Contents	ii
List of Figures	iv
List of Tables	x
1 Vehicle Drivetrain Electrification	1
1.1 Constraints and Performance Criteria	3
1.2 Survey of Topologies	6
1.3 An Unconventional Approach	9
2 Inverter Module Electrical Design	13
2.1 Structure for Computational Analysis	18
2.2 Loss Modeling	19
2.3 Startup Dynamics	32
2.4 Mitigation of Commutation Loop Inductance	33
2.5 Experimental Results	37
2.6 Chapter Remarks	45
3 Thermal Management and Mechanical Integration	47
3.1 Cooling Technologies and Media	47
3.2 Modular Heat Sink Design	48
3.3 Analysis using Computational Fluid Dynamics	52
3.4 Experimental Results	56
3.5 Chapter Remarks	63
4 Inverter Array Design and Operation	65
4.1 Hierarchical Control	66
4.2 Multi-Module Harmonic Content Analysis	70
4.3 Current Sharing & Parameter Sensitivity	73
4.4 Experimental Results	78
4.5 Chapter Remarks	84

5	System Reliability	85
5.1	Short-Circuit Fault Mitigation	88
5.2	Cell-Level Component Failures and Redundancy	104
5.3	Parallel Module Redundancy	107
5.4	Chapter Remarks	118
6	Future Work & Conclusion	121
6.1	Extended Component Characterization	121
6.2	Advanced Gate-Drive Techniques	121
6.3	Validation of Sense and Measurement Schemes	123
6.4	Online Health-Monitoring	124
6.5	Folded & Three-Dimensional Module Packaging	125
6.6	Improved Cooling and Integration	127
6.7	Coupled Inductors for Interleaved Modules	127
6.8	Modeling of Converter Dynamics	128
6.9	Conclusion	129
	Bibliography	130

List of Figures

1.1	Depictions of more electric aircraft propulsion leveraging boundary layer ingestion, as adapted from concepts and figures in [12] and [13]. For illustrative purposes, contributions to drag from the wings of the aircraft are assumed to be less than that of the fuselage and are ignored.	2
1.2	The common, 2-level bridge topology, with corresponding single-phase output waveform.	6
1.3	The 3-level topologies highlighted in this chapter, with corresponding output waveform.	7
1.4	The schematic above shows major power-processing components in a 4-level FCML (omitting input and output filters). Note that the nominal flying capacitor voltages show how device stress is reduced to $V_{dc}/(N - 1)$. The waveforms at right show how PSPWM creates a high effective output frequency at the switched-node, v_{ac} – a factor of $(N - 1)$ greater than the switching frequency. .	10
2.1	Fabrication, installation and high voltage testing of a custom busbar for C ₈ . . .	15
2.2	The schematic of the 10-level, dual-interleaved FCML inverter module of this work, accompanied by the top and side views of a hardware prototype. The converter is shown partially assembled to illustrate the mechanical assembly and thermal management, with the exception that screws used for heat sink mounting are were replaced with an improved design in Chapter 3.	17
2.3	Sampled converter waveforms used to simulate and model FCML losses.	20
2.4	Typical nonlinear Q_{oss} characteristics for an EPC2034C gallium nitride transistor.	23
2.5	Typical nonlinear E_{oss} characteristics for an EPC2034C gallium nitride transistor.	24
2.6	Typical reverse conduction characteristics for an EPC2034C gallium nitride transistor.	25
2.7	Turn-on and turn-off waveforms.	26
2.8	Comparison of loss models without (gray series) and with (blue series) an iterative solver that compensates for converter losses in the output power and efficiency calculations.	30
2.9	Modeling accuracy validated across various operating points.	31
2.10	Startup simulation in PLECS.	32
2.11	Commutation loop between a given pair of complementary switches in a single 9-level FCML inverter, shown here for a 9-level converter.	33

2.12	Cross section showing commutation loop between a given pair of complementary switches. Internal layers and vias create an electrically-thin vertical layout on a single PCB; the cross-sectional area inside the PCB (shaded red), and thus the corresponding loop inductance, can be greatly reduced with this approach. . . .	34
2.13	3D model of a switching cell with buried vias, as imported and meshed in Q3D.	35
2.14	Spice model of commutation loop implemented as a component in LTSpice using parameters extracted from Q3D finite-element analysis.	36
2.15	Experimental characterization of drain-source voltage overshoot during commutation as a means of comparing loop parasitics between designs.	37
2.16	Impedance analyzer and probe used to measure the commutation loop inductance. After calibration is carried out to the adapter tip endpoint (inset), the loop impedance between the drain to source pads can be measured with high accuracy..	38
2.17	Experimental setup for 1 kV dc bus startup on a 10-level FCML inverter module. Waveforms were captured on three synchronized Keysight MSO-X 4024A oscilloscopes with several models of high voltage differential probes of 25 MHz bandwidth or greater.	39
2.18	Startup of the dc bus from 0 to 1 kV, after which output power is ramped to 5 kW. Waveforms were captured on three synchronized Keysight MSO-X 4024A oscilloscopes with several models of high voltage differential probes of 25 MHz bandwidth or greater. High-frequency content present on a given waveform can be attributed to the specific probe used, or the higher fidelity of ac content capture for the signals with lower dc voltage.	40
2.19	Experimental setup for high power module tests.	41
2.20	10-level converter performance at 1kV dc bus with varied switching frequency. .	42
2.21	Experimentally-measured interleaved output current waveforms at 16.7 kW output power. Waveforms are captured digitally on an oscilloscope and plot to enhance details.	43
2.22	Interleaved output voltage waveforms at 1 kV dc bus and 16.7 kW output power. Waveforms are captured digitally on an oscilloscope and plot to enhance details, but are shown next to the oscilloscope screen capture for additional context. . .	44
2.23	Summary representation of how the design of this work meets aggressive performance targets for more electric aircraft compared to the state-of-the-art. . . .	46
3.1	Comparison of previous and proposed modular heat sink mounting strategies. Note, the threaded insert in the right figure is only conceptual; the interference fit described later in this section made this specific consideration obsolete. . . .	49
3.2	Examples of the snap-fit concept with spring alignment and interference features.	50
3.3	Installation of heat sinks on an entire face of a module with dual-sided cooling. .	51
3.4	Illustration of dual-inlet cooling of an FCML-based power module. One inlet is provided for each interleaved phase, with exhaust air designed to exit out the middle of each module.	52

3.5	Meshing and results of CFD for a single heat sink. A finer mesh is used in the region around the heat sink. Fin width, fin pitch and fin count were parametrically varied and the system was systematically re-meshed and solved to find the best configuration.	53
3.6	Results of the parametric sweep of heat sink fin count, pitch and width.	54
3.7	Composite rendering of air particle tracks from the quarter-symmetric fluid simulation.	55
3.8	Composite image of the inlet at one side of the module, where the color map indicates static pressure at the inlet of each individual heat sink. Up to a 14% variation in pressures yields commensurate variation in mass flow, and therefore cooling capacity.	56
3.9	Miniature thermocouple jig that enable consistent fixturing for chip-scale GaN devices.	57
3.10	Schematic for connection of thermal surrogate. The gate of each device is shorted to the respective source, and a current is driven through the string of devices to provide a thermal load through the reverse conduction losses. The string current is measured with ammeter A at the input terminal, while individual source-drain voltages are measured across each device. Note, the “diodes” are only symbolic of the reverse conduction and do not exist in the physical system.	58
3.11	Common- and differential-mode ferrites installed on the temperature sensors. . .	59
3.12	Custom test fixtures, instrumentation and scripted acquisition equipment for measuring thermal performance of a complete inverter module. Variable speed blowers provide cooling air, drawing air through a custom mass flow meter. The module (inset) is wired with thermocouples and drain-source voltage monitors per the test procedure described in this section.	61
3.13	Simulated and experimentally measured thermal performance across varied cooling air inlet mass flow. The simulation does not model TIM resistance; a constant impedance value for the TC5022 TIM from Table 2 of [155] is added to the simulated values to yield the dashed lines.	63
4.1	Schematic representation of how M interleaved 10-level FCML inverters can be interleaved within a phase-leg.	65
4.2	High-level illustration of the hierarchical control structure, configured here to drive a three-phase load.	66
4.3	Illustration of a time-base synchronization scheme common in microcontroller PWM peripherals.	68
4.4	Illustration of a synchronization scheme using a communication link and ADPLL.	69
4.5	Calculated output current spectra in the region of interest (i.e., harmonics of f_{eff} for $N = 10$). The unit circles are normalized to the switching period, where arrows indicate the phase shift of interleaved modules. Green ticks mark each full cycle at the effective output frequency per switching period.	71

4.6	Calculated input current spectra in the region of interest (i.e., harmonics of f_{sw}) for $N = 10$	72
4.7	Calculated input current harmonic content using double Fourier analysis and assuming sinusoidal output current.	73
4.8	PLECS Blockset schematic of a dual-interleaved, 10-level FCML for Monte Carlo simulations. The two interleaved converter phases are the modules labeled Phase A and Phase B and contain the switches and flying capacitors whose parameters are varied throughout the simulation. Measurements of interest (pertaining to current sharing) are highlighted: the two respective phase currents, as well as their respective common- and differential-modes. The RMS values of these measurements, as well as the average of output power into Rload , are computed in a top-level Simulink simulation.	74
4.9	Correlation between key converter measurements from the Monte Carlo simulation.	76
4.10	Statistical results of simulations recording current sharing imbalance between phases of a dual-interleaved converter module given parameter variation.	77
4.11	Electromagnetic compatibility pre-compliance experimental setup for measuring CISPR-25 conducted emissions.	79
4.12	Comparison of common-mode emissions measured at the HV LISNs with and without a common mode choke installed on the connection between the ILM array and the LV LISN.	80
4.13	Hardware photographs for the dual-interleaved module (ILM, on left) and three-phase, 18 converter array (right).	81
4.14	Six v_{ds} waveforms of the S_{9B} low-side switches (closest to the DC bus) measured for the converters of Phase A illustrate that 6-fold interleaving is achieved at the input.	82
4.15	Experimental waveforms at 400 V DC and 6 kW output power for the 3-phase array comprised of nine dual-interleaved FCML modules (16 arrayed converters).	82
4.16	Common-mode and differential-mode conducted emissions measurements of the three-phase segmented drive operating at a 400 V DC bus and 6 kW output power, measured at the HV terminals of the test setup.	83
5.1	The bathtub curve.	85
5.2	Panelized and contract assembled array of switching cell daughter boards.	86
5.3	Panelized converter module circuit boards with features (routing and breakaway tabs) to mitigate mechanical stress from post-assembly depaneling.	87
5.4	Single-leg representation of an N -level flying-capacitor multilevel converter with a split dc bus. The output filter capacitor is connected back to the dc-link midpoint. A short-circuit occurs with a sudden connection of fault impedance, Z_F	89

5.5	Switching waveforms of a 5-level ($N=5$) inverter. v_x is plotted over a fundamental period T_1 , while the zoomed view illustrates the high effective frequency and gate signals obtained using PSPWM. The bottom plot illustrates the charge and discharge of a flying capacitor over each switching period, with a short-circuit fault introduced between between T_{sw} and $2T_{sw}$	90
5.6	Analog detection of short-circuit fault from the mismatch between estimated output voltage and actual output voltage.	94
5.7	PLECS simulation of the model-predictive approach for the 5-level converter from [60] at power factor of 0.8. The detection time, t_{FD} , is the model time-step, T_s – after which mitigating, action (described in Section 5.1.3) is taken. For comparison, the dashed line shows the trajectory of the inductor current, i_L , without mitigation.	96
5.8	Representative time-reversed, complementary switching strategy to reset inductor fault current for $N=5$	97
5.9	Simulation of a mitigation strategy that resets the inductor current via reverse capacitor charge/discharge actions.	99
5.10	Experimental setup, including the load, the switch to simulate faults and an air cooling system. The inset highlights the converter hardware from [1] used in the tests.	101
5.11	Experimental results for the hardware demonstration of fault detection and ride-through at 1.1 kW output power.	102
5.12	Zoomed experimental waveforms showing short-circuit transient for the two mitigation actions for a 6-level FCML.	103
5.13	Switching cell reliability structure for analyzing cell-level failures and redundancy.	104
5.14	Simple Markov model illustrating three phases (A,B and C) of six paralleled modules ($P = 6$).	107
5.15	Markov reliability model for a three-phase array for $P = 6$ and $F = 2$, where failures in a single phase-leg have complex relationship with the overall system reliability. State 1 is the <i>healthy</i> state and state 11 is the <i>failed</i> state, with the remaining states indicating some degree of system degradation.	109
5.16	Markov reliability model for a phase-leg implementing active redundancy.	110
5.17	Markov reliability model for a phase-leg implementing cold standby.	111
5.18	Markov reliability model for a repairable phase-leg implementing active redundancy.	111
5.19	Markov reliability model for a repairable phase-leg implementing cold standby.	112
5.20	State probabilities over time using (5.21) and Γ from (5.17).	113
5.21	A phase-leg of the three-phase, 3-level ANPC for illustrating switch count per module.	115
5.22	Module series reliability structure, dominated by the switching cell failure rate.	116
5.23	Reliability and power density of a 1 MW array with various levels of redundancy.	117
5.24	Temperature-dependent failure rate coefficient, π_T , for the high-k ceramic capacitors considered in this work.	119

6.1	Highlight of the large area occupied by the gate drive circuitry in this work. . .	122
6.2	Instrumentation amplifier used to capture the high-frequency and high-dynamic-range waveform at the switched-node of the FCML. Gain can be toggled between 1 and 2 to enhance resolution near zero-crossings or at low modulation.	123
6.3	Current amplifier for the LMT01 temperature sensors on the modules in this work.	125
6.4	A conceptual illustration of how layouts that implement remote flying capacitance can improve the routing of switching currents in a 9-level FCML.	126
6.5	Conceptual mock-up of how the converter array of this work might be scaled and integrated with the motor of [227]. In the rendering, a liquid cooling system and cold-plate is used to cool the module versus the individual, air-cooled heat sinks presented in Chapter 3.	127
6.6	A root locus plot for the input-to-state transfer function for one of the flying capacitors of the 10-level converter of this work, as the switching frequency is swept from 20 kHz to 115 kHz. The arrows indicate the direction of the pole and zero movement as frequency increases, while the colors indicate pole-zero pairs. .	128

List of Tables

1.1	NASA key performance parameters for an electric drivetrain [27].	4
1.2	Key 2025 targets in the US Drive 2025 electric powertrain roadmap [32].	5
1.3	Key benefits of the FCML topology [57].	10
2.1	Capacitor selections per interleaved phase.	14
2.2	Converter bill of materials.	15
2.3	Converter mass breakdown.	16
2.4	Experimentally measured commutation loops for various switching cells.	38
2.5	Key performance metrics for the design of this work.	45
3.1	Selection of appropriately sized miniature springs from Acssess Spring.	51
3.2	Components used in the orifice meter for measuring mass flow.	60
5.1	Table of key parameters and critical fault clearing times for recently published converters.	92

Acknowledgments

Much of this work was both motivated and funded by the Advanced Air Transport Technology (AATT) Project at the National Aeronautics and Space Administration (NASA) through cooperative agreement NNX14AL79A. Particular thanks are due to Andrew Provenza of NASA Glenn Research Center for providing guidance and feedback during monthly program reviews, and support through the entire length of the program – especially during a major change in research campuses. Additional support and constructive conversations were provided by Ralph Jansen, Dr. Cheryl Bowman and Dr. Nateri Madavan, and their generosity in sharing their time is greatly appreciated.

Additional support has been provided by the Vehicle Technologies Office in the United States Department of Energy through the Electric Drive Technologies Consortium (EDTC) under Award Number DEEE0008712. Advice and guidance from Dr. Vipin Gupta was particularly appreciated as the research in this work began to transition from a proof-of-concept towards hardware that would be ready for integration in demanding automotive environments.

Several members of the Pilawa Group deserve special mention for their direct contributions to the NASA project or the EDTC project that supported this work: Pourya Assem for his work in designing the control cards used in this work, writing the initial versions of code for each, and for his insight in analog and digital circuits; Chris Barth for trailblazing the first implementations of high-level-count, high-voltage, flying capacitor multilevel converters and inverter array architectures (as well as spinning a machine with an early build of the hardware in this work); Thomas Foulkes for leading an effort to characterize dynamic on-resistance (a significant loss mechanism noted in Chapter 2) and fostering collaborations; Samantha Coday for assisting in all aspects of converter module design and test; Joseph Schaadt for his contributions to the thermal design, fabrication and eventual experimental measurements (with fruits of these labors in Chapter 3); Avinash Jois for continuing to advance the control architecture and occasionally getting roped into assembly tasks; Mads Taul for his work on short-circuit fault analysis and mitigation (much of which was successfully tested for Chapter 5); Amanda Jackson for learning so much in such a short amount of time to help design a system that speeds up testing of converters like those in this work; Logan Horowitz for taking over all of the low-inductance commutation loop design, analysis and measurement that has only just begun; and Roderick Bayliss, who joined the group just in time to help with a myriad of conducted emissions tests (presented in Chapter 4) and inherit a pile of mostly working converters. Dr. Tomas Modeer, who as post-doc with the group developed the first precursor to the module and methods of Chapter 2, also provided continuing technical guidance, encouragement and journal edits over the course of this work.

Of course, in addition to the technical support, all of the group members above – as well as Kelly, Rose, Maggie, Rahul, Nathan B., Nathan E., Ting, Francesca, Derek, Zitao, Zichao, Joseph L., Enver, Andrew and Pei – provided a wonderful social circle for hikes, barbecues, and happy hours. Derek and Zitao were especially valuable technical, professional and personal sounding boards as part of the same (if not just a bit earlier) graduating cohort.

Much credit is also due to the professors serving on the committee of this work – beyond the time they volunteered so generously. Teaching power electronics as a Graduate Student Instructor for Prof. Sanders provided valuable experience and perspective with previously familiar concepts, while his comments in both research meetings and graduate seminars stimulated interest in areas that were less familiar. Taking Prof. Callaway’s low-inertia power systems seminar provided valuable insight into both low-level and systems-level challenges – and opened a door to working with his students for a summer; Jose, Ciaran and Rodrigo shared both their technical understanding and modeling expertise, and taught so much through example on how to architect and manage a software development project for research. Additionally, the coursework, review of prior work, and one-on-one meetings with Prof. Dominguez-Garcia provided immensely useful experience with both reliability concepts and modeling dynamical and stochastic systems in general. This informed much of Chapters 5 and 6, as well as much of the work to follow in the professional career of the author.

Finally, Prof. Pilawa-Podgurski: for carefully reading every digest and manuscript, then providing feedback to improve precision and clarity of technical language, add motivation and background that appeals to all readers, and catch a typo here and there; for ensuring funding was never an issue, and that the lab was always well equipped; for providing personal support and understanding when research, grad school, or just life seemed to be going pretty poorly; for organizing or joining lunches, power hours or other social events; for allowing a mutual trust and frankness all discussions, both professional and personal; for the first introductory power electronics lecture on a blackboard in LEES more than a decade ago; for giving an opportunity to a student who did not have the best undergraduate academic record but was willing to make up for it with a lot of hard work in grad school. Thank you.

Chapter 1


Vehicle Drivetrain Electrification

Portions of this chapter are adapted in part from [1] and [2].

In his 2018 book on the future of electric vehicles, *Three Revolutions*, Professor Daniel Sperling notes that, although electric vehicles have been in around since the 1830s, they were not always the “objects of desire” that Tesla, Rivian and other breakthrough manufacturers have made them today [3]. Indeed, after many early pitfalls and the limited range of debut models, it was only through decades of advancement in technologies – primarily in battery energy storage, but also power electronics – that the tide has turned toward mass vehicle electrification. While a complex mix of technological challenges, policy decisions and a fickle consumer market seemed to have “killed the electric car” as recently as 2006 [4], new marketing campaigns by General Motors lauding thirty all-electric vehicle models by 2025 [5] show industry inertia is far greater now than when the single EV1 model was announced in the mid-1990s. At the same time, major transportation network companies like Lyft are announcing goals of migrating their fleets to zero-emission vehicles as early as 2030 [6]. Market moves like these signal a paradigm shift on a drastic scale and aggressive timeline.

Yet, it is possible this shift cannot come soon enough. As of 2017 estimates, transportation accounted for 22% of the global carbon dioxide emissions in the energy sector [7]. Since then, it is estimated that the slow pace taken to address this crisis means global warming is highly likely to reach *at least* 1.5 °C between 2030 and 2052 [8]. As such, an aggressive draw-down in emissions from fossil fuel use in transportation of nearly 90% is required between 2019 and 2070 to reach a sustainable development scenario [9]. While Sperling notes that mass vehicle electrification is only part of the solution, steps taken by fleets, municipalities and motorists can help decarbonize the vast spread of medium- to heavy-duty vehicles to passenger cars alike.

Although electrification of cars, trucks and other road and off-road vehicles has at least become a much more tangible prospect, aviation remains a subset of the transportation sector that is still difficult to decarbonize. Over the past two decades, agencies such as the National Aeronautics and Space Administration (NASA) in the United States and the European Commission have driven research to support more fuel-efficient air travel [10, 11].



(a) Illustration of a traditional commercial aircraft design, with two large combustion turbines providing thrust. The air stream that flows along the boundary layer of the aircraft fuselage is slowed compared to the freestream, causing drag and inefficient flight [13].

(b) Illustration of a turboelectric aircraft based on the NASA STARC-ABL concept [12]. Two large combustion turbines provide thrust as well as electrical power to drive an aft thruster that ingests and re-energizes the slowed boundary layer, reducing net drag.

(c) A potential electrical bus design to support more electric aircraft propulsion. Alternating current from the turbogenerators is rectified and distributed over a high-voltage dc bus. This work addresses the inverter design for the subsequent conversion back to ac.

Figure 1.1: Depictions of more electric aircraft propulsion leveraging boundary layer ingestion, as adapted from concepts and figures in [12] and [13]. For illustrative purposes, contributions to drag from the wings of the aircraft are assumed to be less than that of the fuselage and are ignored.

In particular, through the NASA Advanced Air Transport Technology program, a decades-long research program has already contributed to more efficient combustion turbines, while several innovative concepts have been proposed to further reduce emissions – including bold changes in airframe design and hybrid propulsion architectures [12].

Extensive discussion of aircraft is beyond the scope of this work, but it is worth noting two disruptive technologies in particular: boundary layer ingestion and distributed propulsion. In short, the former is when air along the boundary layer (air at the boundary of the airframe) is ingested in a propulsor and “re-energized.” A leading concept places a 1-2.6 MW electric motor at the aft of the aircraft where drag otherwise caused by the wake of the aircraft can be reduced, increasing propulsive efficiency [14]. This concept of boundary layer ingestion (BLI) is illustrated in Fig. 1.1. Distributed propulsion, in which many propulsors are installed across a wing (instead of the typical two large engines) can also provide an aerodynamic boost and enable drastically different shapes in airframe design [15]. Both of these approaches are best facilitated through the use of electric motors, which can be more compact and allow for better design freedom than their combustion counterparts [12, 16]. Yet, while small electric aircraft are in commercial development, only concepts exist for larger aircraft (e.g., greater than 150 passengers) that carry a majority of all air traffic – and performance of electrical systems needs to advance nearly an order of magnitude to realize them [17]. Nonetheless, with industry giants like General Electric and Rolls Royce even engaged in some of the most promising research in commercial electric aircraft, this challenge is being taken seriously [18].

Partially or fully electrifying vehicles while still meeting the demands of industry roadmaps

and consumers alike poses both engineering challenges as well as many human factors. This thesis will focus on just one but very important aspect of vehicle electrification: the inverter drivetrain. This important building block of the vehicle is what converts stored energy into energy that actually spins a motor. This may mean converting energy stored in batteries in an electric car or energy from a high voltage dc bus supplied by generators in a turboelectric aircraft [12]. While it is worth noting that some concepts explore superconducting systems [19] with motors [20] and power electronics [21, 22] operating at cryogenic temperatures, this work is confined to realizations at standard atmospheric conditions. As such, both automotive and aircraft systems share key design requirements, manufacturing challenges and scalability goals for meeting aggressive electrification needs. In addition to meeting specific input voltage, peak output power and power density targets codified in industry roadmaps (which are reviewed in the following section), the concurrent goals of thermal management and reliability are also paramount to these applications.

This work thus proceeds as follows: the remainder of this chapter will review the key performance requirements of future electric propulsion systems. Several prominent and promising topologies will be reviewed, highlighting recent advances in both wide band-gap semiconductor development and integration. Then, an unconventional approach will be introduced that contrasts with much of the prior work – but holds promise to achieve the order-of-magnitude improvements required by the industry targets, and spur the rapid innovation necessitated by the climate crisis [23]. After that, the remaining chapters will provide an in-depth look at the design and analysis of the electrical (Chapter 2) and thermal (Chapter 3) performance, followed by scalability (Chapter 4) and reliability (Chapter 5) of the proposed architecture. Finally, Chapter 6 concludes the thesis and presents potential extensions to the work in future studies.

1.1 Constraints and Performance Criteria

To focus on reducing emissions of single-aisle, twin-engine aircraft, which make up a significant proportion of all air travel, NASA has conducted several studies to estimate how the various interconnected systems on an aircraft might benefit from full, or at least partial, electrification [24–26]. The primary figure of merit is the fuel burn rate, which is a factor of both the weight of the aircraft (more thrust is necessary to keep a heavier aircraft aloft) and efficiency of the propulsion system (which allows travel of greater miles per gallon, as it were). Lower efficiency also means more fuel is required for a given flight, and is thus coupled to the aircraft weight. These studies produced key performance parameters for the electric drivetrain and identified a target design space – in spite of the series of energy conversions involved in a hybrid architecture (chemical to mechanical to electrical back to mechanical) – where a drive designed beyond the break-even frontier would lead to an increase in aircraft propulsive efficiency upwards of 10%. This frontier is defined as a trade-off between efficiency and gravimetric power density (lightweight, yet still supplying the required output power), and culminated in program targets specified in Table 1.1. The minimum targets correspond

Table 1.1: NASA key performance parameters for an electric drivetrain [27].

	Minimum	Target	Stretch
Specific Power	12 kW/kg	19 kW/kg	25 kW/kg
Efficiency	98.0%	99.0%	99.5%

to the break-even frontier; technology here would demonstrate the concept is plausible, while meeting the reach goal would indicate propulsive efficiency gains are indeed achievable.

In addition to these targets that directly drive vehicle efficiency, two higher-level constraints must also be met. First, a sophisticated electrical distribution system is required for any more electric aircraft, and the weight and efficiency of this system must also be optimized. Several studies considered both ac and dc bus systems up to 6 kV, where higher voltages can be traded against lower currents (and therefore lower conductor mass) at the same power level [28, 29]. A synchronous, high-voltage, ac grid had been shown to be the most efficient implementation, but the synchronous nature tightly couples the generator and propulsor speeds and removes a significant degree of freedom in system design. Conversely, a high-voltage, dc bus, with independent rectifiers and variable speed drives providing power conversion as needed, but trades some efficiency for added flexibility [30]. Nonetheless, flexibility was preferred for the designs pertaining to this work, so a dc grid was selected as the lowest mass option, illustrated in Fig. 1.1c. While dc voltages of up to 3 kV will continue to provide system level benefits, 1 kV represents the knee of the performance curve relating input voltage to power density [28] and was established as an appropriate target for this work in [27].

Finally, the high specific power (HSP) motors for this application also pose unique challenges: lightweight, “iron-less” (typically permanent magnet) machines have very low inductance and require a low total harmonic distortion (THD) drive current, while brush-less dc and switched reluctance variants require a carefully shaped current waveform. Some inverter designs in Section 1.2 have made exciting progress towards addressing this application and these motors. However, unconventional hybrid architectures may ultimately provide the great leaps in performance necessary for electric flight.

Recently, the U.S. Department of Energy launched a similar research program to accelerate electrified propulsion through the Advanced Research Projects Agency – Energy (ARPA-E) program titled ASCEND (Aviation-class Synergistically Cooled Electric-motors with iNtegrated Drives) [31]. This aggressive program demands similar performance to that defined by NASA, but also provides a target flight profile to benchmark against (as done in Chapter 2). Additionally, this program further requires compliance with DO-160, a standard covering everything from environmental conditions to electromagnetic emissions. While the former is outside of the scope of this work, the latter is addressed in Chapter 4.

In the past few years, there has also been a similar push to advance electric vehicle technologies. To cater to diverse consumer preferences and use cases, and to forge a path towards energy efficiency and reliability of transportation choices, the U.S. DRIVE (Driving Research

Table 1.2: Key 2025 targets in the US Drive 2025 electric powertrain roadmap [32].

Power Density	Peak Power	Average Power	Efficiency	Voltage Rating	Fundamental Frequency
100 kW/L	100-200 kW	55-110 kW	>98 %	900-1200 V	2 kHz

and Innovation for Vehicle efficiency and Energy sustainability) consortium has sought to foster collaboration between government organizations, private companies and academic institutions to establish a vision for vehicle research and development [33]. Through regular consortium meetings facilitated by the Vehicle Technologies Office in the United States Department of Energy, an Electrical and Electronics Technical Team Roadmap was produced to guide such efforts [32]. This document ties many of the high level goals of emissions reductions, vehicle range and peak power, price point and lifetime to key requirements of future electric drivetrains, several of which are listed in Table 1.2.

The first target of note is the power density. While gravimetric power density and volumetric power density are both important objectives to maximize, the volumetric density is more heavily weighted for automotive applications. This is because the space in automobiles is limited; where it is not used for passenger or cargo, it would otherwise be used for batteries to extend the range of the vehicle. The high power density of 100 kW/L represents an order of magnitude increase compared to what has been possible with extant automotive technologies. The related goal of a high output power is also necessary to ensure desirable acceleration characteristics as well as satisfy towing and other high-torque applications served by conventional utility vehicles. Additionally, the powertrain must be highly efficient to get the most range out of each charge.

The remaining milestones are again, slightly more nuanced. For instance, while current-generation electric vehicles typically host a 400 V dc bus, higher voltages are increasingly more appealing – especially with the advent of wide-bandgap power semiconductors that support these voltage ratings [34]. Although there are system and energy density benefits gained by moving to higher bus voltages, the targets upwards of 800 V dc are actually in support of fast dc charging; there is less loss “at the pump,” as it were, when pushing charge rates above 350 kW [35]. That said, high-voltage drives have also shown performance boosts in heavy-duty vehicles – with a recent demonstration by John Deere of a 200 kW industrial loader featuring a dc bus over 1 kV [34]. As for the motors, most high-specific-power machines have tended to feature a high pole count, and therefore high electrical frequency – upwards of several kilohertz [36–38]. As such, the inverter system must provide the corresponding modulated waveform – necessitating either a much higher switching frequency or additional harmonic filtering.

1.2 Survey of Topologies

The challenges outlined above have motivated an exploration of diverse converter technologies and topologies. While this section is not meant to provide an exhaustive list, it does highlight many inverter implementations that have achieved significant performance against the targets, or otherwise have desirable attributes. In general, these systems have been enabled by wide-bandgap devices [39], and 2- and 3-level designs are ubiquitous – mostly improving in step with devices and packaging technology [40].

1.2.1 Two-Level Half-Bridge

The traditional half-bridge, or six switch module in a three-phase converter, is the 2-level system common to extant electric vehicle drivetrains. This topology is illustrated in Fig. 1.2b, where the ac output node can toggle between the positive and negative rail of the dc bus (i.e., the levels of the 2-level waveform). While silicon insulated-gate bipolar transistors (IGBTs) have been the workhorses of such converters at both high-voltage and high-current for some time, these devices are accompanied with high switching losses and thus limit the switching frequency. As such, significant filtering of the output of each line phase would be required to drive an HSP machines described above.

However, silicon carbide devices – which allow for very high blocking voltages with commensurate conduction losses and significantly lower switching loss as silicon IGBTs – have expanded the operating range for these converters [41]. Additionally, very high output currents (and therefore torque) can be achieved when the devices are paralleled [42]. This is evident in the Tesla Model 3 powertrain [34], the first commercial electric vehicle to use silicon carbide. However, other than paralleling novel devices, there is little room for innovation on this straightforward, but limited topology. Additionally while these devices have reduced intrinsic, dynamic parasitics (e.g., gate charge and output capacitance), switching transition speeds and frequencies are still typically limited to under 100 kHz. While the lower parasitics do reduce switching energy, the fact that the devices must charge and discharge




Figure 1.2: The common, 2-level bridge topology, with corresponding single-phase output waveform.

to the entire bus voltage each commutation is the dominant factor, as switching loss scales with the square of this voltage. This large rate of change in voltage, dv/dt , also means a significant amount of effort must be spent on output filter design if it is meant to interface with any HSP machines requiring sinusoidal input current. As such, much of the literature is spent optimizing magnetics and switching parameters to find the frontier of this space [43].

A significant amount of effort has also been spent minimizing dc bus inductance to allow for faster switching. This parasitic inductance on the dc side of the switches stores energy while the device is in the on state. When the device switches off, this energy charges the switch drain-source capacitance. If the parasitic inductance is large enough, at high currents this stored energy can cause detrimental overshoot of the device. Unless this inductance can be reduced, the only other recourse is to slow the speed of the device transition, though this would incur greater overlap loss (see Chapter 2). As such, the best designs use in-depth modeling, careful routing and innovative manufacturing methods to achieve a low-inductance, simulated (though not measured) to be as low as 6 nH for 1 kV designs [44].

1.2.2 Three-Level T-Type

The 3-level T-Type converter phase-leg is illustrated in Fig. 1.3a. Here, the dc bus is divided into three levels, $+V_{dc}/2$ and $-V_{dc}/2$ as well as a neutral midpoint. The four switches in each phase leg (12 switches total in a three-phase converter) connect the ac output to these levels in a sequence defined by the specific modulation strategy to generate the waveform shown Fig. 1.3c. Note, a four-quadrant switch is needed at the midpoint connection to prevent self-commutation (through the diodes) to the wrong state when inductively loaded. For example, a positive current flowing out of the v_{ac} node would always turn on D3 before D4 without the




Figure 1.3: The 3-level topologies highlighted in this chapter, with corresponding output waveform.

reverse switch S2 to selectively prohibit this path as needed when driving positive current at the negative half of the line cycle. Since the four-quadrant switch only sees half of the dc bus voltage, lower voltage devices may be used – but the other two devices will need to block the full dc bus voltage. Additionally, complementary switch pairs S1/S3 and S2/S4 only commute during half of the cycle, so implementations with silicon IGBTs can still yield improved performance [45].

Of course, switching losses can be significantly reduced by using silicon carbide, as seen in a recent 250 kW scale demonstration [46]. This system used commercially available devices to achieve 98.5% efficiency and 25 kW/L with liquid cooling. To leverage the reduced costs and improved current handling of IGBTs, [47] further proposed paralleling a silicon IGBT with a silicon carbide MOSFET to make a hybrid switch for S1 and S4 to reduce conduction losses; the silicon carbide MOSFET would be used for fast turn-on and turn-off, while the IGBT would be switched during the rest of the on-time to carry the majority of the current and reduce conduction losses. To ensure correct current sharing between the two device types, busbar design and module arrangement was again a significant design consideration. Inductance in the dc link was as high as 12 nH, though the differences between paralleled switches was on the order of 2 nH [48].

1.2.3 Three-Level Active Neutral-Point-Clamped

The active neutral-point-clamped (ANPC) is the second promising three-level design, shown in Fig. 1.3b. The dc bus is again divided into three levels, with six switches in each phase leg. The switching scheme for generating the three level waveform ensures that each device only sees half of the bus voltage, allowing the use of lower voltage devices and operation at reduced switching loss [49]. Additionally, complementary switches groups (S1+S3)/S2 and (S2+S4)/S3 only commute during half of the cycle, while the switch pair S5/S6 only commutates once per line cycle. This effectively means S5 and S6 can be devices with better conduction characteristics at the cost of worse dynamic parasitics. Indeed, in the megawatt-scale demonstration for NASA, General Electric used proprietary 1700 V, 500 A silicon carbide half-bridge modules for both S1/S2 and S3/S4 half-bridges and a single 1700 V, 600 A silicon IGBT half-bridge for S5/S6 [50]. Depending on the commutation loop (i.e., the loop between silicon carbide or silicon devices), the reported inductance ranged from 6.5 nH to 17.5 nH for this 2.6 kV design [51]. As it were, this high bus voltage meant that even divided among three-levels, the output still experienced extremely high dv/dt . As such, a major caveat remains given that significant amount of the total inverter mass was from to the ac output filter alone (a factor of three greater than the contribution from the power modules).

Nonetheless, it a remarkable milestone for this system to have delivered 99% peak efficiency at 12 kVA/kg – as demonstrated at a full megawatt. While this system might not be scalable, in that the power modules, filtering considerations and thermal management were purpose built for a single power level, it highlights the benefit of a new architecture and new semiconductors for the application. Indeed, to process that much power a specially designed

“pump-back” system was also engineered to circulate current between two prototypes of the three-phase converter [52]. This test setup can thus serve as a useful reference when the array of Chapter 4 is tested at higher power levels in future work.

1.2.4 Higher-Level Topologies

Multi-level versions of the T-Type and ANPC converters can be found in the literature, but require significantly more switches and series-stacked dc capacitors as the number of levels increases. Additionally, routing of the commutation loops becomes much messier, as many more midpoint connections and half-bridge nodes need to overlap and intersect. Thus, the overhead from routing these networks becomes impractical for this power-dense application.

The cascaded H-bridge and [53] and modular multilevel converter [54] are other (typically) higher level-count topologies that have been applied to high fundamental frequency drives. However, in both topologies, the capacitors must be sized to support the full fundamental current. This imposes a significant constraint on the size of the capacitors used in each application: for large output currents, these capacitors must also be large – on the order of hundreds of microfarads to millifarads, making them necessarily bulky. The exact analysis of this power density implication is explored further in [55]. Additionally, in the modular multilevel converter, the phase-legs must manage current between the upper and lower halves of each phase-leg to guarantee the appropriate output and minimize circulating currents that do not deliver power to the load. Both topologies have been able to show good harmonic performance and exhibit cell-level redundancy (see Chapter 5), but in the end, fail to deliver on power density.

1.3 An Unconventional Approach

This work eschews more traditional approaches to search for the breakthrough needed to meet the demanding performance targets. As such, identifying desirable system characteristics has informed this process. In addition to the requirements in Section 1.1, several important aspects may be considered from the previous discussion:

- Multilevel designs divide the output waveform into smaller, discrete steps. This reduces the THD of the waveform, as well as the dv/dt – both desirable effects that reduce filtering requirements and support integration with HSP motors.
- Multilevel designs that can reduce off-state voltage stress of the transistors are preferable, as they dramatically reduce switching losses. Also, as both switching parasitics and conduction losses decrease with lower voltage devices, the lower voltage allows the use of higher figure-of-merit devices [56].
- High switching frequencies are preferable, as these frequencies are more easily filtered by smaller passives, improving power density.

- While several designs above have performed well at the high-power operating point they were designed for, true scalability might be obtained by designing a building block from which several configurations of output power might be obtained.
- Novel topologies and devices provide new opportunities.

By framing various design insights in this way, one unconventional topology stands out as a compelling candidate: the flying capacitor multilevel (FCML) converter [57], shown in Fig. 1.4. Originally proposed nearly two decades ago, this circuit also uses a number of charged capacitors to provide the discrete steps available in the output waveform. However, these capacitors only need to support ripple switching frequency [55]; while the peak of these currents tracks the current at the ac output, it is only during the switching period – not the entire fundamental. At the same time, phase-shifted (carrier) pulse-width modulation (PSCPWM, or simply PSPWM) schemes mean each switching pulse is interleaved through the series connection of cells. This leads to a high effective output frequency, where the switching frequency seen at the switched-node v_a is a factor of $(N-1)$ greater than the switching frequency (also shown in Fig. 1.4), where N is the number of levels. The benefits

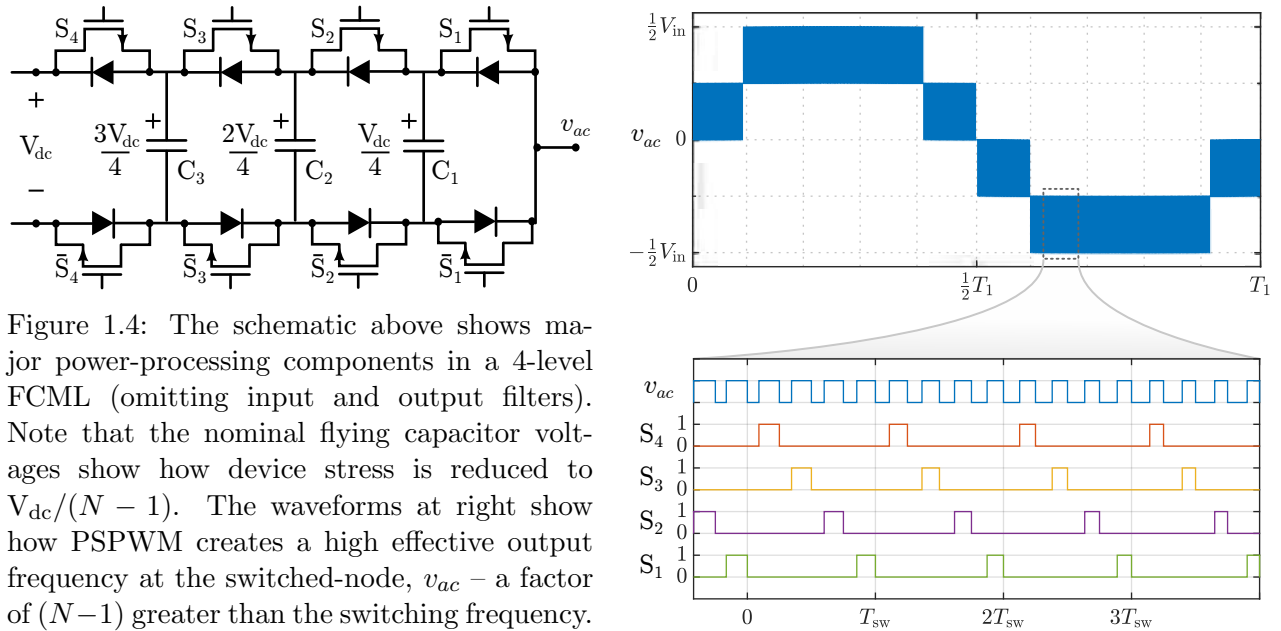


Figure 1.4: The schematic above shows major power-processing components in a 4-level FCML (omitting input and output filters). Note that the nominal flying capacitor voltages show how device stress is reduced to $V_{dc}/(N-1)$. The waveforms at right show how PSPWM creates a high effective output frequency at the switched-node, v_{ac} – a factor of $(N-1)$ greater than the switching frequency.

Table 1.3: Key benefits of the FCML topology [57].

	2-Level	FCML
Switch Stress	V_{in}	$V_{in}/(N-1)$
v_{ac} Ripple Amplitude	V_{in}	$V_{in}/(N-1)$
v_{ac} Ripple Frequency	f_{sw}	$f_{sw} \cdot (N-1)$
Output Inductance	$L_{2-level}$	$L_{2-level} \cdot (N-1)^2$

of this effect are highlighted in Table 1.3, where the key value proposition is highlighted by the dramatic scaling of the required output filter inductance.

Since the original conception, and indeed most implementations in the years since conception, were targeting stationary, grid-frequency applications, designs were constructed with slow-switching IGBTs and bulky film capacitors on the order of tens of microfarads. However, advances in wide-bandgap power transistors – gallium nitride specifically – and ceramic capacitors have driven a revival of hybrid, switched-capacitor architectures such as FCML [58]. For instance, the class II dielectrics that have been ubiquitously adopted for these applications simultaneously support high voltages and high capacitance in a modest package, leading to extreme energy densities [59]. As clear in Fig. 1.4, capacitors comprise a bulk of the power processing components, applying this gain in energy density to produce extreme power density in the converter.

Additionally, though younger than silicon carbide devices, gallium nitride offers significantly lower switching and conduction loss in the 100 V to 300 V device voltage ratings considered in this work [56]. The series-stacked nature of the devices in the topology means that device voltage rating is also divided by a factor of $(N-1)$, allowing lower-voltage, high figure-of-merit transistors – even for a design with a 1 kV dc bus. Compared to the topologies above, this series connection also means routing is more straightforward and reasonably scales with level count. The simultaneously high figure-of-merit (specifically low capacitances that would lead to switching loss) and lower per-device voltages means switching frequencies can be increased dramatically. While the designs above were largely constrained to kilohertz or low tens of kilohertz, switching frequencies of hundreds of kilohertz to upwards of a megahertz are well within the design space for these devices. For instance, a 10-level FCML can leverage 200 V GaN devices switching at 115 kHz to produce a greater than 1 MHz output waveform with one ninth of the dv/dt of a comparable 2-level converter. This high frequency output lessens the impact of weight penalties associated with the output filter, while also allowing for control of high frequency motors. Furthermore, the devices also support higher current densities, meaning if the individual converter module is rated appropriately, using a multitude of devices can still result in a compact design.

As for the notion of a converter *module*, this approach is unconventional in the way that it interprets *scalability*. The fundamental approach is to create a modular design that, while still capable of tightly integrating with the rest of a thermoelectrical drivetrain, divides a high output power requirement into smaller, more tractable design elements. As such, a 160 kW three-phase converter could be comprised of nine paralleled, 18 kW modules (see Chapter 4). This also means that a common module design, or at least design approach, could be leveraged across applications and target power levels – from the vehicle drivetrains discussed in this work, to solar inverters [60] and even space applications [61]. The notion of scalability also extends to the manufacturing process, where electrical design uses standard, mass-production printed circuit technology compared to the proprietary modules in the prior works. As such, gains in both scaling and cost can be achieved.

In conclusion, this chapter highlighted the pressing need for innovative technology to decarbonize the transportation sector. Several traditional topologies, many with commendable

performance, where highlighted to provide context and contrast to the proposed strategy in this thesis. While the discussion above presented the general benefits of the architecture in this work, the remaining chapters will further highlight the innovations in the electrical and thermal domains, as well as the progress towards achieving scalable and reliable performance.

Chapter 2

Inverter Module Electrical Design

Portions of this chapter are adapted in part or in whole from [1] and [62].

While the FCML was originally developed for large, stationary applications [57, 63, 64], advances in high dielectric constant (high-k) capacitors and wide-bandgap devices have driven miniaturization of converter systems to achieve extreme performance [56]. Over the past few years, power dense inverter designs based on the FCML have transformed from proof-of-concept demonstrations [65–68], to hardware platforms that focus on control and scalability challenges [69–72]. This chapter details an approach for the electrical design and testing of a modular inverter for use in multi-phase, paralleled, and/or segmented operation alongside other modules [73–75].

The foundation of this FCML module electrical design is based on a previous generation of FCML inverter module: the dual-interleaved, 9-level converter presented in [62]. That design was the result of a Monte Carlo optimization; different combinations of extant components (inductors, capacitors and transistors) were selected at random and used to define a computational converter object similar to that presented in Section 2.1. If a randomly discovered converter configuration passed key constraints, such as rated component voltages and currents, the estimated losses for this design candidate were calculated using a the loss model that is a precursor to the one in Section 2.2. Many lessons were learned from this initial design effort, including the frequent necessity to characterize component parameters beyond what is available from the manufacturer in order to have an accurate loss model. As such, rather than re-run the same simulation with potentially inaccurate component parameters, in this work, the following advances what was already an exceptional design candidate and refines the loss model – fundamental to the optimization process – to culminate in an extensive study of a single, robust design.

Indeed, there are several key changes in the converter of this work compared to [62]. First, a recent discovery indicated that FCML converters with an even number of levels tend to exhibit better natural balancing than designs with odd level counts. This is particularly notable when the dc input to the converter has a nonzero source impedance [76]. Since increased conduction losses in the 9-level design (due to dynamic on-resistance, see Section

Table 2.1: Capacitor selections per interleaved phase.

	DC Bus		C_8	C_7	C_6	C_5	C_4	C_3	C_2	C_1
Technology	Film	Series	NPO 630V	Series X6S 450V						
Nom. DC Voltage [V]	1000	1000	889	778	667	556	444	333	222	111
Voltage Derating [%]	–	77	86	74	61	49	49	37	25	12
Component Cap. [μ F]	–	0.1	0.1	—————			2.2	—————		
Num. Parallel/Series	–	20/2	25/2	6/2	—————		4/2	—————		
Nom. Cap. [μ F]	28.2	1	1.25	6.6	—————		4.4	—————		
Derated Cap. [μ F]	28.2	1	1.25	1.3	1.0	1.2	1.4	1.8	2.4	3.5

2.2 and [77]) are approximately equivalent to the loss introduced by an additional series device, the 10-level topology was preferable. Additionally, the reduced voltage stress at 10 levels allowed for possible use of lower voltage (i.e. 150 V) devices with better figures-of-merit, albeit with reduced overhead. In the end, transistor losses for this design were minimized by using newly-released EPC2034C devices. These feature better performance and voltage ratings were chosen for the experiments and results presented later in the chapter.

A 10-level design also provided better voltage derating of the flying capacitors. As energy dense capacitors facilitate much of the high specific power conversion in this topology, appropriate selection of the flying capacitors is vital. An influential survey on capacitor technologies [59] originally indicated that the 450 V, 2.2 μ F X6S ceramic capacitor by TDK would yield the best power density. However, use as flying capacitors in the higher voltage nodes of [67, 73] conflicted with the component voltage rating, while increasing series connections would decrease overall energy density. Early on, emerging 500 V and 630 V, X7R capacitors seemed to provide an suitable alternative. Yet, subsequent intensive component characterization in [78, 79] revealed that, while the X7R capacitors provided a comparable energy density for this voltage rating, they also had significantly lower quality factor – contributing to increased losses in the converter of [80].

As such, this work again leverages the X6S component for a majority of the flying capacitance. Because the capacitance of this dielectric decreases with applied dc voltage bias, the highest X6S node (C_7) also includes 50% more physical capacitors than the others. Since this capacitance decrease is so significant at high voltage (i.e. -80% of the rated value), 630 V NPO capacitors are used for the remaining flying node (C_8) and local dc bypassing with little loss in energy density. This change also lead to a more practical design, as and the 900 V rating of the series X6S capacitors fell below safe operating limits at the highest voltage nodes of both the 9- and 10-level design, while the 1300 V rating of the series NPO components is more than adequate. Table 2.1 breaks down the capacitor selection for this design, and includes derated capacitance at 1 kV operation as well as the nominal voltage at each level. The series connection of NPO and X6S capacitors means the rated voltage for the given flying capacitance is nominally twice that of the component rating. Note, the film capacitance listed is part of the backplane where the module is installed, and is not included in module power density calculations.




Figure 2.1: Fabrication, installation and high voltage testing of a custom busbar for C_8 .

The use of NPO capacitors at the highest voltage nodes was facilitated by the release of NPO capacitor arrays that allow for a more efficient use of board area. Due to the limited capacitance values available with this dielectric, many more discrete components need to be paralleled to achieve the same order of flying capacitance as the rest of the levels. These arrays fixture the same form factor capacitors in a vertical orientation, allowing approximately twice as much capacitance for the same board area. For C_8 , a custom busbar was also fabricated to parallel additional arrays, as shown in Fig. 2.1. Routing this connection off-board allowed for a more efficient power path, as well as better noise isolation and high-voltage clearances to signals routed on the PCB itself.

The full converter bill of materials can be found in Table 2.2. Here, part numbers for the

Table 2.2: Converter bill of materials.

Component	Part Number	Parameters
X6S capacitors	TDK - C5750X6S2W225K250KA	2.2 μF , 450 V
NPO capacitors	TDK - C5750C0G2J104J280KC	0.1 μF , 650 V
	TDK - CAA572C0G2J204J (2-up)	0.2 μF , 650 V
	TDK - CAA573C0G2J304J (3-up)	0.3 μF , 650 V
Output capacitors	TDK - C5750C0G2J104J280KC	0.1 μF , 650 V
	- 6 in parallel per phase	\rightarrow 0.6 μF
Output inductors	Coilcraft XAL1510-153	15 μH
	- 2 in parallel per phase	\rightarrow 7.5 μH
GaN transistors	EPC - EPC2034C	8 m Ω , 200 V
Gate driver	TI - LM5114	GaN
Gate resistors	Turn-on	15 Ω
	Turn-off	0 Ω
Digital isolators	Silicon Labs - SI8610BC-B-IS	
Isolated power supplies	Analog Devices - ADUM5010ARSZ	

Table 2.3: Converter mass breakdown.

Component	Mass [g]
aluminum heat sinks	85.0
3D-printed assembly	107.0
mounting hardware	14.1
- springs	2.8
- #0-80 screws	4.7
- #0-80 brass inserts	5.9
- #0-80 SMT nuts	0.7
flying capacitor cells	95.7
- X6S capacitors (x 120)	48.0
- NPO capacitors (x 100)	38.7
- FETs and isolation	9.0
filter passives	89.4
- input capacitors	31.5
- output inductors	54.3
- output capacitors	3.6
component balance	20.3
module PCB	76.1
local controller	4.9
<i>TOTAL MASS</i>	492.5 g

capacitances described in Table 2.1 are shown, as well as the parameters for other key power-processing components. Additionally, it should be noted that this specific implementation of the FCML uses an isolated power supply for each transistor voltage domain. While this is largely considered a legacy design (compared to newer systems leveraging a cascaded bootstrap to power each floating gate driver [81–83]), this vastly simplifies debugging and circuit analysis. Likewise, a separate digital isolator is used in conjunction with a discrete gate driver, although newer isolated gate drivers may be preferable in newer designs. As such, future work will nonetheless benefit from the decreased mass and component area required by these newer approaches, and is expected to further the power density beyond what is achieved in this work. Indeed, the mass breakdown of this converter – valuable in determining which components most affect gravimetric power density – is reported in Table 2.3. Here, the mass is shown to be fairly evenly split across hardware, heat sinks, filter passives and flying capacitors, so an opportunity to reduce mass in any category will provide a relatively comparable improvement – and may indicate a lower level count design is worth additional consideration.

Finally, an annotated schematic and photograph of the converter module used as the research platform for the bulk of this work is shown in Fig. 2.2. Here it is evident that in ad-




Figure 2.2: The schematic of the 10-level, dual-interleaved FCML inverter module of this work, accompanied by the top and side views of a hardware prototype. The converter is shown partially assembled to illustrate the mechanical assembly and thermal management, with the exception that screws used for heat sink mounting are replaced with an improved design in Chapter 3.

dition to creating a compact design to maximize power density, with components populated to best use the converter area while still complying with clearance constraints. Mechanical structures were integrated into the design to fixture heat sinks and facilitate thermal management (further illustrated in Section 3.4.2), as well as provide support and handling of modules when inserted into an array backplane. The card edge connectors also facilitated easy parallelization with other modules, discussed further in Chapter 4.

2.1 Structure for Computational Analysis

This section introduces modeling variables in the format of a MATLAB structure. The overall goal of this format is to bridge the equations used to describe the theory of operations with the actual implementation of the physical quantities in the calculations performed in this work. This will enable readers to extend this work to their own applications, or allow this work serve as a companion to any future releases of code. Several key a converter parameters must be captured and stored computational structure in order to analyze performance of a single converter configuration, as well as support future Monte Carlo studies. While later generations of the software may benefit from a fully defined MATLAB class (as is the case for several of the the inverter components), the structure data type is sufficient for the research presented here. As such, an `inv` structure is configured to describing a given inverter based on operating and hardware configuration parameters. The operating parameters are:

- `inv.fs` - Transistor switching frequency, typically expressed as f_{sw} .
- `inv.td` - Programmed dead-time between turn-on and turn-off of complementary switches. Symmetric for both turn-on and turn-off.
- `inv.fmod` - Fundamental output frequency, used in output waveform generation.
- `inv.M` - Inverter modulation index, used in output waveform generation.
- `inv.Vd` - DC bus voltage setpoint, used in output waveform generation, nominal drain-source voltages, as well as flying capacitor biasing and derating.
- `inv.Tamb` - Ambient temperature in degrees Celsius. Applicable if device operating temperatures are to be determined based off of a target thermal impedance using an iterative solver (described in Section 2.2.4).
- `inv.Rth` - Thermal impedance from transistors to ambient. Informed by work elaborated upon in Chapter 3, this parameter allows for calculation the device temperature given an ambient operating temperature in tandem with loss calculations.
- `inv.Ttrans` - Transistor device temperature. Overrides any closed-loop calculations that include thermal impedance by fixing the device junction temperature, T_J .

The converter structure parameters are:

- `inv.N` - Inverter level count, denoted simply as N .
- `inv.C` - An array of `Capacitor` objects, with each `inv.C[k]` where $k \in 1 \dots N-2$ is an instance of the `Capacitor` class configured to describe a given flying capacitance; this object can reflect the total effective capacitance and ESR (at correct bias) of several series and parallel connected components. Additional parameters and functions exist within this object and will be invoked in the loss modeling section.
- `inv.L` - An `Inductor` object, used to describe the output inductance. This object can reflect the total inductance and soft-saturation characteristics (at correct current) for several series and parallel connected components within a single logical inductance.
- `inv.T` - A `Transistor` object, used to describe the transistor used across the design; assumes that only one device type and configuration will be used in the converter. Additional parameters and functions exist within this object and will be invoked in the loss modeling section.
- `inv.Rg,on` - Extrinsic gate resistance, used to slow the turn-on of the device defined in `inv.T`. Affects overlap losses calculated in the `inv.T` loss function.
- `inv.Rg,off` - Extrinsic gate resistance, used to slow the turn-off of the device defined in `inv.T`. Affects overlap losses calculated in the `inv.T` loss function.
- `inv.Lcom` - Nominal switching cell commutation loop inductance, either estimated from typical board geometries or dynamically by integrating the method in [84]. Can be used to estimate device overshoot.
- `inv.RPCB` - Circuit board trace resistance, either estimated from typical board geometries or dynamically by integrating the method in [84]. Used for calculating ohmic losses; typically relevant only for high-power operation.

For clarity, the `inv.` prefix will be omitted in the following sections. This should be noted when parameters and functions within `inv.T` and `inv.C` are referenced in the respective loss calculations.

2.2 Loss Modeling

In this work, the low passive component values and high-speed transistors in the multilevel topology all provide many system-level benefits. However, moving away from more traditional two or even three-level designs requires a more nuanced approach to loss modeling. Where averaged models may have been previously suitable [85], fluctuations in state values

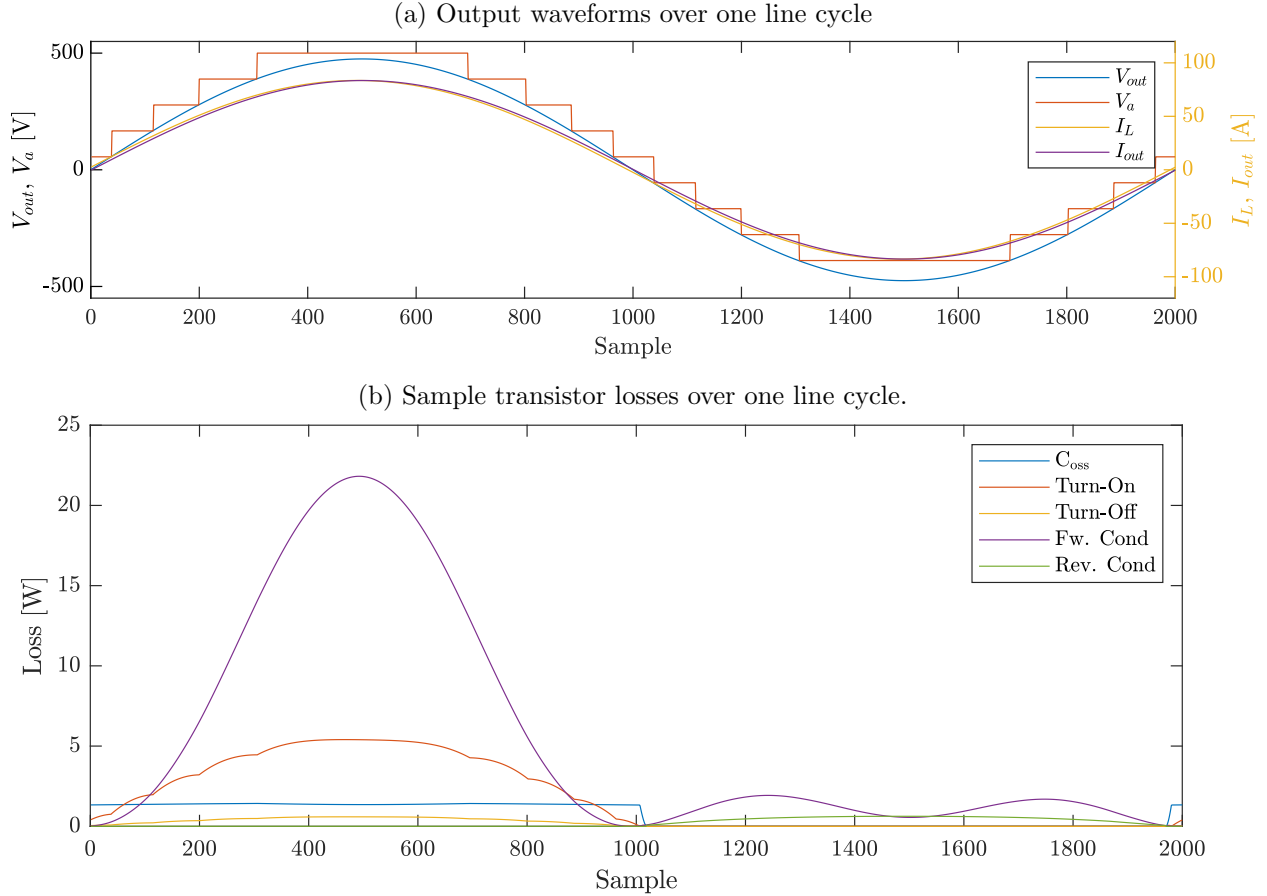


Figure 2.3: Sampled converter waveforms used to simulate and model FCML losses.

(e.g., flying capacitor voltages), high ripple currents, and partial soft-switching in the FCML and drastically affect losses as the converter moves through each period of the fundamental.

Therefore, to model converter losses under inverter operation, steady-state losses are calculated at each operating point along the output waveform then averaged over one fundamental cycle. As such, much of the methodology that follows is similar to [62], but reflects additional considerations, such as high output current ripple contributing to increased RMS losses and the effect of incorporation of dynamic on-resistance, dR_{on} . Note, this methodology allows for calculating losses given an arbitrary reference. That said, though FCML inverter is capable of generating an arbitrary output waveform (contingent upon harmonic content of the reference residing well below the switching frequency), the following discussion is limited to sinusoidal modulation for simplicity. Fig. 2.3a shows the simulated output waveforms for an 18.9 kW operating point. Here, I_{out} is calculated as

$$I_{out} = \frac{V_{out}}{Z(2\pi \cdot \text{inv.fmod})}, \quad (2.1)$$

where V_{out} is the nominal inverter output voltage (driven by a sinusoidally modulated duty

cycle, D), and Z is the load impedance evaluated at the modulation frequency.

Next, at each point sampled along this line cycle, component losses are calculated as if the converter were in steady-state operation at any of these points. For example, the plot in Fig. 2.3b shows the magnitude of different kinds of transistor losses at each of these sampled operating points. Note, this is not a plot over time – it is not a waveform as such; these are the losses associated with an operating point described by the output current, voltage and duty cycle at that instant in the line cycle. Then, for steady-state *inverter* operation, it's the average value of these losses over the full line cycle. The total converter losses are found by summing these averages for all components in the inverter.

While Section 1.3 detailed many of the benefits and operating principles of the FCML inverter, some of the preliminaries with respect to converter operation are necessary to review. First, for an N -level converter operating at a switching frequency f_{sw} under PSPWM, the effective frequency seen at the switch node v_x is

$$f_{\text{eff}} = (N - 1) \cdot f_{sw}. \quad (2.2)$$

Correspondingly, as the duty cycle, D , of each of the switches is, e.g., sinusoidally modulated, the effective duty cycle seen at the switch node, D_{eff} , will vary according to

$$D_{\text{eff}} = (N - 1) \cdot D \bmod \frac{1}{N - 1}. \quad (2.3)$$

These effective quantities are useful for determining additional current stress that contributes to component losses (not captured in [62]). Specifically, in addition to the fundamental current where the magnitude is attributed to the specific loading at the output, a significant current ripple will be generated in the output inductor at f_{eff} :

$$i_{L,\text{pk}} = \frac{v_h - v_{\text{out}}}{L_{\text{out}}} \quad (2.4)$$

where

$$v_h \approx [(N - 1) \cdot D] \cdot V_{\text{DS,nom}} \quad (2.5)$$

and

$$V_{\text{DS,nom.}} = \frac{V_d}{N - 1}. \quad (2.6)$$

. A total current value seen by the transistors and inductor can then be approximated by:

$$i_{L,\text{RMS}} \approx \sqrt{I_{\text{out}}^2 + \left(\frac{i_{L,\text{pk}}}{\sqrt{3}}\right)^2} \quad (2.7)$$

2.2.1 Transistor Losses

The transistor losses are broken down into five categories: C_{oss} , dead-time (i.e., reverse “diode” conduction), turn-on, turn-off, and (forward) conduction. These five loss types are

illustrated in Fig. 2.3b, where the instantaneous loss is plotted for a single device over a complete line cycle. While [86] also examined nuances in loss distribution among paralleled devices, the modeling in this section only considers aggregate parameter scaling when paralleling; for example, halving peak device currents, halving on-resistances and doubling output capacitances.

C_{oss} losses are switching losses associated with the charge and subsequent discharge of the MOSFET drain-source capacitance under hard-switching conditions. From [87], this energy falls into two categories: 1) the energy stored in the off-state device that is discharged during hard turn-on:

$$E_{oss} = \int_0^{v_{DS,eff-oss}} v_{DS} \cdot C_{oss}(v_{DS}) dv_{DS}; \quad (2.8)$$

and 2) the energy dissipated in this device due to the charge flowing through it as the complementary device turns off (and the respective complementary C_{oss} is charged):

$$E_{Qoss} = \int_0^{v_{DS,eff-Qoss}} (v_{DS,eff-Qoss} - v_{DS}) \cdot C_{oss}(v_{DS}) dv_{DS}. \quad (2.9)$$

In both equations, $v_{DS,eff}$ is the effective voltage a turn-on to account for the fact that an inductive load current will allow for soft-switching of half of the devices, depending on the half-cycle of the fundamental. While recent work has shown that soft-charging of the C_{oss} of Si [88], GaN [89], and SiC [90] semiconductors is not an altogether lossless process, that second order effect is most apparent in very high-frequency resonant applications. Since C_{oss} is nonlinear, the $v_{DS,eff}$ used for the two switches, i.e., for the separate E_{oss} and E_{Qoss} calculations, are computed slightly differently:

$$v_{DS,eff-oss} = \min(\text{T.VossFn}(\text{T.QossFn}(v_{DS}) + (i_{on})/2 \cdot t_{de}), v_{DS}) \quad (2.10)$$

$$v_{DS,eff-Qoss} = \min(\text{T.VossFn}(-i_{on}/2 \cdot t_{de}), v_{DS}), \quad (2.11)$$

where $i_{on} = I_{out} - i_{L,pk}$. In (2.10), the calculation is determining how much of the transistor charge has been removed by half of the current i_{on} over the course of the effective dead-time, t_{de} . The effective $v_{DS,eff-oss}$ is then calculated by determining the drain-source voltage associated with this remaining charge. In (2.11), the calculation determines how much the other device has charged with the remaining half of i_{on} over the same time. The min function is purely a computation check to ensure that the drain-source voltage does not go *above* the starting off-state value during the part of the line cycle without soft-switching. For reference, an example of the nonlinear transfer function for this off-state, drain-source charge is shown in Fig. 2.4.

The effective dead-time, t_{de} , is used to compensate for the turn-on characteristics arising from gate threshold dynamics of a specific device:

$$t_{de} = \text{inv.td} + \left(\frac{\text{T.Q}_{g,th} \cdot (\text{inv.R}_{g,on} + \text{T.R}_g)}{\text{T.V}_{drv} - \text{T.V}_{th}(T_J)} \right) - t_{d,off}, \quad (2.12)$$




Figure 2.4: Typical nonlinear Q_{oss} characteristics for an EPC2034C gallium nitride transistor.

where `inv.td` is the programmed dead-time, $T.Q_{g,th}$ is the gate charge at the threshold voltage (datasheet parameter), $T.R_g$ is the gate resistance internal to the device (datasheet parameter), $T.V_{drv}$ is the nominal voltage the gate is driven to when on (which can vary between inverter designs, but typically chosen to be near the maximum allowable for a given device) and $T.V_{th}(T_J)$ is the device threshold voltage as a function of junction temperature (datasheet parameter). The middle fraction in (2.12) represents the additional time beyond the programmed dead-time where the off-state device remains off while the gate is being charged to the threshold voltage. Conversely, $t_{d,off}$ represents a turn-off delay caused by the opposite effect taking place on the on-state device, which remains on until the gate is discharged to the plateau region [91]. This period is calculated as:

$$t_{d,off} = (\text{inv}.R_{g,off} + T.R_g) \cdot T.C_{iss} \cdot \log\left(\frac{T.V_{drv}}{T.V_{plat}}\right), \quad (2.13)$$

where $T.C_{iss}$ and $T.V_{plat}$ can both be read from the device datasheet.

Combining all of these details, the total losses arising from hard-switching the device output capacitance can be computed using:

$$P_{oss} = (T.EossFn(v_{DS,eff-oss}) + \{T.EossFn(v_{DS}) - T.EossFn(v_{DS,eff-Qoss})\}) \cdot f_{sw} \quad (2.14)$$

The calculations described in (2.10) – (2.13) describe the prerequisites necessary to compute the integrals in (2.10) and (2.11). For arbitrary values of v_{DS} these integrals may be either pre-computed or read from the datasheet then entered into a function, $T.EossFn$, as done for (2.14). An example of this function is shown in Fig. 2.5. The first term in (2.14) thus represents the integral (2.10), while the group of terms in curly braces represent the integration in (2.11). Multiplying these two energies by the switching frequency provides the corresponding switching loss.




Figure 2.5: Typical nonlinear E_{oss} characteristics for an EPC2034C gallium nitride transistor.

Dead-time losses refer to the losses associated with *reverse conduction* in the device. This occurs during the effective dead-time, before the gate driver has been engaged to fully turn on the device. In silicon, this reverse path through the off-state switch is through a physical PN diode junction. However, in GaN, this diode-like effect is actually a consequence of enough source-drain (i.e., gate-drain) voltage building up to enhance the channel and allow conduction [92]. Behavior of this current-dependent, source-drain voltage can be referenced from the device datasheet and defined in the transistor object as `T.revFn`. The typical characteristic of the EPC2034C device are shown in Fig. 2.6 and illustrates why – at high current – dead-times are ideally minimized. Then, losses can be calculated as the product of the drain current magnitude at the time of turn-on or turn-off and the corresponding reverse voltage, scaled by the actual time spent in this state each switching period:

$$P_{\text{dead-time}} = (t_{\text{de,on}} \cdot \text{T.revFn}(i_{\text{on}}, T_J) \cdot (i_{\text{on}}) + t_{\text{de,off}} \cdot \text{T.revFn}(i_{\text{off}}, T_J) \cdot (i_{\text{off}})) \cdot f_{\text{sw}}, \quad (2.15)$$

where $i_{\text{on}} = I_{\text{out}} - i_{L,\text{pk}}$ and $i_{\text{off}} = I_{\text{out}} + i_{L,\text{pk}}$ and the effective dead-times at turn-on and turn off are, respectively:

$$t_{\text{de,on}} = \max\left(t_{\text{de}} + \min\left(\frac{\text{T.QossFn}(v_{\text{DS}})}{i_{\text{on}}/2}, 0\right), 0\right) \quad (2.16)$$

$$t_{\text{de,off}} = \max\left(t_{\text{de}} + \min\left(\frac{\text{T.QossFn}(v_{\text{DS}})}{i_{\text{off}}/2}, 0\right), 0\right) \quad (2.17)$$

Just as the C_{oss} loss calculation accounted for soft-switching, the speed at which the self-commutation under inductive currents charges and discharges the two output capacitances is also included here to determine the effective time spent in reverse conduction. This time is calculated by dividing the charge on C_{oss} in the off-state by the current at turn-on or




Figure 2.6: Typical reverse conduction characteristics for an EPC2034C gallium nitride transistor.

turn-off – divided between the two commutating devices. The min and max criteria ensure that only the relevant half-cycle of the fundamental is considered; capacitors are not charged above the nominal V_{DS} or below 0, and the dead-time is not longer than otherwise possible based on circuit and timing constraints.

Note, expressions for i_{on} , i_{off} , $t_{de,on}$, and $t_{de,off}$ for the top switch and bottom switch will be slightly different: the top switch will always be turning on at the valley of the inductor current ripple and off at the peak, with the bottom switch experiencing the opposite. This specific consideration is only necessary for examining losses of individual devices along the line cycle. If the average losses at the end of a fundamental period are the only desired results, only one may be calculated and simply doubled at the final loss tabulation; on average, the losses will be equivalent.

Turn-on and turn-off losses, often referred to as *overlap losses*, are switching losses that arise when transistors are hard-switched. Depending on the direction and magnitude of output current, switching dynamics for the half of the FCML experiencing hard-switching can be approximated by the waveforms illustrated in Fig. 2.7. For turn-on, Fig. 2.7a shows how a switch that was previously blocking begins to turn on and carry current once v_g reaches the threshold voltage, $T.V_{th}$. However, the device still simultaneously sees the off-state voltage until the drain current, i_D , equals the inductor current. Then, v_{DS} declines as the gate driver continues to charge the gate-source and gate-drain capacitances through the plateau region. In the bottom plot, the area of the shaded region indicates the energy lost from the voltage and current overlap during this transition. The computation of turn-on losses effectively requires calculating this area and multiplying by the switching frequency. An illustration of turn-off in Fig. 2.7b shows a similar process for determining energy lost during that transition.

To calculate the energy, the three parameters that define the triangular area in both plots

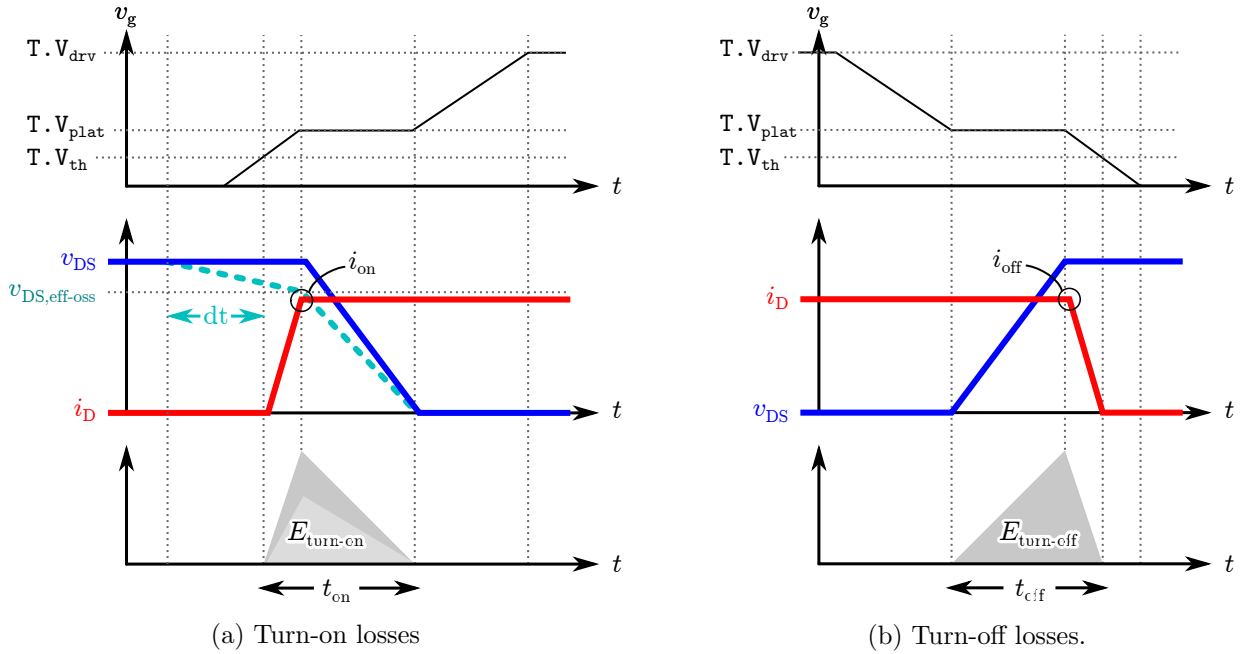


Figure 2.7: Turn-on and turn-off waveforms.

need to be determined. For turn-on, these parameters are $v_{DS,eff-oss}$, i_{on} and t_{on} . The first two have been determined in calculations for the previous two transistor loss mechanisms. Correspondingly, note that Fig. 2.7a illustrates how v_{DS} could be discharged to partial or full soft-switching if the output current is the right sign and magnitude during commutation dead-time. This leaves t_{on} , i.e., the time it takes for the gate current to carry the device from threshold across the plateau region. From [87], this can be approximated by:

$$t_{on} = \frac{(T.Q_{gd} + Q_{g,sw}) \cdot (T.R_g + inv.R_{on})}{T.V_{drv} - T.V_{plat}}, \quad (2.18)$$

Intuitively, this is the time it takes for a current supplied by the (approximately constant) voltage drop $T.V_{drv} - T.V_{plat}$ across the gate resistance to deliver the necessary charge $T.Q_{gd} + Q_{g,sw}$ to the gate. While $T.Q_{gd}$ and $Q_{g,sw}$ both may be read from the datasheet, it may be more accurate to calculate the Miller charge $Q_{g,sw}$ depending on how much the operating point in the design differs from that in the datasheet. This is also adapted from [87] as:

$$Q_{g,sw} = T.C_{iss}(v_{DS,eff-oss}) \cdot T.V_{plat} \frac{\frac{i_{on}}{g_m}}{\frac{i_{on}}{g_m} + T.V_{th}(T_J)}, \quad (2.19)$$

where g_m is the transconductance, read from the datasheet as $g_m \approx \Delta i_D / \Delta V_g$.

For t_{off} , a similar process is carried out:

$$t_{off} = \frac{(T.Q_{gd} + Q_{g,sw}) \cdot (T.R_g + inv.R_{off})}{T.V_{plat} - T.V_{th}(T_J)}, \quad (2.20)$$

Finally, turn-on and turn-off losses can be determined:

$$\mathbf{P}_{\text{turn-on}} = \frac{1}{2}t_{\text{on}} \cdot v_{\text{DS,eff-oss}} \cdot i_{\text{on}} \cdot f_{\text{sw}} \quad \text{and} \quad \mathbf{P}_{\text{turn-off}} = \frac{1}{2}t_{\text{off}} \cdot v_{\text{DS}} \cdot i_{\text{off}} \cdot f_{\text{sw}} \quad (2.21)$$

Note that, if C_{oss} is partially or fully discharged during turn-on, the area of the loss energy shrinks (shown by the lighter triangle in Fig. 2.7) and can be entirely eliminated. This is evident in Fig. 2.3b, where turn-on losses (and C_{oss} losses for that matter) are only seen in half of the line cycle – when the current is in the right direction to soft-switch at least half of the FCML devices. Additionally, it is worth noting that controlling (i.e., slowing) the turn-on transition is most crucial as that will always be a source of a fast, hard commutation that can lead to detrimental v_{DS} overshoot. However, turn-off is typically not as critical; $\text{inv.R}_{\text{off}}$ is often chosen to be zero or at least much smaller than inv.R_{on} . As such, turn-off losses are minimal.

Conduction losses dominate converter losses at high power due to the high output current. This is notable in Fig. 2.3b. For Si MOSFETs, this calculation has been the somewhat straightforward product of current squared and (a temperature dependent) on-resistance. However, for GaN devices, the notion a *dynamic* on-resistance, dR_{on} , has been introduced to capture the dependence on off-state voltage [77, 93, 94] and switching conditions [95–97]. While [96] presents a framework that can capture several aspects of operating point that influence impact of dR_{on} , such characterization is not available for a range of devices, and it is yet unclear whether such a complex model will be necessary for the proposed inverter operation. Conversely, [77] provides a simple methodology that leverages the general linear dependence of dR_{on} to off-state drain-source voltage. This approach incorporates the effects of temperature and off-state voltage through respective coefficients, T.tempCoeff and T.dynRonCoeff , that scale the nominal (datasheet) on-resistance R_{on} :

$$\mathbf{P}_{\text{conduction}} = \text{T.tempCoeff}(T_J) \cdot (\text{T.dynRonCoeff} \cdot \left(\frac{v_{\text{DS,eff}}}{\text{T.V}_{\text{DS,rated}}} \right) + 1) \cdot D \cdot i_{L,\text{RMS}}^2, \quad (2.22)$$

where T.tempCoeff is read from datasheet plots for a value of T_J , and T.dynRonCoeff may be determined from academic work, e.g., Fig. 9 in [77]. For the EPC2034 device, a coefficient of 1.154 is indicated and has been shown to significantly address modeling error that arises without considering dR_{on} . If future work determines stronger dependence on other parameters, such as off-time, additional coefficients as a function of, e.g., $D \cdot t_{\text{sw}}$ may be introduced into (2.22) as well.

2.2.2 Capacitor Losses

As noted previously, capacitors are the primary power processing component. Until recently, losses in these components have largely been attributed to a constant equivalent series resistance, or ESR. However, recent work has shown that certain ceramic capacitors have more complex and nonlinear loss mechanisms. Much of the literature focuses on applications in

power pulsation buffers, with comprehensive characterization of X6S [79, 98], X7R [99] and PLZT [100] for frequencies in the tens to hundreds of hertz. However, given that the flying capacitors in the FCML see only high-frequency currents (typically tens to hundreds of kilohertz), only [98] and [78] have appropriate data for the operating regime considered in this work. Additionally, only [78] also investigated the additional effect of voltage bias, known to affect rated capacitance, as discussed earlier.

The first step towards applying the data in [78] is to translate the operating conditions at each sample point on Fig. 2.3a into the duty ratio seen by each flying capacitor. Then, coefficients corresponding to RMS currents through the flying capacitors for the two switching frequencies can be calculated. These coefficients are used to derive RMS values from the nominal I_L and $i_{L,\text{pk}}$ at each sampled point on the output waveform, which affect components of the capacitor current at the switching frequency and effective frequency, respectively. Given that the switching frequency is much higher than the fundamental, the inductor current I_L can be assumed constant during each switching period. This implies that the current seen by each capacitor will be a rectangular pulse with a nominal amplitude of I_{out} and a pulse width that varies with D . The expression for the RMS capacitor current for such a waveform is given by:

$$i_{C,\text{RMS-}f_{\text{sw}}} = I_{\text{out}} \times \begin{cases} \sqrt{\frac{2D_{\text{eff}}}{N-1}}, & D \leq \frac{1}{N-1} \\ \sqrt{\frac{2(1-D_{\text{eff}})}{N-1}}, & D \geq \frac{N-2}{N-1} \\ \sqrt{\frac{2}{N-1}}, & \text{otherwise} \end{cases} \quad (2.23)$$

The second current component is at the effective (output) switching frequency. This is to compensate for the fact that, in addition to the output nominal value above, there is a superimposed ramp due to inductor ripple each time the capacitor is connected. This is treated separately here owing to the frequency dependent loss mechanism of the ceramic capacitors. Considering the ramp profile of the current at this frequency, a $\sqrt{3}$ associated the triangular waveform makes an appearance:

$$i_{C,\text{RMS-}f_{\text{eff}}} = i_{L,\text{pk}} \times \begin{cases} \sqrt{\frac{2D_{\text{eff}}}{3(N-1)}}, & D \leq \frac{1}{N-1} \\ \sqrt{\frac{2(1-D_{\text{eff}})}{3(N-1)}}, & D \geq \frac{N-2}{N-1} \\ \sqrt{\frac{2}{3(N-1)}}, & \text{otherwise} \end{cases} \quad (2.24)$$

While modeling frameworks using Steinmetz parameters are in development for ceramic capacitors [78, 98], the accuracy of these relations in the regime of interest for this work is

still under investigation. Therefore, two dimensional lookup tables and linear interpolation were used to directly leverage data collected in [78] for the ESR at a given bias voltage, V_{cap} , and frequency. The power dissipated in a given flying capacitor can then be calculated as:

$$P_{C_{\text{fly}}} = \sum_{k=1}^{N-2} i_{\text{RMS},C-f_{\text{sw}}}^2 \cdot C[k] \cdot \text{esrFn}(f_{\text{sw}}, V_{\text{cap}}) + i_{\text{RMS},C-f_{\text{eff}}}^2 \cdot C[k] \cdot \text{esrFn}(f_{\text{eff}}, V_{\text{cap}}), \quad (2.25)$$

where k is the index of a specific flying capacitance.

2.2.3 Inductor Losses

The inductor losses in this work are treated similarly to most extant modeling studies. These losses arise from both resistances in the copper windings as well as other phenomena in the magnetic media itself. Ample literature is available on general methods for modeling losses in the conductors and magnetic material [101–103], including the fundamental application of Steinmetz parameters [104–108]. These works provide a useful foundation for loss calculation and optimization, especially when working with custom magnetics, where winding and core geometries are known.

Although custom magnetics provide several degrees of freedom for the designer, this work seeks to leverage economical, off-the-shelf, discrete inductors to simplify assembly and drive down system costs. This leads to two caveats: 1) the design space is discrete – only existing part numbers with specified mass, current and loss profiles can be selected, and 2) the loss model is only as detailed as the manufacturer wishes it to be. Fortunately, major suppliers provide this information in some sort of equivalent form [109, 110], with core losses generally expressed using the Steinmetz equation:

$$P_{L\text{-core}} = K \cdot f^\alpha B^\beta \quad (2.26)$$

where K, α and β are constants from fitting empirical loss data against known conditions of the excitation frequency, f , and peak flux density, B . Additionally, some manufacturers appear to include the conversion from inductor current to flux density (i.e., $B := I$) into the parameter fitting – as appears to be the case for the discrete inductors chosen for this work.

2.2.4 Efficiency Calculation

Ideally, the efficiency of a converter under study can be calculated from the output power divided by the input power, i.e. the output power plus the losses described above. However, the output voltage and current waveforms generated in Fig. 2.3 were under ideal conditions assuming a transfer function of a lossless converter. In the real system, the converter losses – especially conduction losses – will introduce a voltage drop across the converter and decrease the output power delivered to the load. In other words, the output voltage in physical hardware is lower than what the model, as defined, can predict. This difference in experimental and modeled voltages is plotted as the gray series in Fig. 2.8a over various converter

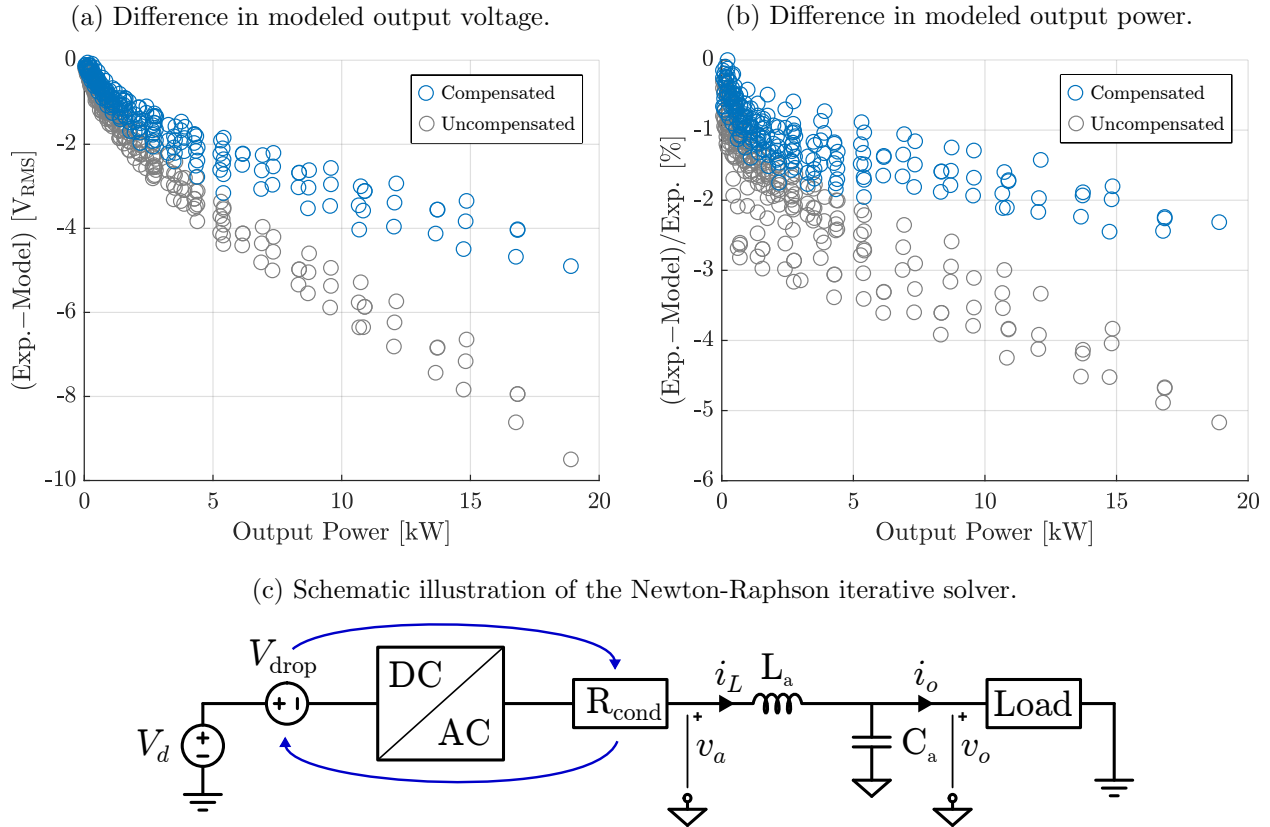
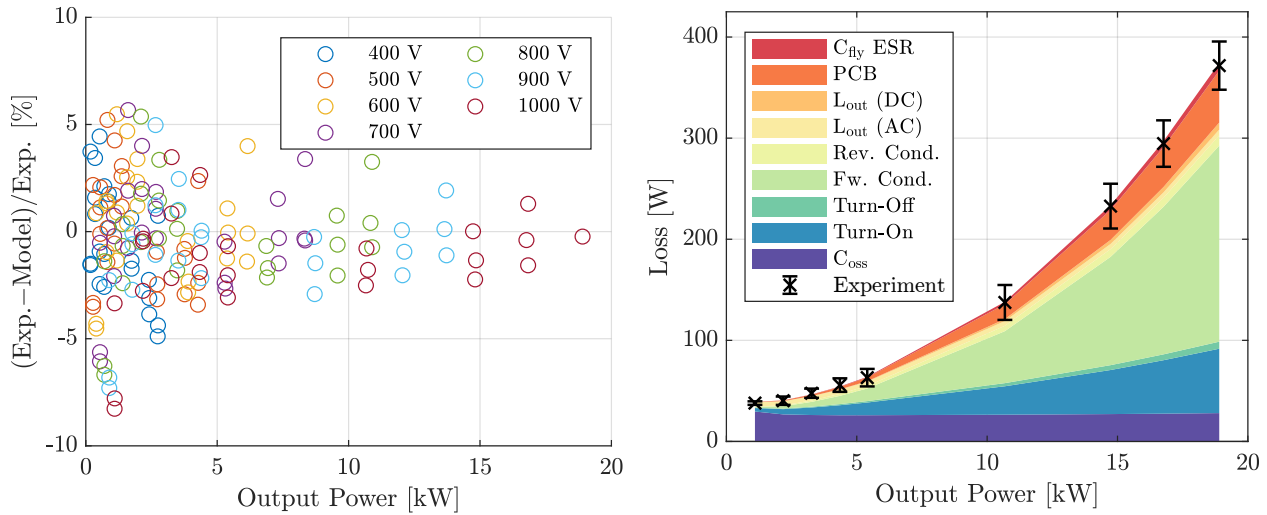


Figure 2.8: Comparison of loss models without (gray series) and with (blue series) an iterative solver that compensates for converter losses in the output power and efficiency calculations.

operating points, with greater voltage deviation accompanying greater converter loss. The corresponding decrease in output power is shown in Fig. 2.8b. Therefore, failing to address this issue leads to an inaccurate loss model.

Fortunately, this problem can be addressed numerically. With the understanding that a voltage drop will occur largely due to conduction losses, a Thevenin resistance was defined by this loss divided by the output current. The drop across this term was then subtracted from the nominal voltage used to determine output waveforms on the next iteration, thereby compensating the generated waveforms and output power. A Newton-Raphson iterative solver, graphically illustrated in Fig. 2.8c was implemented to arrive at the compensated series shown as the blue plots in Fig. 2.8a and Fig. 2.8b. This approach was able to dramatically improve the accuracy of the model, with remaining error due to unmodeled behavior of the test fixture, and some other dynamics not captured using simplified loss models. Note, though junction temperatures were known to reasonable accuracy for model validation in this work (see Chapter 3), this method can also be used to compensate for the temperature dependence of losses: if the thermal impedance of the proposed cooling system is known, the losses can again be iterated over to find the operating temperature of the device.



(a) Ensemble loss modeling error (with compensation). (b) Experimental and modeled losses at 1 kV.

Figure 2.9: Modeling accuracy validated across various operating points.

To evaluate the robustness of the loss model, the converter described at the beginning of this chapter was swept across various operating conditions, including dc bus voltage, output power, load, and switching frequency. The experimentally measured losses were recorded along with the exact operating parameters for each experiment. Then, the loss model was simulated at each operating point and compared to the observations. The plot of error between experiment and simulation is shown in Fig. 2.9a for the ensemble of measurements. Here, it is evident that good agreement is observed (less than $\pm 5\%$ error) across most operating points, with remaining differences owed to uncertainty in component parameters and dynamics not modeled (e.g., analytically exact turn-on/-off transients or flying capacitor voltage fluctuations) in this simplified approach.

For operation at the peak dc bus voltage rating – the most informative operating point – the model agreement and loss breakdown is especially important. Fig. 2.9b shows the 1000 V dc operation as the load is increased until the converter reaches peak power at 18.9 kW. Here, the experimental performance (with associated measurement error) is plotted alongside the model losses at each loading. Notably, very good agreement is observed across the range, with the exception being some additional mismatch at light load – where minute differences in parameters affecting switching losses can have a larger impact. A heuristic compensation factor was used in this work to improve agreement, while future work might examine instantaneous voltage balance across modulation for more accurate results. Another key takeaway is that forward conduction is the dominant loss mechanism – emphasizing the need to both minimize on resistance in the face of dR_{on} , as well as effectively cool these tiny devices. Turn-on losses are also significant, indicating that any design improvements leading to faster device turn-on without adverse affects from commutation loop parasitics (see Section 2.4) will continue to produce a pronounced reduction in converter losses.




Figure 2.10: Startup simulation in PLECS.

2.3 Startup Dynamics

While the FCML topology provides many desirable benefits, operation under transient conditions deserves specific scrutiny. In addition to dynamics associated with motor acceleration and failure modes, start-up is regularly identified as a major concern. For this application, it is assumed that i) logic power will be available before the dc bus is energized and ii) that the bus will not be energized to the full voltage instantaneously, but will ramp up at a given slew. Breakers are proposed for both the dc and ac side of aircraft propulsion drive systems [37], and a conceivable configuration might also include a temporary connection through a pre-charge resistance. Furthermore, startup slew for some applications through such means has been identified as plausibly on the order of 100-500 ms [111].

Much like the topology analyzed in [111], this converter also features an LC output filter that can double as an “external balance booster” [63]. By operating the converter at 0% modulation and a frequency that leverages this filter during startup, the self-balancing startup approach described in [112–114] provides a simple means to charge the flying capacitors. As much of the flying capacitance varies with voltage, exact analysis of startup dynamics using the analytical methods of these works is difficult. However, a simulation (with the non-linear capacitances) of start-up is plotted in Fig. 2.10 and shows balanced pre-charging can occur quite quickly, with all transistors staying within 75% of their rated voltage during the transient. This startup behavior is experimentally validated in Section 2.5.2. Note, the frequencies of oscillations present in the simulated startup waveform can be analytically determined [115], however determination of the actual damping of the circuit is only possible with adequate knowledge of the high-frequency ac resistance within the circuit [116].

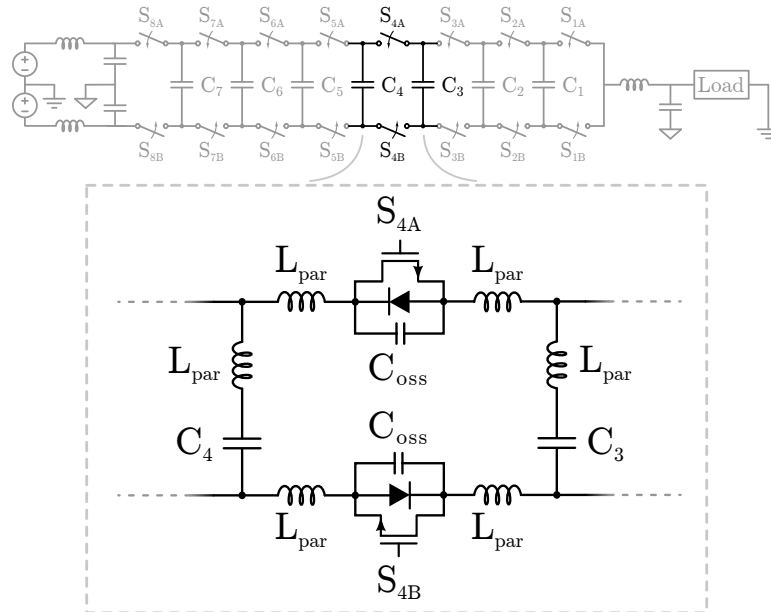


Figure 2.11: Commutation loop between a given pair of complementary switches in a single 9-level FCML inverter, shown here for a 9-level converter.

2.4 Mitigation of Commutation Loop Inductance

Much like the two-transistor half-bridge, the FCML topology is built from pairs of complementary transistors. In the schematic of Fig. 2.11, S_{4A} and S_{4B} represent two of these complementary pairs. Also like the half bridge, parasitic inductances in the commutation loop arising from PCB layout similarly hinder performance; during fast transitions, energy stored in this inductance can lead to overshoot in v_{DS} to the detriment of EMI or device ratings, while slowing or damping the commutation (typically done by increasing the gate resistance) would lead to additional turn-on loss (as noted in Section 2.2.1). The GaN switches in this work enable designs that are simultaneously low loss and low weight due to their high-speed switching capability. However, leveraging these devices in FCML converters requires mitigation of the parasitic inductance [117]. Converter designs using a lateral layout (i.e. all switching devices are on a single PCB side) do so by employing local bypass capacitors that effectively short the commutation energy through a low inductance path [70, 118]. Lateral layout converters are easier to assemble and debug, with most components confined to a single side.

However, a vertical design – where complementary components are mirrored on opposite sides of the PCB [67, 73, 80], as illustrated in Fig. 2.12 – can improve thermal management and high-voltage capability. For instance, dual-sided cooling is straightforward, while the high-voltage standoff capability of the PCB substrate helps address clearance requirements. While local bypass capacitors are precluded in this configuration, the use of internal copper layers and via stitching can effectively null much of the inductance arising from the com-




Figure 2.12: Cross section showing commutation loop between a given pair of complementary switches. Internal layers and vias create an electrically-thin vertical layout on a single PCB; the cross-sectional area inside the PCB (shaded red), and thus the corresponding loop inductance, can be greatly reduced with this approach.

mutation loop passing through the converter PCB. This “electrically-thin” approach [119] reduced the loop inductance of [73] from 8 nH to 2.7 nH. However, as that design used “switching cell” daughter boards soldered to the main PCB, both blind and buried vias were necessary to create the electrically-thin routing without overlapping mounting pads.

This work uses a vertical layout but eliminates daughter boards, and instead places the switching cells on the main converter PCB. A custom 6-layer stackup and buried vias are then used to create the electrically-thin path within this PCB. This custom stackup requires the thickness of the prepreg between the outer copper and first inner layers to be as thin as possible, as illustrated in Fig. 2.12. The general intuition is that all current paths should be routed to tightly couple (i.e., to be in close proximity) with their corresponding anti-parallel return paths. Eliminating the daughter boards does remove an element of modularity, but improves manufacturability, power density (less substrate material), and reduces loop inductance by decreasing path lengths.

2.4.1 Parameter Extraction using Finite-Element Analysis

To analyze the benefits of the above approach, finite-element analysis was conducted in Ansys Q3D. Trace and pad shapes were imported into Autodesk Inventor, where three-dimensional geometries with thicknesses corresponding to copper and laminate specifications in the stackup were defined. While Ansys does have a toolchain to import PCB files directly, many vias were used to stitch the inner layers, and their inclusion would have created a complex body to process. Instead, via stitching was simplified by forming a solid copper mass where the vias connect layers. Previous work noted that this does not decrease accuracy significantly; the resulting discrepancy is on the order of the precision at which the physical system can be experimentally measured [62].

The model of the switching cell is shown in Fig. 2.13. To reduce the computational domain, only a subset of the volume around a single switching cell is within the simulation geometry. Additionally, extraneous copper from adjacent gate drive circuits is omitted for simplicity. While there is likely some coupling into these traces, inclusion in this analysis




Figure 2.13: 3D model of a switching cell with buried vias, as imported and meshed in Q3D.

would likely only be of interest for EMI purposes and have little effect on the commutation loop concerns of this section. Additionally, the multilayer ceramic capacitors are treated as ac short circuits (though the pads for capacitors on adjacent cells are shown for illustrative purposes). While the physical real will have a more nuanced interaction with the magnetic fields present in the commutation loop, the short circuits allow the analysis to be approximately limited to trace geometries and routing for the purposes of evaluating switching cell design. However, the lumped models of the ceramic capacitors may be added in series with the relevant segments of the computed model for a more practical number. Finally, to ensure the current loop is correctly closed, the bottom transistor is replaced with an ideal short circuit while the top switch is left open, as indicated in Fig. 2.13. This way, a source and sink stimulus can be established on the respective left and right edges of the two, adjacent top copper layers in the boxed region. The red and blue arrows are used to indicate the general flow of commutation current from the source to the sink of the electrical net, matching the convention in Fig. 2.12.

This analysis was solved at a frequency of 100 MHz, which is the order of magnitude matching the frequency of oscillation between device capacitances and commonly observed loop inductances. The primary loop inductance from the drain of the top switch, through the PCB and to the source of the top switch, was calculated to be 0.939 nH. For the two sets of two series, four parallel ceramic capacitors serving as flying capacitance for most of the switching cells, the lumped component inductance amounted to an additional 0.65 nH [120]. Thus, an overall inductance of 1.59 nH was expected for this design, representing a further substantial decrease in parasitics for the state-of-the-art.

2.4.2 Simulation of Commutation Loops in SPICE

The inductance reported above represents a lumped quantity; However, the finite-element analysis also generates a full impedance matrix for any number of ports defined in the model.

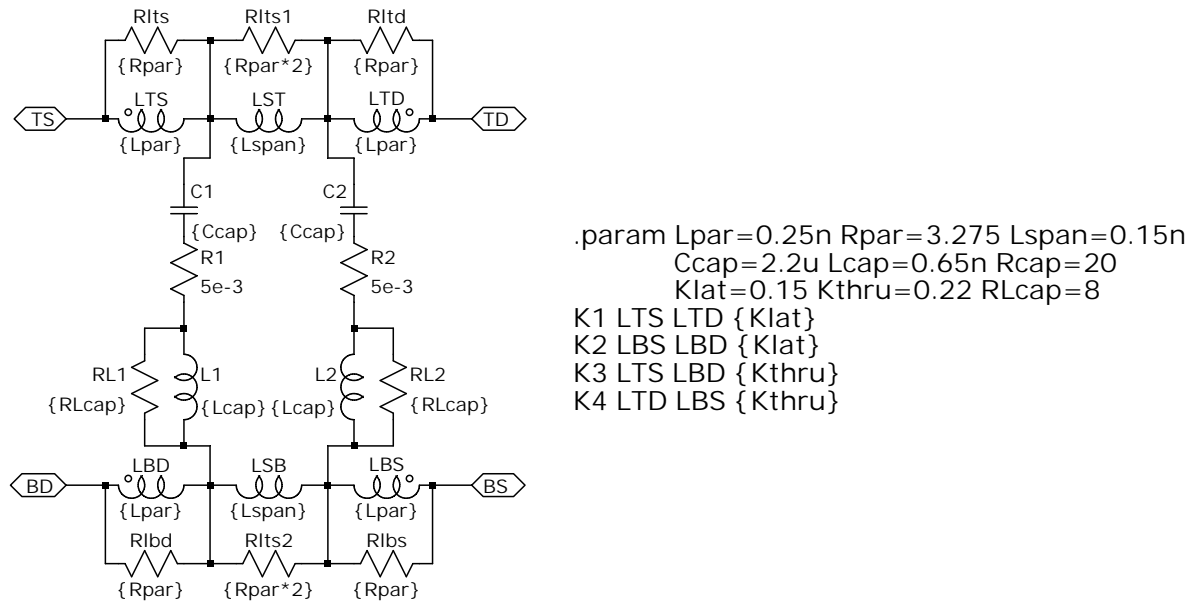
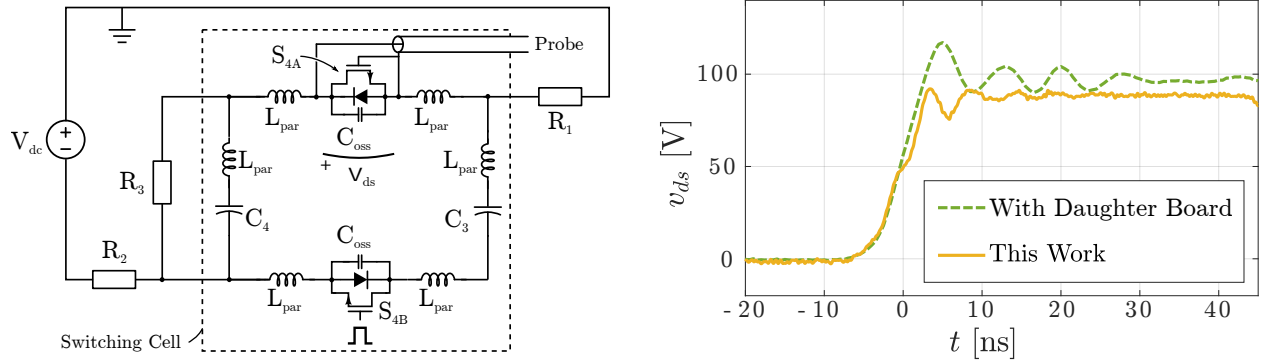


Figure 2.14: Spice model of commutation loop implemented as a component in LTSpice using parameters extracted from Q3D finite-element analysis.

Additionally, ac resistances associated with each current branch are also calculated, and useful for informing the damped behavior at high frequencies. This is valuable when simulating the effects of the commutation loop, as using the lumped inductance alone produces an underdamped response when the only circuit resistances are the nominal on-resistance of the transistors and the ESR of the capacitors. Additionally, it may be useful to examine the inductance between adjacent switching cells, especially when operating under high module output current. Therefore, a multiport equivalent model was also extracted.

Though not always ideal for simulating switching circuits, Spice is the appropriate platform for nodal analysis involving these parasitics. In Fig. 2.14, the equivalent model for half of a switching cell (one flying capacitor) is represented in LTSpice schematic format. The values in the `.param` declaration are mostly computed and transferred from Q3D directly, specifically the inductance and the coupling terms (with prefix `K`). Here, rather than defining the unit cell to contain the two complementary devices and the adjacent flying capacitors, as illustrated in Fig. 2.11, only the path through a single flying capacitor is modeled. Choosing this configuration as the building block allows the model to be easily repeated between each switch pair for the whole converter. On the contrary, choosing the unit cell as illustrated in Fig. 2.11 would only correctly capture the dynamics of every other switching cell. Furthermore, both intuition and the results of finite-element analysis indicate stronger coupling between copper traces on a single side of the devices (in parallel planes) versus those on the opposite side (also parallel, but with significant offset). As such, the coupling terms corresponding to the latter are safely omitted in this model.

Finally, it is important to note that the ac resistances calculated in the finite-element



(a) Test setup for measuring overshoot. A coaxial connection mitigates high-frequency pickup.

(b) Results for the proposed commutation loop, showing greatly reduced drain-source overshoot.

Figure 2.15: Experimental characterization of drain-source voltage overshoot during commutation as a means of comparing loop parasitics between designs.

element model are dependent on frequency, whereas resistances modeled in spice are not. To emulate this effect, a series-parallel transformation was used to add this damping resistance in parallel with the partial inductances to provide the desired resistances at high frequency. The step-response of the system is thus damped as expected, but without the use of excessive series resistances.

2.5 Experimental Results

2.5.1 Commutation Loop

The hard-switched overshoot, arising from the loop inductance oscillating with the capacitance of the EPC2034 GaN FET of this design, was obtained using the circuit shown in Fig. 2.15a. Resistors R_1 and R_2 , both $470\ \Omega$, isolate the commutation cell from the parasitic capacitance of the voltage source, while R_3 ($33\ \text{k}\Omega$) provides a path for a low but non-zero current ($\sim 4\ \text{mA}$) to bias the complementary switch (S_{4A}) in the on-state prior to the turn-on of S_{4B} . An ultraminiature coaxial connector was affixed in tight proximity to the drain and source terminals of S_{4A} and, when used with a Keysight N2894A 700 MHz high-bandwidth passive probe and associated coaxial accessories, provided a high-fidelity measurement of v_{ds} .

Careful measurements of the specified EPC2034 GaN FETs and a $22\ \Omega$ gate resistance produced the hard-switched v_{ds} waveforms in Fig. 2.15b, shown near the nominal operating point of a 10-level design with a 1 kV dc bus. Here, the overshoot of v_{ds} in this work is drastically lower than that of the previous electrically-thin design under comparable driving conditions. To best leverage this improvement, the gate resistance was set to $15\ \Omega$ (compared to the $22\ \Omega$ indicated in [73]) for faster device turn-on and reduced overlap losses.

Direct measurements of the loop inductance with capacitors in stalled were also performed using a high frequency impedance analyzer are available in Table 2.4. The test frequency

Table 2.4: Experimentally measured commutation loops for various switching cells.

S_1	S_2-S_7	S_8	S_9
2.76 nH	2.57 nH	2.34 nH	2.3 nH




Figure 2.16: Impedance analyzer and probe used to measure the commutation loop inductance. After calibration is carried out to the adapter tip endpoint (inset), the loop impedance between the drain to source pads can be measured with high accuracy..

of 100 MHz was chosen to both assist in measuring the very small inductances, as well as match the order of magnitude at which hard-switched ringing is commonly observed. The inductance measurement setup is shown in Fig. 2.16, with a Keysight E4990A impedance analyzer and 42941A impedance probe; the inset photo shows how the tight coupling of the probe tips allows for precise measurement of the circuit while minimizing the impact of the test setup. Notable error between the simulated results arises from the simplifications in simulation geometry, discrepancies in stackup thicknesses, and behavior of the physical ceramic capacitors. It is also worth noting the increased decreased inductance of the S_8 and S_9 loops, which have many more paralleled flying capacitors – reducing their combined effective series inductance.

With all transistors unpopulated, the probe tips are connected to corresponding drain and source pins of one of the complementary pads for a given switching cell. The opposite device is then shorted with thin, copper foil soldered across the corresponding source and drain pads. Then, the resulting impedance measurement produced the approximate contribution from the routing (if flying capacitors are shorted, as shown in the inset of Fig. 2.16) or the total loop inductance – including non-negligible contributions from passives (if the flying capacitors are populated).

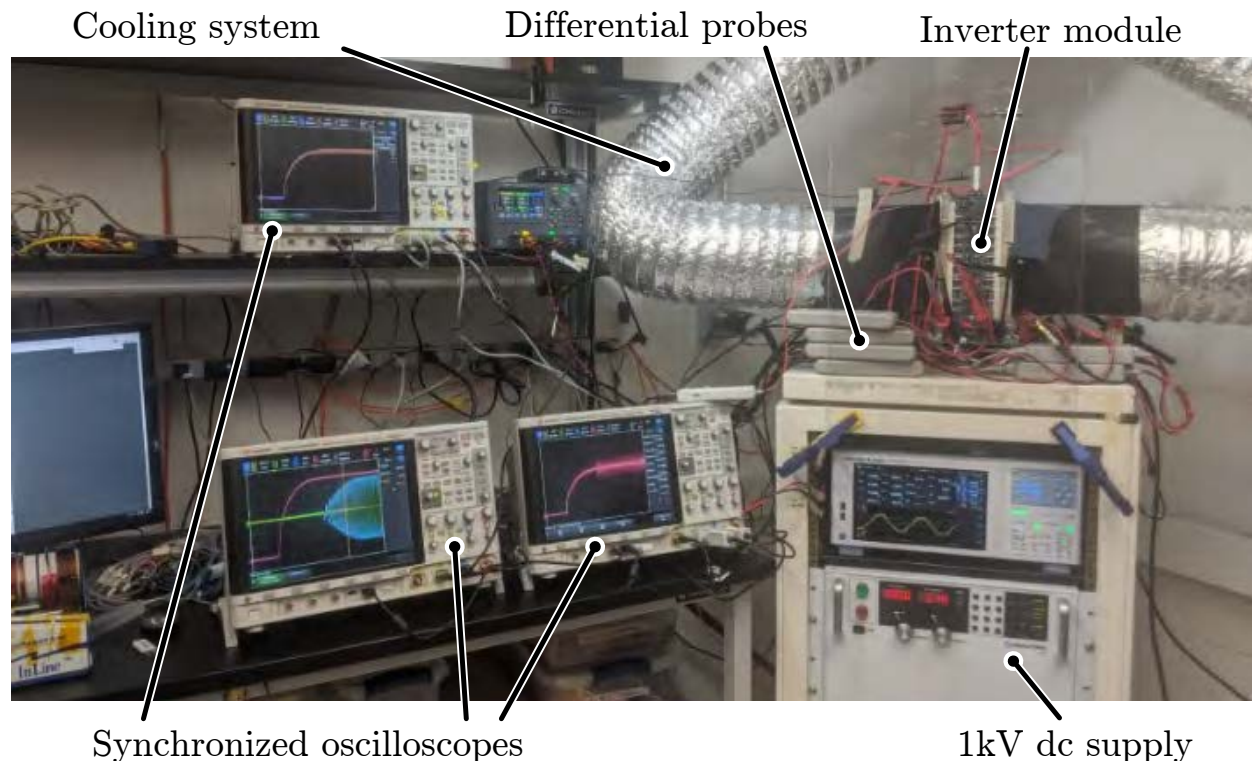


Figure 2.17: Experimental setup for 1kV dc bus startup on a 10-level FCML inverter module. Waveforms were captured on three synchronized Keysight MSO-X 4024A oscilloscopes with several models of high voltage differential probes of 25 MHz bandwidth or greater.

2.5.2 DC Startup

Observing startup on a 10-level converter – which amounts monitoring 8 high voltage flying capacitors and a thousand volt dc bus is quite challenging. Each capacitor represents a floating voltage source with respect to earth ground, and accessing each electrical node is challenging when designing for high power density. Fortunately, the surface mount nuts adjacent to each transistor (see Chapter 3) provided electrical test points. Nonetheless, a total of 10 high-voltage differential probes were required to measure each capacitor voltage, the dc bus voltage and the output voltage waveform, as shown in Fig. 2.17. As each available oscilloscope provided a maximum of four channels, three independent oscilloscopes were synchronized against a common trigger to capture the full startup dynamic. While a future design may incorporate analog sense circuitry at either the global or local (module) level (see Chapter 4) to detect the dc bus state, for this experiment, dc bus startup was detected using an auxiliary oscilloscope that provided an external trigger for the remaining three oscilloscopes once the dc bus rose to a certain threshold.

To experimentally evaluate this proposed startup procedure described in Section 2.3, the converter of this work was programmed to power-on at a switching frequency of 20.4 kHz (approximately half of the resonant frequency of the output filter). Next, the high voltage

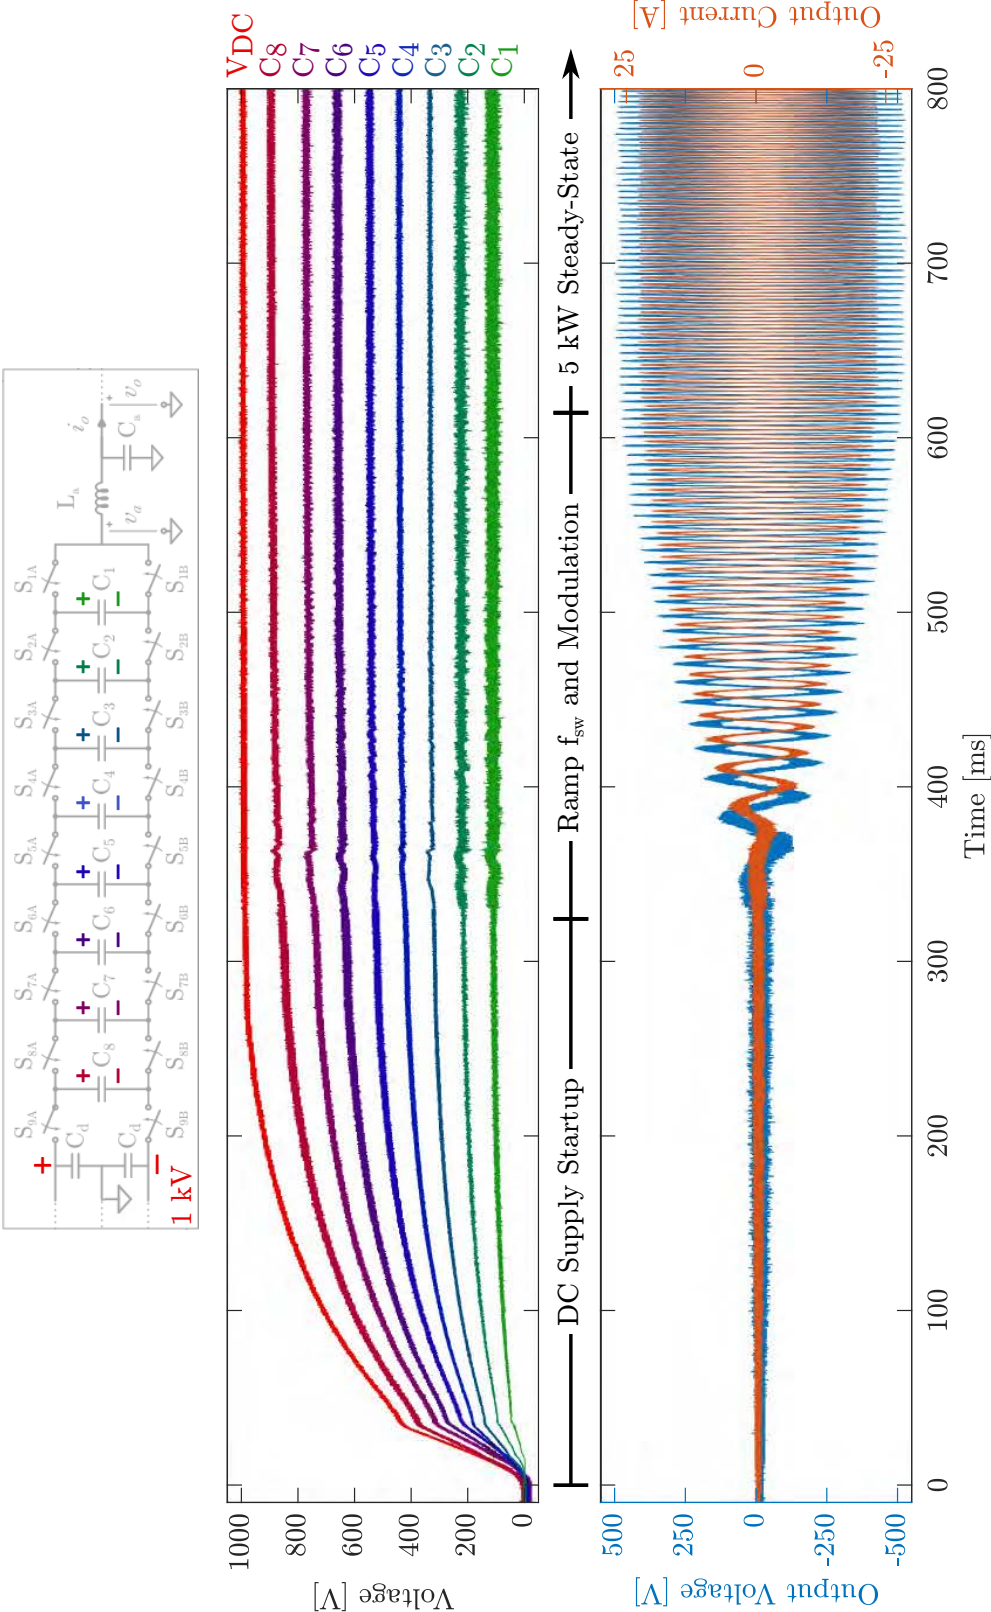


Figure 2.18: Startup of the dc bus from 0 to 1kV, after which output power is ramped to 5kW. Waveforms were captured on three synchronized Keysight MSO-X 4024A oscilloscopes with several models of high voltage differential probes of 25 MHz bandwidth or greater. High-frequency content present on a given waveform can be attributed to the specific probe used, or the higher fidelity of ac content capture for the signals with lower dc voltage.

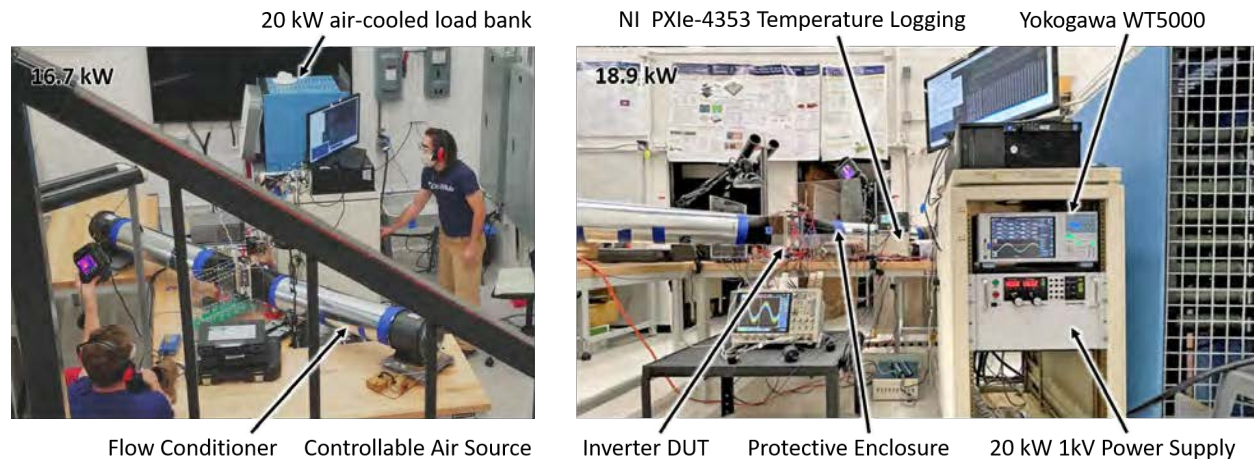


Figure 2.19: Experimental setup for high power module tests.

supply was commanded to ramp to a 1 kV set-point at maximum slew. Finally, once the dc bus was detected to settle near the rated value, the controller was commanded to increase the switching frequency to 115 kHz for steady state operation and ramp both modulation frequency and amplitude to a target power level. The resulting behavior of the dc bus and capacitor voltages, along with the output voltage and current, are shown in Fig. 2.18. Here it can be seen that the dc bus can power up and the converter can safely balance flying capacitor voltages, as well as ramp output power, in less than 500 ms (compared to [111], this is more than sufficient). Additionally, while the Magna Power TS Series supply used for this project could not be made slew to full voltage faster than shown, a fast initial slew of 10 V/ms is exhibited. Even balancing during this interval indicates this approach may indeed be leveraged for even faster startup times.

It is worth noting that startup to the full 1 kV was only attempted after several successful tests conducted at lower dc bus voltages. Furthermore, to avoid any unexpected interaction between phases, interleaving was disabled for these experiments. However, this is not expected to significantly impact startup dynamics. As it were, interleaving can be dynamically re-engaged after startup – similar to the way the switching frequency is ramped from the lower startup configuration to the nominal 115 kHz.

2.5.3 Module Efficiency and Power

The main deliverable of the approach set forth in this work is a high-efficiency and high-power-density inverter module. As such, experimental results across the target operating points are necessary for both validating the loss models discussed earlier and demonstrating the competitiveness with the state-of-the-art. Although the instrumentation requirements are less cumbersome than those of the start-up experiment in Section 2.5.2, the high-voltage and high-power test setup shown in Fig. 2.19 nonetheless required careful design and con-




Figure 2.20: 10-level converter performance at 1kV dc bus with varied switching frequency.

figuration. High-voltage, dc power is supplied by a water-cooled, Magna Power TS Series 20 kW, 1000 V supply. A $0.5\ \Omega$ power resistor was inserted in series to both rails to attenuate spurious noise or ripple from the supply. The output from this filter was run through one 30 A, high-power element of a Yokogawa WT5000 to sense the current of the input power measurements. Then, a split-bus was generated using a capacitor-divider across the dc-link, with the midpoint tied to earth ground. Correspondingly, all tests at 1000 V dc correspond to ± 500 V.

This dc connection supplies power to a backplane that can support up to 15 parallel modules for full power, three-phase operation. While only a single module (and thus single phase) was tested for the performance measurements here, Chapter 4 has additional results for three-phase operation with nine paralleled modules. Nonetheless, the backplane serves a vital role in delivering dc power to the card-edge ports on the single module under test (which is the point where the Yokogawa measures the module input voltage), as well as providing mechanical fixturing and 5 V logic power. This backplane also provides additional film capacitance for high frequency decoupling. However, the midpoint of this capacitance does not return to the earthed midpoint to address potential conducted emissions concerns.

The module itself was instrumented with K-type thermocouples to both spot check transistor temperatures as well as record operating conditions for comparison with electrical and thermal models; further details on this measurement are discussed in Chapter 3. For these tests, unlike the startup experiment, it was impractical to instrument each of the flying capacitor voltages for monitoring. However, the two switched-node voltages were recorded using high-voltage differential probes referenced to the dc midpoint, and the degree of balancing can be inferred from this measurement. Additionally, the filtered output voltage, phase currents and total output current were also recorded. Due to the high amplitude (over 40 A) and high frequency (over 1 MHz) content of interest in the current waveforms at the phase outputs, Rogowski coils were used for all of the current measurements.

The ac output of the module was routed to two additional 30 A elements in the Yokogawa for the current measurement of the output power calculation. Here, two elements needed to

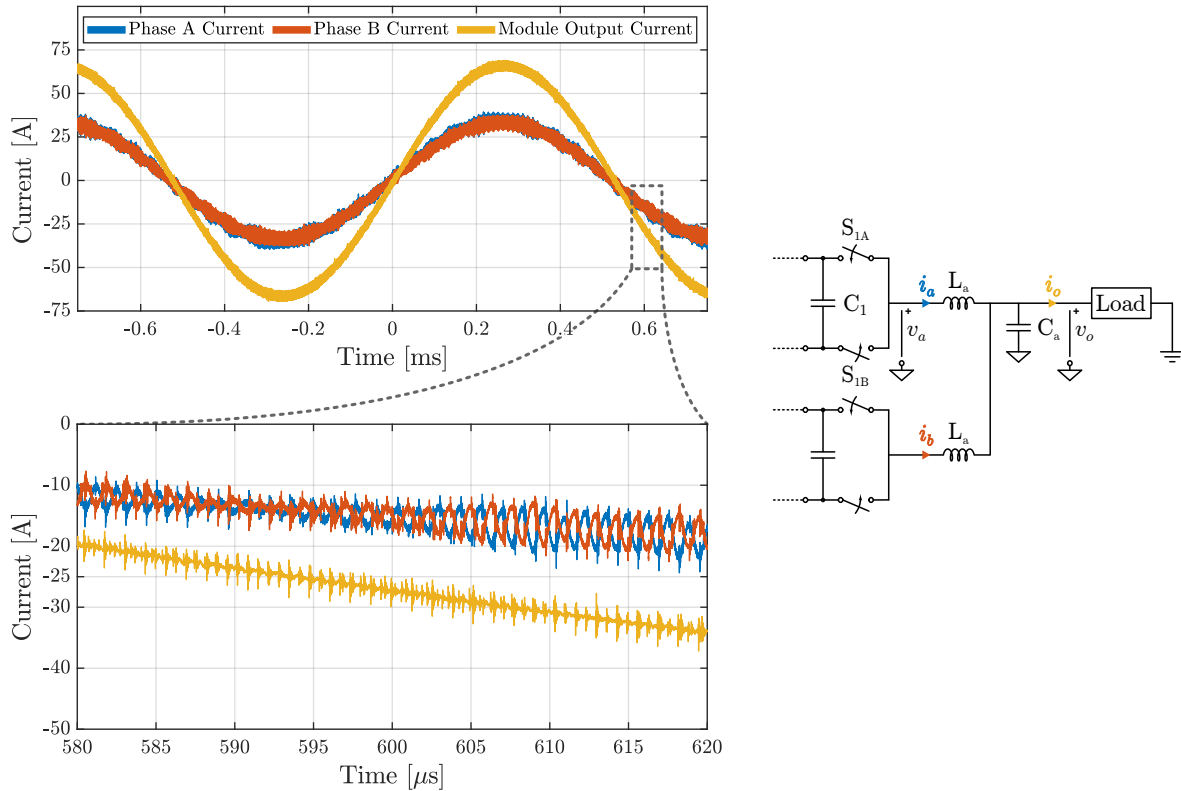


Figure 2.21: Experimentally-measured interleaved output current waveforms at 16.7kW output power. Waveforms are captured digitally on an oscilloscope and plot to enhance details.

be paralleled as the peak output current of the module would exceed the rating of a single element. This, in turn, fed a high-power, air-cooled load bank built from up to eighteen $100\ \Omega$ resistors [121]. Each resistor is rated for 1500 W, allowing for reasonable overhead for 20 kW-scale tests. As the high-power tests in this section were conducted using open-loop modulation, dc content in the output was not actively controlled. While such currents would likely be small in magnitude, the resistive load bank was connected in series with approximately 3 mF of ac coupling capacitance to block any dc path through the load. Finally, the output of this capacitance was returned to the dc midpoint provided by the split-bus capacitor. The voltage for the output power calculation was measured from the output terminal of the module to the midpoint return connection just after the ac coupling capacitors.

The output power was varied by changing the modulation index and by increasing the load through connecting additional resistors in the load bank. While the dc bus voltage was also swept, tests at voltages other than 1 kV were primarily to validate the modeling of different loss mechanisms. Therefore, only operation at the rated voltage of 1 kV will be highlighted in this section. A summary plot of converter performance is shown in Fig. 2.20. Here, the dc voltage was kept constant and power was statically varied by increasing the resistive loading

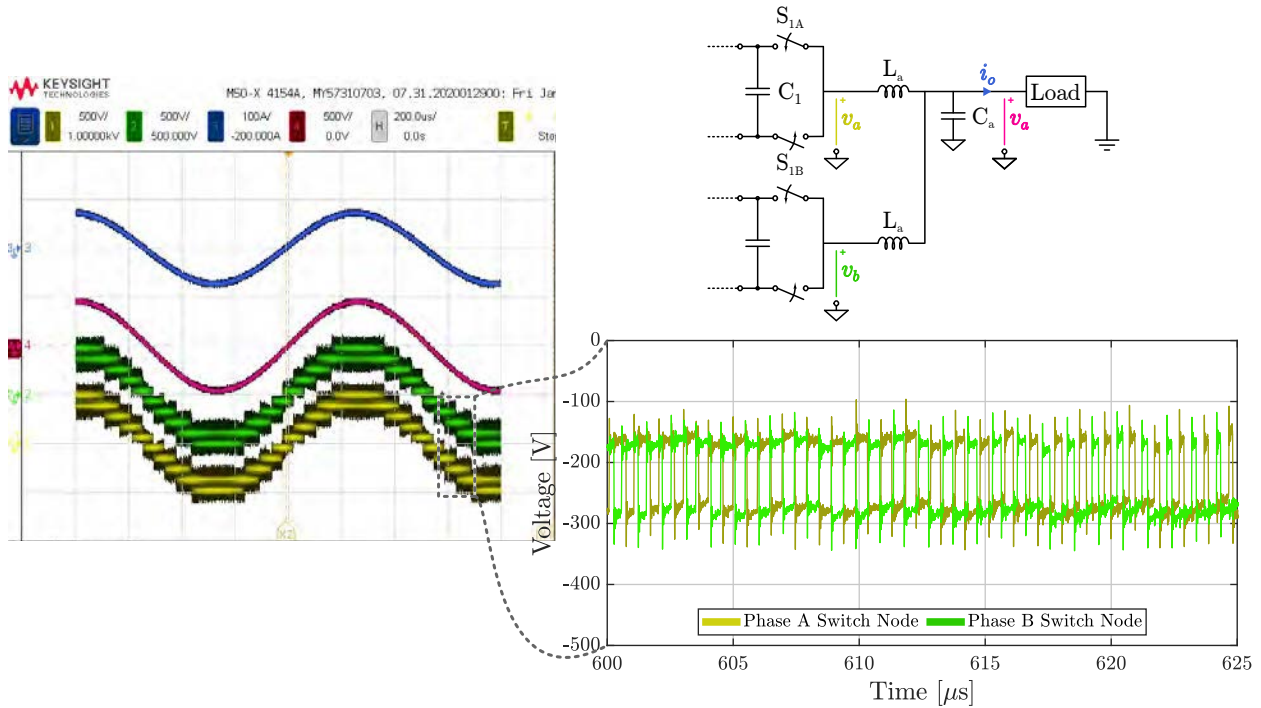


Figure 2.22: Interleaved output voltage waveforms at 1 kV dc bus and 16.7 kW output power. Waveforms are captured digitally on an oscilloscope and plot to enhance details, but are shown next to the oscilloscope screen capture for additional context.

between measurements. Additionally, three switching frequencies were tested to examine the trade-off between efficiency and voltage margin on the flying capacitors and devices. The peak performance trend mostly tracks operation at the lowest switching frequency (with the lowest switching losses) until just above the 15 kW operating point. Then, there is a clear point where the higher switching frequency is preferable. This change is likely due to the higher RMS currents associated with the lower switching frequency, though the increase in peak drain-source voltage magnifies the effect of the dynamic on-resistance at these operating points. Nonetheless, this plot highlights the impressive peak output power of 18.9 kW and peak efficiency of 98.95% – just shy of 99% after including logic and gate drive losses.

In addition to the power and efficiency measurements above, the output waveforms shown in Fig. 2.21 and Fig. 2.22 demonstrate both the general operation of the FCML, as well as the robust implementation in this work. While Chapter 4 provides a detailed analysis of interleaving many converters, Fig. 2.21 highlights how just dual-interleaving can significantly reduce switching harmonics in the output; although significant ripple is present in each phase, the summing at the output leads to nearly sinusoidal current delivered to the load. This is further affirmed by total harmonic distortion (THD) measurements by the Yokogawa: even with the highest current ripple, THD remained under 0.7% – well under the 2% to 5% typical for converters at this target power level.

Finally, Fig. 2.22 shows the unfiltered, switched node voltage for both phases. Here, the relatively even levels of the high and low edges of the respective waveforms indicates that the flying capacitors remain balanced, even at high output power. Note, at high output current the switch node waveforms are not perfectly rectangular due to high of charging and discharging of the relatively small flying capacitors value of capacitance necessary for operation. Regardless, operating conditions remained within acceptable limits.

2.6 Chapter Remarks

While Section 1.3 outlined the promises of a modular approach using the FCML topology, realizing the extreme performance in hardware can be challenging. For instance, to determine the optimal component selection and converter configuration, the loss mechanisms of the converter need to be well understood yet modeled in a way that permits rapid computation and iteration. This chapter demonstrated a method of capturing complex loss mechanisms through tractable functions. The resulting model showed good agreement with experimental results across a range of operating conditions, indicating applicability to a broad range of configurations in future design studies.

Nonetheless, this chapter also illustrated the dependence of loss models on component parameters that are still largely uncertain. Dynamic on-resistance in gallium nitride remains a challenging phenomenon to account for. The literature shows that the effect varies significantly across devices and operating points, while proposals on the best way to account for this variation are still evolving. Manufacturers release little to no characterization of this effect across model numbers, while experimental characterization in the literature usually focuses on a singular part number (with the exception of [77]). Similar concerns exist with power passives: inductors and capacitors are typically only subject to small-signal characterization by the manufacturer, while operating at high power often demonstrates other loss mechanisms that are not well-represented through small-signal characterization alone. Therefore, it is apparent that to be truly confident a design lies on the Pareto-optimal front, the designer might likely need to embark on a parallel, component characterization effort to fully understand the available materials.

In spite of the evolving understanding of these emerging devices and materials, this work leveraged the component characterization that is available in the literature to push the frontier of extreme performance power conversion, with key results summarized in Table 2.5. For perspective, Fig. 2.23a shows the design of this work compares with competing designs against the NASA performance targets for more electric aircraft [24]. While some aspects

Table 2.5: Key performance metrics for the design of this work.

Peak Efficiency	Peak Power	Grav. Power Density	Vol. Power Density
98.95%	18.9 kW	38.4 kW/kg	24.4 kW/L



(a) Comparison of this work (Gen 4) performance against NASA electric aircraft targets [24] and competing designs.

	Time [min]	Power [kW]	Power/ILM [kW]	Efficiency [%]
Takeoff/Climb	0	250	16.7	98.36
	1	250	16.7	98.36
	5	175	11.7	98.80
	10	125	8.3	98.93
	15	100	6.7	98.93
Cruise	20	83	6.9*	98.93
	300	83	6.9*	98.93

(b) Theoretical performance on ARPA-E ASCEND mission for an array consisting of 12-15 modules. An asterisk denotes module shedding to maintain peak efficiency at light load.

Figure 2.23: Summary representation of how the design of this work meets aggressive performance targets for more electric aircraft compared to the state-of-the-art.

of the system balance in a commercialized system will reduce power density below what is presented here, it is worth noting that leveraging liquid cooling – as some of these designs already do – will advance the power density of this work by at least an order of magnitude. In addition to excelling against the general NASA targets, the work also excels on the more recent mission targets set forth by ARPA-E. A full mission profile for a future electric aircraft is shown in Fig. 2.23b, where an array constructed with the module described in this section exceed requirements at each stage of the demonstration flight. Although implementation, manufacturing and operation of designs based on this approach will continue to evolve, these results show this work has established a strong conceptual and practical foundation for future work.

Chapter 3

Thermal Management and Mechanical Integration

The previous chapter discussed how topologies like the flying-capacitor multilevel converter can simultaneously enable high-power-densities and high-efficiencies. Advances in wide-bandgap devices – gallium nitride in particular – mean that current densities previously only possible in large TO-247 packages (approximately 1680 mm³) can be achieved in new, much smaller, flip-chip footprints (less than 8 mm³). This exciting development also means that for comparable converter efficiencies, and corresponding per-device losses, the same amount of heat must be dissipated from a significantly smaller area. This chapter reviews several promising technologies using both air and liquid cooling media, then proposes a modular scheme that is well aligned with the theme of this thesis – and any converter implementation built from a multitude of independent, discrete devices.

3.1 Cooling Technologies and Media

Recent, comprehensive reviews in [124] and [125] have shown that cooling of power electronic device and modules is a prolific area of research. While the different solutions can be classified in increasingly granular sub-categories, a practical dichotomy exists between the choice of liquid or air to remove heat from the device. Given the high heat capacities of liquids, or refrigerants operating at the boundary of a phase change, it is understood that the highest power density designs will be facilitated through liquid cooling of some sort [126].

The most straightforward implementation of liquid cooling uses a cold-plate (heat sink) where a metal slab conducts heat into a cooling loop circulating through pipes within the plate [127]. A more advanced version of this concept replaces pipes with microchannels in the heat exchanger. This popular – but somewhat sophisticated – manufacturing process drastically increases the interface surface area in contact with the fluid, increasing heat transfer to 10-80 kW/m²-K for a significantly smaller volume of heat sink and fluid mass [128, 129]. Similarly, engineered surfaces have been used to enhance heat transfer through

both wetting and wicking action [130, 131], as well as mass transport through hydrophobic surfaces that facilitate “jumping droplets” [132]. Both solutions target heat transfers on the order of 20 kW/m²-K. Alternatively, jet impingement techniques, where a cooling fluid is directed at the surfaces of a thermal load, focuses less on surface structure and more on nozzle configuration and impact area. This approach allows for enhanced mixing of the fluid at the interface, and subsequent ejection of the heated media from the surface being cooled. Heat transfers of up to 60 kW/m²-K are possible, and designs applied to automotive power modules have received significant attention [133, 134]. Finally, direct immersion cooling, in which electronics are directly immersed in the coolant, presents an attractive method of leveraging high cooling of the various media capacities without the need for complex plumbing interfaces to the thermal loads. While this concept was originally proposed for aircraft modules using a dielectric refrigerants in direct contact with the circuitry [135, 136], recent advances in surface coatings allow the use of more stable and benign media like deionized water [137]. Although this technology is still somewhat nascent, thermal performances up to 50 kW/m²-K are on the horizon.

Indeed, each of the above configurations provide high heat flux solutions that will likely enable power densities beyond what is demonstrated in this work. However, the current lack of design experience, potential of leaks for coolant leaks, potential for fouling of the closed-loop fluids and interfaces, and cost can present a significant barrier to integrating liquid cooling into power electronics systems. In this way, air-cooled thermal management still provides an appealing solution. Advanced manufacturing techniques and materials have allowed for new means for optimal heat sink design, fabrication and augmentation [138–140]. A paradigm shift in design has also spurred engineers towards three-dimensional cooling architectures, such as dual-sided [141] or three-sided [142], forced-air cooling. This work recognizes the value of an air-cooled solution, and again applies a modular approach to heat sinking the multitude of devices on an FCML based converter.

3.2 Modular Heat Sink Design

Other designs using single, large area transistors or power modules often need to contend with local hot spots and rely on extrinsic heat spreaders to help mitigate this issues. However, FCML implementations leveraging smaller devices – though having a higher device count – benefit from an intrinsic heat spreading effect[80]. However, removing the heat from each transistor is still required for operation at high power. Previously, FCML inverters employing chip-scale GaN transistors used a compressible thermal interface material (TIM) to mate the devices to a solid heat sink or cold plate [66–68, 143], or otherwise required the heat sink be machined to precise tolerance [118]. Thus, a modular design approach, where heat sinks are attached via spring compression to each chip-scale GaN device, presented an appealing opportunity.

A modular heat sink design for FCML or any other converters requiring the individual heat sinking of multiple, discrete power transistors was originally proposed in [80, 144]. As




Figure 3.1: Comparison of previous and proposed modular heat sink mounting strategies. Note, the threaded insert in the right figure is only conceptual; the interference fit described later in this section made this specific consideration obsolete.

such, the approach will only be summarized here with the remainder of the section reserved for advances in the design, manufacturing and performance. Essentially, this approach focuses on cooling each power device individually rather than attempting to provide cooling to the entire converter with a single heat sink. This can provide an opportunity to minimize heat sink mass, reduce manufacturing tolerances (especially the need to exactly match variations in component height), and allow for the use of a higher performance TIM that might otherwise be impossible with a single heat sink. Additionally, just as a modular approach is applied at the system level to mitigate design complexity and support scalability, modularity in heat sink design can also be leveraged for similar benefits.

3.2.1 Single Switching Cell

The original conception of the modular heat sink intended the heat sinks to mount to the converter PCB through screws threaded into surface-mount nuts [80]. As this scheme continued to evolve, incremental advances in pad design and registration features were included to properly align the surface mount nuts during reflow in [1]. However, thermal performance continued to fall short of the expected improvements, and heat sink installation proved to be tedious and yielded inconsistent quality. Major issues arose from nut misalignment (in spite of improvements to the pad design), the use of plastic (electrically-insulating) screws that were prone to strip and cross-thread, and the narrow spring selection available for the form factor of the design. This last issue was driven by three competing constraints: a short enough free length to allow the screw to seat in the corresponding hole in the heat sink; or few enough turns so as not to go “solid” (completely compressed, i.e., rigid and no longer applying compressive force) when the screw was fully threaded; and a suitable spring constant to apply the correct contact pressure. Assembly was further complicated by the

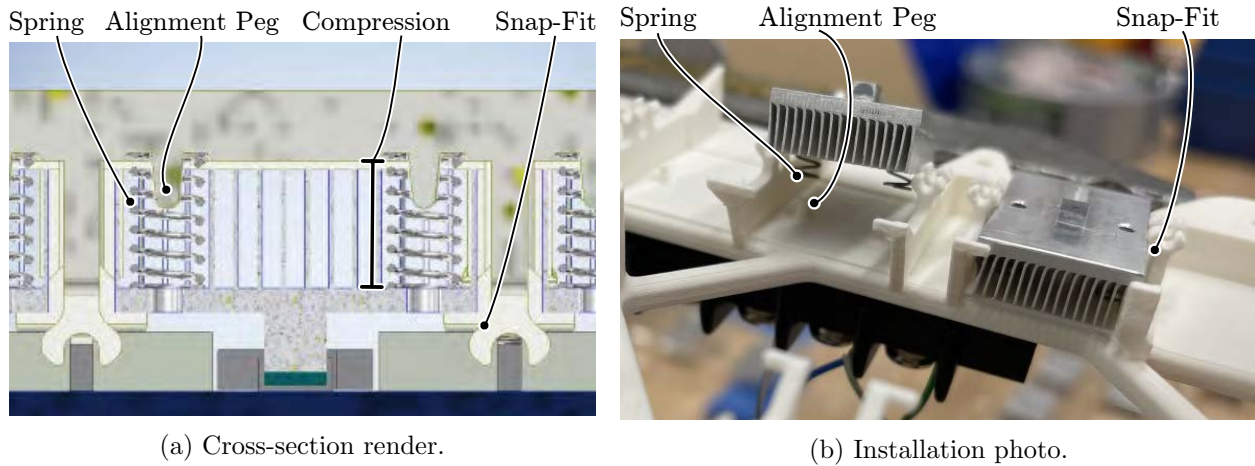


Figure 3.2: Examples of the snap-fit concept with spring alignment and interference features.

ongoing need to apply careful alignment and threading to two screws per heat sink for each device on the PCB – up to 72 manual insertions per converter for a dual-interleaved, 10-level module.

As such, this chapter addresses both the manufacturability and overall performance issues of the previous work. The first major update changed the plane by which the spring was compressed from the PCB surface (by way of threading the screw into the surface-mount nut) to the surface of heat sink ducting that was already present in the design. This update is illustrated in Fig. 3.1, where Fig. 3.1a illustrates the previous approach and Fig. 3.1b illustrates the approach developed in this work. By doing so, the range of acceptable spring free lengths was expanded, increasing the quantity of applicable, commercially available springs, listed in Table 3.1. Here, beyond mechanical fit, the primary characteristic was whether the spring, when compressed, would exhibit a contact pressure corresponding to the test conditions of the TC-5022 TIM that would be used [145]. TC-5022 was identified to provide superior thermal performance to other, compressible TIMs used in prior work. Additional considerations included spring cost, the extent to which the difference between free length and compressed length could absorb manufacturing tolerances, and whether the ends of the spring were milled flat to ensure a flush mate with the compressing surfaces. The highlighted row in Table 3.1 is the part number that best fit this criteria.

The second major change was to retain heat sinks (via captive attachment) in the converter enclosure housing by means of an interference, or snap-fit. This design is illustrated in the mechanical model of Fig. 3.2a, where registration pegs align each spring to the compression plane of the enclosure and interference features retain the heat sink when springs are uncompressed. Fig. 3.2b shows one heat sink in mid-assembly, where springs are being aligned with registration pegs on the enclosure, with the other heat sink retained with the deformable interference features.


Table 3.1: Selection of appropriately sized miniature springs from Access Spring.

Part Number	Length [mm]			Rate [N/mm]	Force [N]
	Free	Compr.	Delta		
PC356-3175- 6000-MW-6756-CG-N-MM	6.76	5.70	1.06	1.77	3.73
PC356-3048- 6130-MW-6350-CG-N-MM	6.35	5.43	0.92	1.97	3.61
PC305-2388- 9000-MW-6350-C -N-MM	6.35	5.43	0.92	1.35	2.48
PC356-2388- 7000-MW-5994-CG-N-MM	5.99	5.43	0.56	3.78	4.24
PC406-2896-10300-MW-5944-CG-N-MM	5.99	5.43	0.56	2.11	2.37

3.2.2 Complete Converter

The design modifications described above drastically improved converter manufacturability of the modular heat sink design. In this way, when the converter is leveraging dual-sided cooling, an entire half of the converter could be assembled at once by first installing heat sinks on a single enclosure side (shown in Fig. 3.3a), then mounting the enclosure to the PCB. Fig. 3.3b shows how this process further streamlined TIM application where, when using a stencil, application could proceed quickly from heat sink to heat sink. Note, the polyimide stickers visible in Fig. 3.3b are part of a pre-assembly process to electrically insulate heat sinks from high-voltage nodes present at the flying capacitor terminals.

Finally, an illustration of the proposed air-cooling configuration for the full converter module is shown in Fig. 3.4. Here forced-air is provided at inlets on either side of the converter. This cooling air runs across heat sinks on the outside of both the front and back of the PCB, then exhausts from the middle of the converter. While a single flow path from left to right or vice versa (i.e., across both sets of heat sinks on a single converter face and in



(a) Fully installed heat sinks.



(b) Stenciling of TIM.

Figure 3.3: Installation of heat sinks on an entire face of a module with dual-sided cooling.




Figure 3.4: Illustration of dual-inlet cooling of an FCML-based power module. One inlet is provided for each interleaved phase, with exhaust air designed to exit out the middle of each module.

one direction) would potentially provide a more ideal flow path, such a configuration would cause one phase to receive warmer cooling air and thus run hotter than the other. The implications of such an imbalance in cooling are discussed further in Section 4.3, while the consequences of the proposed cooling configuration are described in the following section.

3.3 Analysis using Computational Fluid Dynamics

Various simulations were conducted to provide expectations on pressure drops, fluid flow paths and overall thermal performance for the strategy proposed. Any computational fluid dynamics (CFD) study requires, at minimum, the following model definitions:

- *mesh geometry*, typically derived from a solid model that dimensions the system assembly. Fine meshes must be used for fine features.
- *material properties*, including specific heat, thermal conductivity, and density, as well as whether the material is a fluid or a solid.
- *boundary conditions*, for example fluid inlets and outlets or heat sources (i.e., heat fluxes), as well as wall definitions (interfaces between liquid and solid zones).

In the application considered in this work, a forced-air supply will provide the necessary fluid flow through the system. An ideal air supply would either be able to provide a constant pressure or a constant mass flow at the inlet to the system under study. In a circuit analogy, the two types of sources might be considered an ideal voltage or current source, respectively.

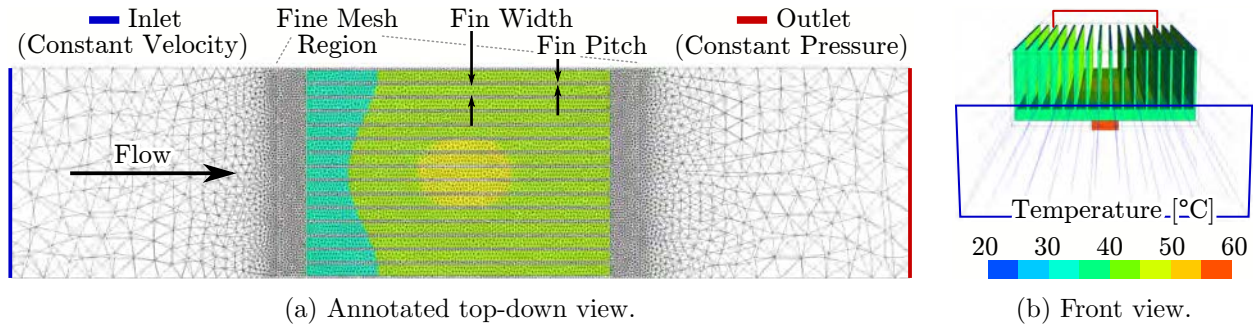


Figure 3.5: Meshing and results of CFD for a single heat sink. A finer mesh is used in the region around the heat sink. Fin width, fin pitch and fin count were parametrically varied and the system was systematically re-meshed and solved to find the best configuration.

Similarly, restrictions in the flow path can be thought of as resistors, where increased mass flow will induce a pressure drop across them. However, in a practical implementation the air source will have a finite capacity to provide flow or pressure; this could be considered the output impedance of the source, and the corresponding property is represented by a “fan curve” that links the dynamic flow the source can provide at a given pressure [146].

This property informs the boundary conditions used for the study in this work. Except at some combination of low fluid flow and high pressure-drop across the heat sinks, it is unlikely that a practical air source will be able to supply an arbitrary static pressure. On the contrary, the exhaust ports of the converter in the configuration under investigation would be open air; once air has passed through the outlet, the surrounding ambient can sink a practically infinite mass flow. Therefore, it makes sense that the boundary condition at the outlet be configured to constant pressure i.e., one atmosphere for tests at sea level – possibly lower if investigating systems at altitude. Then, considering that both the system curve (relating heat sink flow to pressure drop) is unknown, and that the individual heat sink performance (with respect to mass flow) is known from [80], it follows that the boundary condition at the inlet should be velocity or mass flow at a known (cooling) air temperature. In other words, total mass flow at the inlet can be swept to corroborate performance studied at the individual heat sink level – from which a new system pressure drop can be determined.

Heat flux is another important boundary condition, applied to the wall of the heat sink that interfaces with the transistor. For bottom cooled devices, the lowest thermal impedance to the junction is through the bottom of the device, so heat spreading and cooling through the PCB is typically prioritized [60, 72, 147]. For top-cooled devices, heat from the junction must primarily be removed through a heat sink. In this work, the devices under consideration also have a low thermal impedance path into the PCB, however, the dual-sided design and close component placement limits the efficacy of heat spreading and prohibits the use of such bottom-side cooling. Therefore, the focus here is on top-sided cooling; most of the heat flux through the junction will be directed through the top of the device. As such, the heat flux boundary condition can be safely approximated as the expected power dissipation in the device (up to 12 W) divided by the heat sink contact area (2.6 mm by 4.6 mm). This

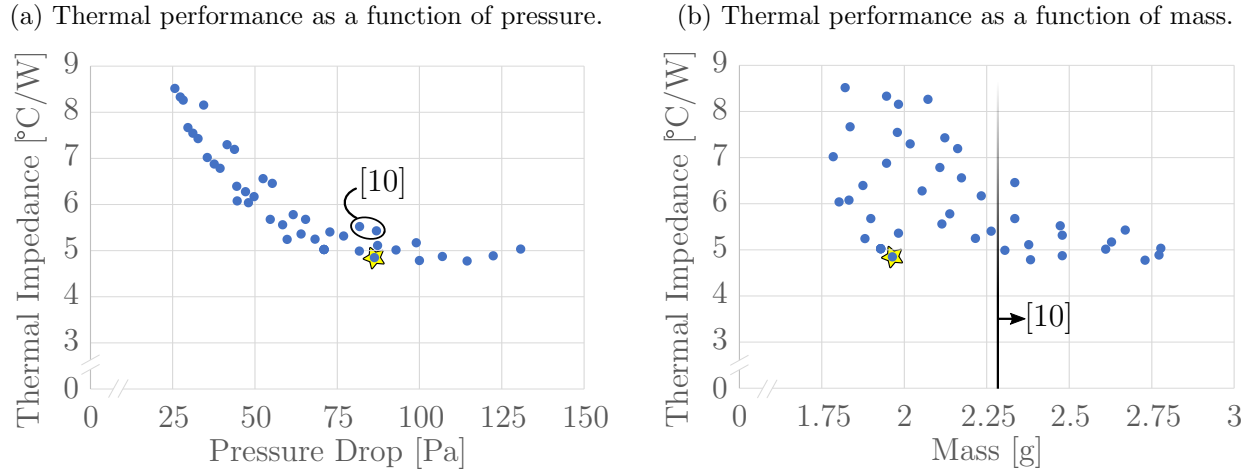


Figure 3.6: Results of the parametric sweep of heat sink fin count, pitch and width.

condition is then applied to the bottom of the heat sink post.

3.3.1 Single Switching Cell

Prior work in [80] established that an airflow velocity of 4 m/s should be sufficient to cool a heat sink that is sized to fit within the volume available in a dual-sided, high-density layout. However, that work only focused on implementations that could be manufactured on an in-house gang saw at the respective institutions' machine shops. Transition to a contract manufacturer enabled the use of skived heat sinks [148], a less constrained design space. While a methodology for constrained, weight-based optimization for air-cooling can be found in [149], this chapter used a parametric sweep across to study the improved but still reasonably tractable design space. The main independent parameters were the fin width, fin pitch and fin count.

First, a parametrically-defined solid model of the heat sink from Autodesk Inventor was imported into Ansys Workbench. These parameters were linked to the model definition and could drive new geometries as the simulation progressed through each iteration. Constant pressure (zero gauge) was specified at the outlet, and constant velocity (corresponding to the criteria from [80]) was defined at the input. Meshing was tetrahedral and largely automatic, with the exception of the refinement required over the region in and around the heat sink. An annotated rendering of this simulation setup is shown in Fig. 3.5 for the proposed configuration: Fig. 3.5a shows the mesh and flow direction while the contours indicate temperature; Fig. 3.5b illustrates the full three-dimensional geometry, with the heat sink post visible at the base. Note, a temperature gradient is visible from the top of the heat sink, particularly where the boss from the post on the bottom of the heat sink is situated [150].

The results of the parametric plot are shown in Fig. 3.6. With the system fan curve left as a free variable, the preferred design candidate was chosen to be near the pressure drop of the heat sink originally proposed in [80]. From Fig. 3.6a, up to a 13% improvement in

thermal impedance was identified; using the skived fin process, fin widths could be reduced from 0.5 mm to 0.2 mm – allowing for 5 additional fins (16 total) in the same volume with little impact on pressure drop. Correspondingly, Fig. 3.6b illustrates how this reduction in extra heat sink mass could produce an improvement in thermal impedance of up to 14%. This latter comparison indicated direct improvement to power density would result from this design iteration.

3.3.2 Complete Converter

A system-level analysis is essential to see what variation might exist across the inlets to each of the heat sinks on the converter. Unlike the study of the single heat sink, the full system design has cooling inlets on opposite sides of the module; exhaust air from both streams collides and must make a right angle turn to exit the converter. This was expected to impact cooling performance, and this section focuses on quantifying this impact. Modeling the full




Figure 3.7: Composite rendering of air particle tracks from the quarter-symmetric fluid simulation.




Figure 3.8: Composite image of the inlet at one side of the module, where the color map indicates static pressure at the inlet of each individual heat sink. Up to a 14% variation in pressures yields commensurate variation in mass flow, and therefore cooling capacity.

system posed significant challenges, both due to the large volume considered, as well as the dynamic range of features in the system. To simplify the computational domain, boundary conditions were defined to leverage the quarter-symmetric properties of the converter, as shown with the XY and YZ planes intersecting the module in Fig. 3.7. Furthermore, it was assumed that the inlet air stream would be well formed and developed from a straight section reduced to an individual converter inlet (i.e., not divided across multiple modules) – at least for initial tests.

Given the larger cross-sectional area of the inlet in this study compared to that of the single heat sink, a mass flow condition was instead applied at the input. Here, the mass flow was extracted from the heat sink CFD and multiplied by the number of heat sinks in the quarter symmetric configuration; 4.02×10^{-3} kg/s. The resulting pressure at the surface of the heat sinks was simulated and is visualized in Fig. 3.8. It is apparent that an appreciable variation in pressure of approximately 14% is present across each of the inlets: approximately 14%. Also, as this pressure is referenced to the ambient (i.e., zero gauge) pressure, it is clear that the full system induces a much greater pressure drop than originally expected: more than twice the increase in static pressure than that of the single switching cell simulations from [80] and Fig. 3.6 in Section 3.3.1.

3.4 Experimental Results

As the prior work in [80] has already provided ample experimental characterization of thermal performance of a single heat sink, this section only addresses system level concerns.

3.4.1 Temperature Measurements

Just as the general cooling challenges imposed by the small, power-dense devices featured in this work, measuring the thermal performance of the proposed technique were similarly compounded. While direct measurement of the junction temperature is preferred, as it directly correlates with device performance, this surface is on the bottom side of the device and is therefore inaccessible. However, [80] was able to demonstrate that measurements




Figure 3.9: Miniature thermocouple jig that enable consistent fixturing for chip-scale GaN devices.

of a thermocouple directly attached to the side of the chip-scale device accurately tracked junction temperature (as measured via infrared thermography through a window in the PCB) to within 1°C . Therefore, the preferred approach for this work was to affix a thermocouple to the side of the device.

This posed two further challenges. First, the thermocouple wire needed to be routed among densely populated components, and a suitable pressure applied to firmly hold the thermocouple junction against the side of the transistor. Second, some of the devices the thermocouple would pass near, as well as the transistor it would contact, connect to a high-frequency and high-voltage net during normal operation. To address the routing concern, an appropriate thermocouple was first chosen; a K-type thermocouple assembled from fine, 36 AWG wire were identified (Omega part number 5SRTC-TT-K-36-36). Then, custom fixtures were designed to safely and securely route the thin wire to the device. This fixture is annotated in Fig. 3.9, and was additively manufactured on a Form 3 stereolithography 3D printer. The interference-fit features indicated are a highlight of the design, as they allow for simplified installation in the space between two flying capacitors and above the gate driver; instrumenting an entire interleaved phase was able to proceed much faster than ad hoc use of epoxy and tape. On each jig, the thermocouple junction is routed across a cantilevered arm that provides a second interference fit against the side of the GaN device, applying the necessary contact pressure for the measurement.

The second challenge, related to the presence of high-voltages near the thermocouples, was addressed in two separate experiments. In the first experiment, all devices were operated in reverse conduction at dc to serve as thermal loads. This “thermal surrogate” approach was introduced in [80], and entails shorting each device gate terminal to the source potential. In previous works, this was a physically soldered short, whereas in this work, the logic power was simply enabled while gate drivers were commanded to the off-state. Then, by applying a reverse voltage across the module dc terminals, the reverse conduction losses described in Section 2.2.1 served as a realistic thermal load. Temperature and power measurements can then be performed at a much lower-voltage, dc operating point than normal operation. Thus, thermocouple junctions were allowed to make direct contact (with the sparing addition

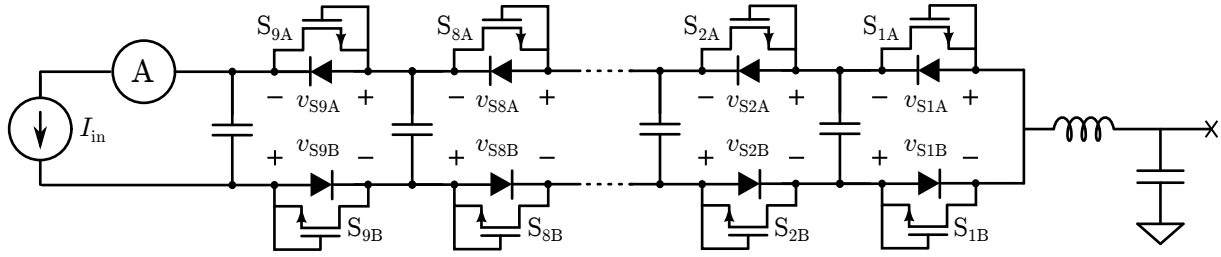


Figure 3.10: Schematic for connection of thermal surrogate. The gate of each device is shorted to the respective source, and a current is driven through the string of devices to provide a thermal load through the reverse conduction losses. The string current is measured with ammeter A at the input terminal, while individual source-drain voltages are measured across each device. Note, the “diodes” are only symbolic of the reverse conduction and do not exist in the physical system.

of Dow 340 thermal interface material) to the GaN devices as long as the common-mode voltage limits of the temperature logging hardware were observed. As such, the temperature measurements for devices across a single interleaved phase needed to be split across three independent logging units: one National Instruments PXIe-4353 module, one Omega HH1384 hand-held unit (with USB connection) and a Pico Technology TC-08.

The thermal surrogate configuration for a single FCML phase is illustrated in Fig. 3.10. Here, a dc current, I_{in} , is driven through each of the series connected transistors and measured through the ammeter (an element of a Yokogawa WT5000 in this work). Then, each source-drain voltage of interest (i.e., where the devices is instrumented with a thermocouple) is measured and multiplied by the series current to determine the power dissipated in each individual device. A Fluke 2635A Hydra Series II Data Bucket was used to make the voltage measurements; this data acquisition unit multiplexes up to 20 isolated channels, so it was necessary to allow for thermal steady-state to ensure accurate measurements before proceeding to the next operating point. Note, for these thermal surrogate experiments, the two phases on the dual-interleaved module were connected in series to ensure the same current was driven through each transistor.

To record temperatures during normal operation, where much higher voltages and frequencies are otherwise in close proximity to the thermocouple junction, additional measures needed to be taken. First, to address the high common-mode voltage (exceeding the allowable range by any of the three loggers), a polyimide tape barrier was laminated to the side of the device. The adhesive side was gently pressed against the PCB and GaN transistor to assure the surface was adequately wetted for the lowest thermal resistance. Additionally, Dow 340 thermal interface material was applied to the non-adhesive side for interfacing with the thermocouple junction. To mitigate high-frequency interference, which had shown to influence temperature measurements in previous experiments, additional electrical filtering was added to each thermocouple channel. This strategy is illustrated in Fig. 3.11. Here, two turns of both thermocouple wires are wound through a Laird LFB095051-100 cylindrical ferrite, while the individual strands are each wound twice around Laird 28B0250-100 cores. The

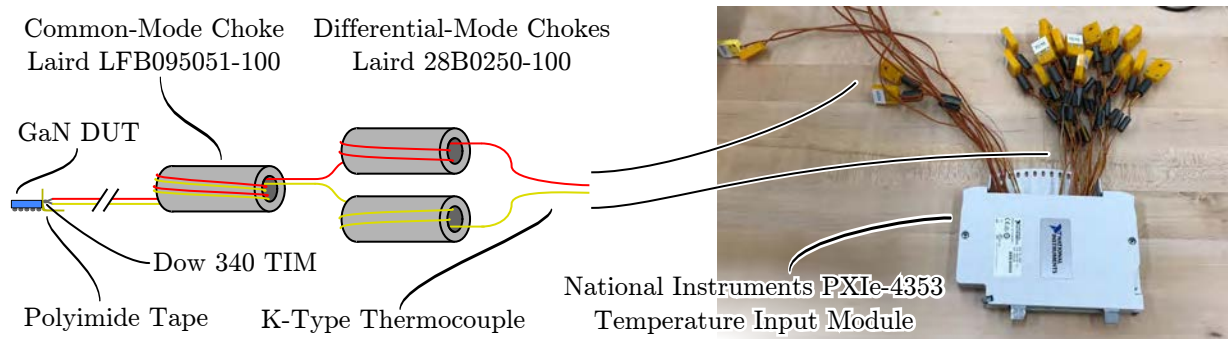


Figure 3.11: Common- and differential-mode ferrites installed on the temperature sensors.

broadband 28 material provides impedances of $> 90 \Omega$ for the frequency range 25-300 MHz while the LF material provides similar attenuation for the 500 kHz-30 MHz range. For reference, switching content exists primarily at 115 kHz and 1 MHz, while higher-order harmonics are also of consequence.

3.4.2 Mass Flow Measurements

To ensure the thermal performance was characterized at the appropriate operating point, it was necessary to determine the mass flow of cooling air from inlet to outlet. Several instrumentation options are available for airflow measurement, and vary in complexity, expense and ease of implementation. Ideally, the measurement technique should not affect the pressure or flow profile at the inlet to the module and should have a well understood means to convert to mass flow. While a turbine anemometer represents a cheap and low complexity solution, the meter will produce a vortex at the output and exact conversions to mass flow generally require generous assumptions on operating conditions at the inlet or outlet, or are otherwise unavailable. Both a pitot or hot wire represent appealing alternatives, but require a uniform and well developed flow profile. Additionally, the pitot tubes for the inlet size typically require high flows for appreciable measurement accuracy, while hot wire equipment is generally a more expensive option.

Therefore, a differential pressure measurement was used in this work – specifically an orifice flow meter. The general principle is based on the correlation between the pressure drop in a constriction in a pipe to the flow through the constriction. The principles of the measurement are well described in [151] and [152], and if the measurement apparatus is constructed to standards set forth in ISO 5167-2 [153], reasonable performance and accuracy can be expected without specialized characterization. To facilitate the straight pipe lengths necessary for well-formed flow, the inlet to the module was initially ducted via 3 inch schedule 40 pipe. The orifice plates for the flow rates in this work were sized using HydrauCalc software [154]. HydrauCalc implements the ISO 5167-2 [153] in an easy to use GUI for quickly checking operating points. This analysis showed that a 1.25 inch orifice would suffice for the flow rates corresponding to the initial CFD study, while a 2.25 inch orifice place

Table 3.2: Components used in the orifice meter for measuring mass flow.

Component	Manufacturer and Part Number	Range/Units
Orifice plate (low flow)	Custom laser-cut 0.03" steel plate	1.25" diameter
Orifice plate (high flow)	Custom laser-cut 0.03" steel plate	2.25" diameter
Static pressure	TE Connectivity 4525DO-SS3AS015AP	15 PSI absolute
Diff. pressure (low flow)	Amphenol DLHR-L01G	1" H2O
Diff. pressure (high flow)	Amphenol DLHR-L02G	2" H2O
Thermocouple	Adafruit MAX31856 Breakout	K-Type

would expand the dynamic range by an order of magnitude.

The orifice plates were laser cut from 0.035 inch mild steel, while custom flange fittings were fabricated from PLA on standard FDM 3D printers. The flanges were subsequently epoxy-sealed to the ends of the schedule 40 pipe to implement the orifice fixture. An orifice plate and custom flange is visible in the left photo of Fig. 3.12a, while the right shows the full measurement fixture. This fixture included several sensors, listed in Table 3.2, for measuring static pressure, differential pressure and inlet temperature. These sensors provided measurement data through a serial SPI link. The two different ranges of differential pressure sensor were built into the meter to further expand the dynamic range without requiring a physical change in orifice plate, though the lower range sensor was typically preferred due to higher span accuracy. All sensors were soldered onto a breadboard and connected to a measurement computer using a USB to SPI adapter.

3.4.3 Full Test Stand and System Measurements

To ensure minimal swirl in the flow measurement, the meter was placed upstream of two Fantech prioAir 8 EC blowers supplying cooling air to each converter inlet. This additionally allowed the use of a single transition from the 3 inch orifice measurement pipe to the 8 inch diameter of the blower and converter inlet duct. Flexible ducting was used for the segment between the orifice meter transition and the blower inlet to accommodate the overall size of the rig in the laboratory, where turbulence induced by the corrugation of the ducting would not impact the test setup. Conversely, smooth, rigid ducting was used in the segment connecting the blower to the converter inlet to ensure stable flow. Unlike a system where each heat sink has a dedicated cooling source, abrupt changes in direction – or the outlet of a fan, like the blowers described above – can add swirl or other non-uniform characteristics to the airflow profile. This, in turn, can create greater than expected discrepancies in the thermal performance across individual heat sink inlets – as the inlet pressure front becomes significantly distorted. Therefore, just like the straight lengths used to ensure fully established flow at the inlet of the orifice meter, additional flow conditioning was installed at each module inlet: a bundle of quarter-inch diameter tubes were packed into the rigid, eight inch duct supplying air from the blower outlet to the converter inlet, as shown in Fig. 3.12b. This





Figure 3.12: Custom test fixtures, instrumentation and scripted acquisition equipment for measuring thermal performance of a complete inverter module. Variable speed blowers provide cooling air, drawing air through a custom mass flow meter. The module (inset) is wired with thermocouples and drain-source voltage monitors per the test procedure described in this section.

measure effectively mitigates swirl and helps normalize the flow profile over a short length at the cost of a minor pressure drop.

The complete system is shown in the photograph in Fig. 3.12. Here, the orifice meter and associated 3 inch Schedule 40 piping is shown mounted to the door frame. The extent of the straight lengths of 3 inch pipe were necessary to ensure stable flow at the orifice. The flow from the two meters feeds two custom printed 3 inch to 8 inch fittings that connect the piping to flexible ducts attached to the blower inlets. The blower drives cooling air through a rigid duct with previously discussed flow conditioner into a final reducer that adapts the blower outlet diameter to the cross-section of the module inlet. The shape of this reducer is included in the full system CFD. The converter module itself is highlighted in the inset. Prior to module packaging for this testing, jumper wires were soldered to the source and drain terminals of each device and thermocouples were fixtured in place. After module assembly, the jumper wires were connected to a terminal block that interfaced with the Fluke data logger for v_{DS} measurement, while the thermocouples were plugged into their respective meter ports.

3.4.4 Comparison of Simulated and Observed Results

Using the experimental setup described above, several series of data were collected across various device dissipations and cooling air flow rates. Fig. 3.13 summarizes the data available at emulated device dissipations of 4 W, 6 W and 8 W. While the simulations indicated that a systemic variation in thermal performance across devices due to flow paths was likely, this was less apparent in the experiment. Indeed, differences in measurement arose mostly from variation in thermocouple contact pressure, heat sink shape or the degrees of leakage in each heat sink channel. As such, only the mean and standard deviation of the measured results for a given flow rate are plotted in the figures, and only the mean for the ensemble of devices is plotted for the simulation series.

As expected, the simulated values demonstrate an optimistic portrayal of the proposed cooling system. In the model, all fluid boundaries are perfectly sealed while thermally-conducting interfaces are perfectly mated. However, the physical realization had imperfections contributing to deviations from the simulated value. For instance, leaks will mean some of the cooling air bypasses the heat sink and reduces the cooling efficacy at a given operating point. Likewise, if the heat sink is not perfectly mated to the device, there may be additional thermal impedance at the interface. However, it is also worth noting that the thermal simulation did not include the impedance of the thermal interface material – simply considering the interface to be an ideal thermal flux path from each device. Once this additional impedance was included, using experimentally determined values from [155], the experimental results track the theoretical performance quite well.

Finally, to determine the peak cooling capability, the mass flow measurement apparatus was removed from the blower inlets. This allowed for increased flow due to the reduction in static pressure each blower needed to overcome at peak flow. To estimate mass flow without the meter, the blowers were set to maximum and the resulting static pressure at the outlet

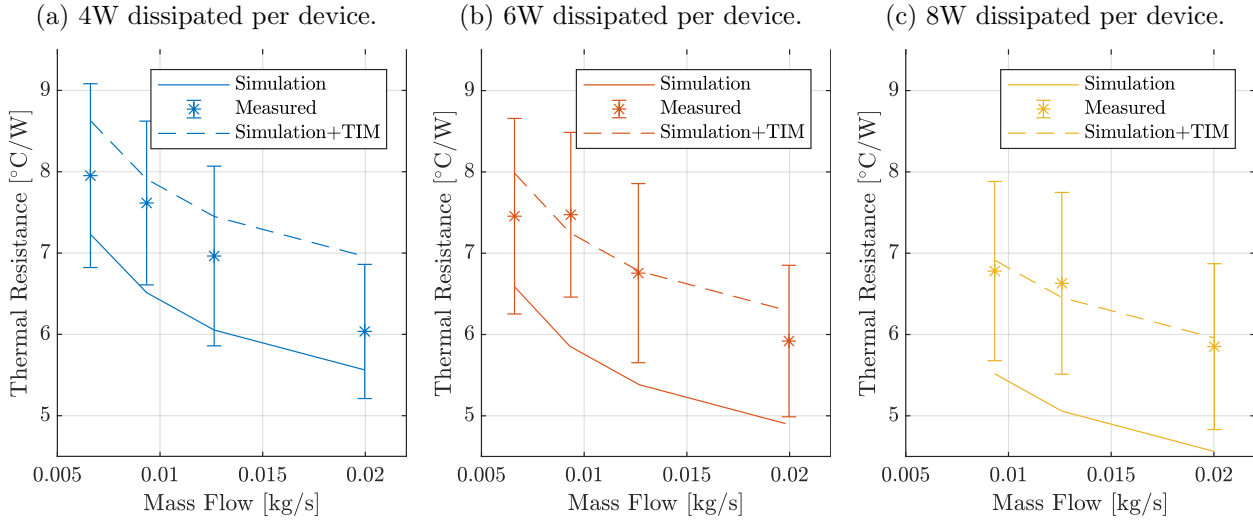


Figure 3.13: Simulated and experimentally measured thermal performance across varied cooling air inlet mass flow. The simulation does not model TIM resistance; a constant impedance value for the TC5022 TIM from Table 2 of [155] is added to the simulated values to yield the dashed lines.

was measured using the same sensors that were part of the flow meter. Then, this pressure was referenced to the blower fan curve [156] to determine the rated flow at that operating point. This ultimately resulted in an experimental value of average total thermal impedance ranging from 4.75 to 5 $^{\circ}\text{C}/\text{W}$ for a mass flow of 0.076 kg/s .

3.5 Chapter Remarks

The results presented above represent a significant advance in the modular heat sinking concept. Compared to prior work, both manufacturability and overall system thermal performance have been drastically improved. Nonetheless, several sources of error between predicted and realized performance exist. First and foremost, given the tolerances on the printed baffles and the heat sinks (including tolerance on the total number of pins), gaps between nominally sealing interfaces (e.g., the duct area around each inlet) can contribute to lost cooling air and decreased cooling efficacy for a prescribed mass flow. In future designs, this mismatch in device tolerance should be accounted for with additional features or sealing mechanisms that can absorb this mechanical variation. Introducing dynamic seals (e.g., which are made as the heat sinks and complete enclosure assembly are installed) will pose further challenges, likely including additional friction forces that may oppose the intended heat sink mating force and complicating device alignment. Nonetheless, avoiding heat sink bypass is essential for the efficacy of this type of air-cooled configuration.

Further considerations will also be needed when scaling this concept to multiple modules. The work presented above highlights the challenges of ensuring that a single module is receiving a nearly uniform pressure front across all of the individual heat sinks. If this is

to be guaranteed across several modules, a decision will need to be made on whether to supply cooling air to the entire array with a single duct, or to route individual cooling air connections to each module in the array. In either approach, the inlet flow path is sure to be much more meandering than what was illustrated in the experiments above, as integration into a more compact system will necessitate elimination of the straight lengths shown in the test setup of this work. As such, guaranteeing a uniform pressure front will likely require the addition of well-designed baffles that can compensate for any swirl, turbulence or non-uniform pressure cross-sections. It is likely that this will be a complex analysis owing to the large dynamic range of fluid channel sizes, however some compromise between design and analysis complexity and uniformity in thermal performance is likely to be acceptable (as was the case in this work).

Additionally, this work represents improvements through revision of an existing mechanical design. This means that, although much of an existing design and material stock was able to be repurposed for this work, further improvements might be realized through a completely new design effort. For instance, the post extending from the heat sink to the chip-scale device was sized to meet the area of the chip exactly; a future revision should add to the dimensions to reduce the alignment precision required during assembly. Or, if the large flying capacitors can be moved or reduced in physical size, the post may be reduced or removed altogether, as it represents an odd manufacturing, mating and thermal path geometry. If the post is retained, the heat sink fixturing scheme should also be revisited. The design presented above showed only two mount points on either side of the post. This both underconstrains alignment of the post, while also leaving open the possibility for the heat sink to tilt towards the front or the back. At minimum, three or four mount points should be considered to apply symmetric downward force.

Finally, any future designs considering air cooling should limit the design to a single inlet and outlet pair. While the original intent of this module design was to save space by grouping two interleaved phases on a single PCB, Section 3.2.2 illustrates that implementing dual-sided cooling with a vertical switching cell design leads to opposing outlet paths. Instead, the cooling air should flow from inlet straight through the heat sinks to the outlet on the other side to avoid excessive pressure drops. A vertical layout may be used with dual-sided cooling for a single phase, or a dual-interleaved approach could again be implemented if lateral switching cells are instead used. Nonetheless, this is likely the most important design insight to be realized in future revisions.

While liquid cooling is likely the only path forward for the highest power density designs yet to come, there are still applications where either the expense or complexity hinder the otherwise impressive performance advantages that come with the technologies discussed in Section 3.1. In these applications, the methodology set forth in this chapter will serve as a guidepost and reference for what can be attained with air-cooling alone. As interest for multi-transistor topologies such as the FCML continues to grow, this modular approach is likely to be of continued interest in future designs.

Chapter 4

Inverter Array Design and Operation

Portions of this chapter are adapted in part or in whole from [2].

The previous chapters have introduced the notion of a segmented or modular approach, and how it can produce a more robust and scalable three-phase drive through the interconnection of individual inverter modules [75, 157]. This chapter more fully explores the design of a this inverter array using modules based on the high-performance flying-capacitor multi-level architecture. While single-phase experimental demonstrations have shown exceptional performance in Chapter 2, multi-module, multi-phase operation represents a significant step forward for the technology. Further interleaving these modules also reduces the harmonic spectra at the converter terminals and filtering requirements necessary to meet vehicle electromagnetic compliance (EMC) standards.

This chapter proceeds as follows: First, Section 4.1 illustrates a hierarchical control interface to coordinate between inverter modules so that multi-phase operation can be demonstrated in the ensemble. Next, interleaving of modules paralleled within a given line phase is proposed – with implementation strategy discussed in Section 4.1.2, followed by analysis of corresponding input and output current spectra in Section 4.2. Then, Section 4.3 discusses some potential implications of unequal current sharing between interleaved phases. Finally, Section 4.4 presents experimental results validating the approach, measured on a CISPR-25 pre-compliance test fixture. Finally, Section 4.5 provides concluding remarks.




Figure 4.1: Schematic representation of how M interleaved 10-level FCML inverters can be interleaved within a phase-leg.

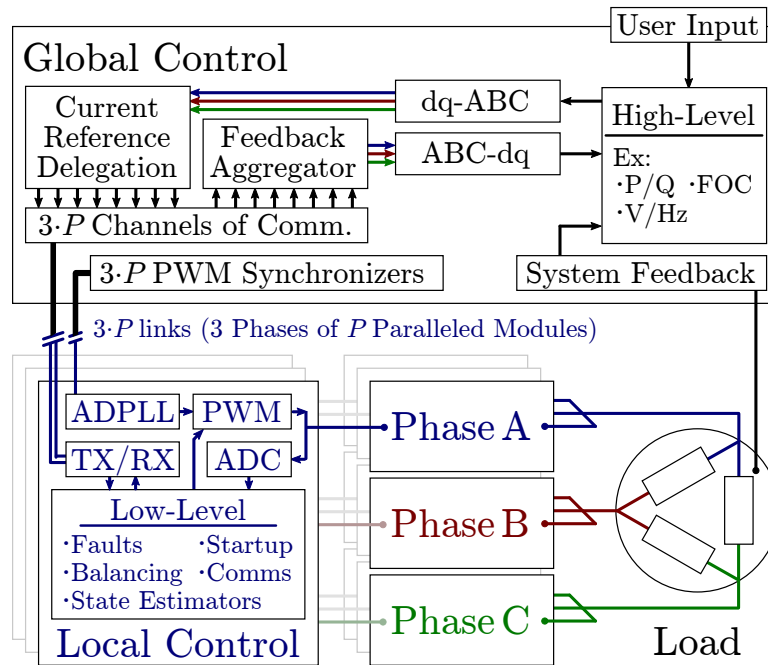


Figure 4.2: High-level illustration of the hierarchical control structure, configured here to drive a three-phase load.

4.1 Hierarchical Control

A drive based on a modular design scales output power through paralleling inverter modules. This increases the overall device count, and therefore gate signals, in the full converter – regardless of phase-leg topology chosen. For multilevel topologies, which feature numerous switches for even a single-phase (e.g., the 10-level FCML in Fig. 4.1 has 18 per phase), this gate signal count would be even higher; driving a 3-phase array of P -paralleled modules of N -level FCML inverters from a single microcontroller would require routing $3 \cdot P \cdot (N - 1) \cdot 2$ high-frequency signals with low-latency in what is typically a high-noise environment. To address this signaling challenge, control tasks can be divided between a global array controller and local controllers across the individual inverter modules.

4.1.1 Global & Local Control Domains

For pulse-width modulation (PWM), gate signal generation represents a relatively low-level task. This is especially the case with phase-shift PWM (PSPWM), where the carriers remain invariant under normal operation. As such, the carrier generation, as well as the comparison to an input reference, can be performed locally and independently on each inverter module. However, the input reference depends on states outside of the individual inverter domain (e.g., motor angle, speed reference, direct and quadrature currents, etc.), and this control must be computed at the system level. These two tasks serve to represent the local and

global domains of the proposed hierarchical control structure for the segmented drive of this work. This concept is fully illustrated in Fig. 4.2.

4.1.2 Control & Measurement Scheme

In the scheme of Fig. 4.2, the local controller is responsible for generating switching signals and collecting output voltage and current measurements – which are used locally for fast-acting fault mitigation and also sent upstream to the global controller via a serial communication link. In turn, the global controller can receive the output current and voltage measurements from each inverter module and, in aggregate with any exogenous inputs, can send a corresponding closed-loop reference back down the communication link to each local controller. While an industry standard protocol may be able to provide this link [158], to simplify development a proprietary protocol comprised of 64 bit Manchester-encoded frames was implemented over a point-to-point optical link in this work. This also serves to defer any signal integrity challenges that may arise from electromagnetic interference (EMI) coupling into an electrical communication link (e.g., RS-485), with the temporary trade-off of higher cost and assembly effort.

Note that while the AFBR-59F2Z transceivers used can support up to 250 Mbit/s, the data-rate was limited to 6.25 Mbit/s (to benchmark for future implementations what will use cheaper and slower electrical interfaces); the system update rate is just under 100 kHz. With this configuration, the multi-phase drive waveform at the fundamental can still be generated at the global level and distributed to local controllers digitally. However, any interleaving of inverters at the switching frequency requires a more nuanced strategy.

4.1.3 Carrier Synchronization

Interleaving of inverter modules provides an attractive method of reducing switching harmonics at the input (predominantly at the switching frequency, f_{sw}) or the converter output (predominantly at the effective output frequency, $f_{eff} = (N - 1) \cdot f_{sw}$). This control change lessens the impact of passive filters on power density. However, interleaving requires the modules to share a common frequency and phase reference around f_{sw} . This can be realized by synchronizing each local controller clock with a global reference.

One common synchronization method utilizes a digital pulse to reset local PWM counters to a nominal value (associated with the requisite phase-shift) each switching period. Such a scheme is often employed to synchronize counters across PWM peripherals within a single microcontroller [159], and is indeed implemented within the local controller of this work for PSPWM generation. It is also possible to choose a source external to the microcontroller (via a configurable input pin) to synchronize PWM counters with an external clock. Such a scenario is illustrated in Fig. 4.3, where the rollover of a global (signed) counter triggers a sync pulse that correspondingly resets the local (signed) counter used in PWM generation.

Here, when the two counters are initially unsynchronized at ①, the synchronization input at ② can lead to a gate signal of arbitrary pulse width on multiple devices instantaneously.




Figure 4.3: Illustration of a time-base synchronization scheme common in microcontroller PWM peripherals.

In a system with a common clock source driving PWM counters, or where the various global and local clock frequencies deviate only slightly, this may be the only abrupt correction that occurs. For example, if system clocks on both the local and global controllers are operating at 200 MHz and derived from 50 ppm oscillators, at most a single digital count of error may accumulate over counter periods used for kHz-range and above switching. In this case, the system would maintain glitch free synchronism shown at ④ and all future sync pulses. However, if there were significant frequency variation in system clocks, or if a (e.g., spurious EMI-induced) glitch occurred on the sync line, shown at ⑤, a similar distortion of the switching waveform would again occur, as seen at ⑥.

To address these concerns, additional filtering or windowed-detection schemes could be employed on the synchronization line. However, this work instead employs an all-digital phase-locked loop (ADPLL) to provide both inherent filtering (via the loop filter), as well as more gradual correction mechanism that varies the counter increment value rather than abruptly resetting the counter as described above. If the global reference is provided over a dedicated synchronization line, as shown in Fig. 4.2, implementation of the phase-frequency detector (PFD) is relatively straightforward [160] and design and analysis can be adapted from [161]. Note that if the two system clocks only differ slightly, as in the case in this work, adding or skipping increments of the local counter each switching period to speed up or slow the counter will overcorrect. As such, a fractional PLL [162] was implemented. This way, these discrete corrections can be spread over several switching periods to provide greater average frequency precision at the price of negligible instantaneous phase jitter.

While this configuration (as illustrated in Fig. 4.2) was verified in initial tests, the routing of the additional sync signals in tandem with the existing communication line—as well as ensuring their signal integrity at high-power operation—represented an additional barrier towards scalability. Instead, synchronization over the single communication link was much more preferable. Thus before sending each communication packet, the modified strategy shown in Fig. 4.4 was developed for further testing. Here, the global controller samples its counter value and appends this signed value to each transmission. Then, when the local controller receives each packet, the error between local counter value and that in the




Figure 4.4: Illustration of a synchronization scheme using a communication link and ADPLL.

packet is directly calculated. This proposed solution worked well for experiments in this work, with results detailed in Section 4.4. Note that in either implementation of ADPLL, a delay will be introduced when transmitting the counter reference from the global controller to each local controller. This delay will play a role in motor control loop development. However, considering that it is both stationary and common to each converter, it does not significantly impact module interleaving. Indeed, any differential delay between modules is limited to the nanoseconds of jitter or propagation delay across individual transceivers. As such, the differential phase between local controller PWM time bases is dominated by the programmed phase that interleaves the modules. As it were, interleaving is thus achieved by correctly phase-shifting PWM synchronization signals the global controller provides to individual modules. Note, loss or corruption of this signal only leads to aberrations in the EMI and does not significantly compromise PWM generation.

This second approach differs from the ADPLL described above in two significant ways. First, this method cannot resolve the full phase error range of -2π to 2π [160]. Unlike the PFD, which generates an error signal based on the timing and order of clock edges, the direct difference as implemented cannot differentiate between an error less than $-\pi$ and one greater than π . Thus, the “PFD” error here only approximately correct when the global and local controller frequencies are close. This also means the ADPLL may lock π out of phase – though this issue can be ignored as that the modules in this work are dual-interleaved locally.

Second, the conventional loop analysis used to characterize PLL frequency response and stability is not directly applicable. This is primarily because the two counters in this work are under-sampled; the update-rate of the communication link described in Section 4.1.2 is under 100 kHz, while the switching frequency (and respective counter frequencies) are closer to 115 kHz. Although coherent sampling, where the communication link would be synchronized with the switching frequency [163], may mitigate these issues, it would place additional constraints on the communication interface. Another potential solution would be to boost the communication frequency to above the Nyquist rate of the maximum permissible switching frequency.

4.2 Multi-Module Harmonic Content Analysis

The benefits of multi-phase power, including the reduction of input current ripple at the *fundamental* and cancellation of *line frequency* harmonics at the output, are generally well understood and available in the literature. Therefore, this section focuses on reducing harmonics of the *switching frequency* in the input and output currents through the synchronization and interleaving of modules as previously discussed.

Prior work has shown that a PWM process can be described by a double Fourier series [164]. Indeed, the Fourier coefficients and switching functions necessary to study the PSPWM operation of an FCML inverter can be readily found in [63, 116, 165]. The general idea behind the double Fourier series is that a summation of terms composed of harmonics of both the carrier and fundamental can describe the modulated waveforms in an inverter. The Fourier coefficients for such a study are doubly subscripted as A_{mn} and B_{mn} , where m specifies the carrier harmonic and n specifies the harmonic of the fundamental. As this work used trailing-edge carrier natural sampling (as opposed to double-edge, i.e. up-down, used in [116, 165]) a separate derivation is provided here.

First, A_{mn} and B_{mn} are recalled from (3.16-3.18, 3.21, 3.25) in Chapter 3 of [164]:

$$A_{mn} + jB_{mn} = \begin{cases} 2 & m = 0, n = 0 \\ M & m = 0, n = 1 \\ 0 & m = 0, n > 1 \\ j \frac{2}{m\pi} (\cos(m\pi) - J_0(m\pi M)) & m > 0, n = 0 \\ \frac{2}{m\pi} J_n(m\pi M) (\sin(n\frac{\pi}{2}) - j \cos(n\frac{\pi}{2})) & m > 0, n \neq 0 \end{cases} \quad (4.1)$$

for a modulation index, M , and where J_n is the Bessel function of the n^{th} kind. The next step is to substitute these values into the expression for the output current. This leads to a phasor representation of the output current as a function of m and n , adapted from (15) in [116]:

$$I_{o,mn} = \frac{V_{dc}}{Z_{mn}} \left(\frac{1}{2} Y_{mn} e^{j\gamma_m} - \sum_{k=1}^{N-2} \frac{k}{N-1} X_{mn} e^{j\theta_{k,m}} \right) \quad (4.2)$$

where k is the index of the k^{th} flying capacitor and, adapting (9) from [116],

$$X_{mn} = 2 \sin\left(\frac{m\pi}{N-1}\right) (A_{mn} + jB_{mn}) \quad \text{and} \quad Y_{mn} = 2(A_{mn} + jB_{mn}). \quad (4.3)$$

The angles, γ_m and $\theta_{k,m}$, are used to introduce phase shifts for harmonics of the switching frequency arising from PSPWM as a function of k , where k is the index of the capacitor such that $k \in [1, N-2]$. These terms are modified in this work to accommodate for interleaved operation. Essentially, to model a converter x where $x \in [0, P-1]$ in a group of P interleaved

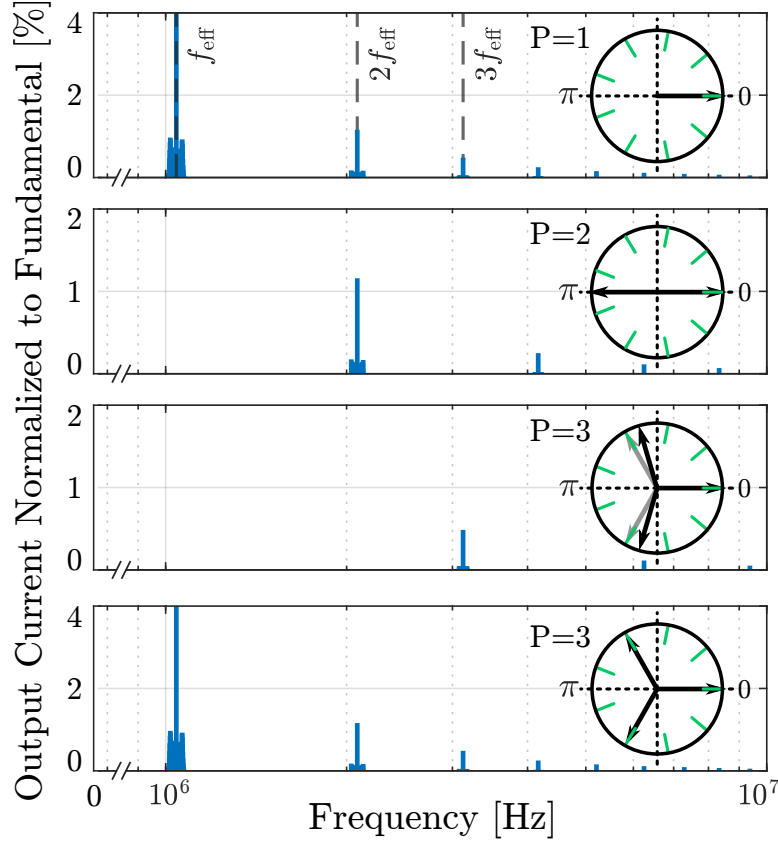


Figure 4.5: Calculated output current spectra in the region of interest (i.e., harmonics of f_{eff} for $N = 10$). The unit circles are normalized to the switching period, where arrows indicate the phase shift of interleaved modules. Green ticks mark each full cycle at the effective output frequency per switching period.

N -level FCML inverters, a phase shift term of $\frac{mx2\pi}{P}$, is added to the switching function angles defined in [116, 165]:

$$\gamma_m = \frac{2m(N-2)\pi}{N-1} + \frac{mx2\pi}{P} \quad (4.4)$$

$$\theta_{k,m} = \frac{m(2k-1)\pi}{N-1} + \frac{pi}{2} + \frac{mx2\pi}{P} \quad (4.5)$$

The harmonic magnitude for a single converter can already be directly calculated from (4.2) by taking the magnitude of the complex expression at each frequency. Note, the phase shifts in the complex sum of harmonic contributions from each level are what lead to the reinforcement of the “effective switching frequency” and elimination of the switching frequency at the output. A further summation of (4.2) for γ_m and $\theta_{k,m}$ from (4.4) and (4.5) with $x = 0, 1, \dots, P-1$, then taking the magnitude, produces the output current harmonics for an entire interleaved phase-leg. The resulting spectra for a single 10-level FCML inverter

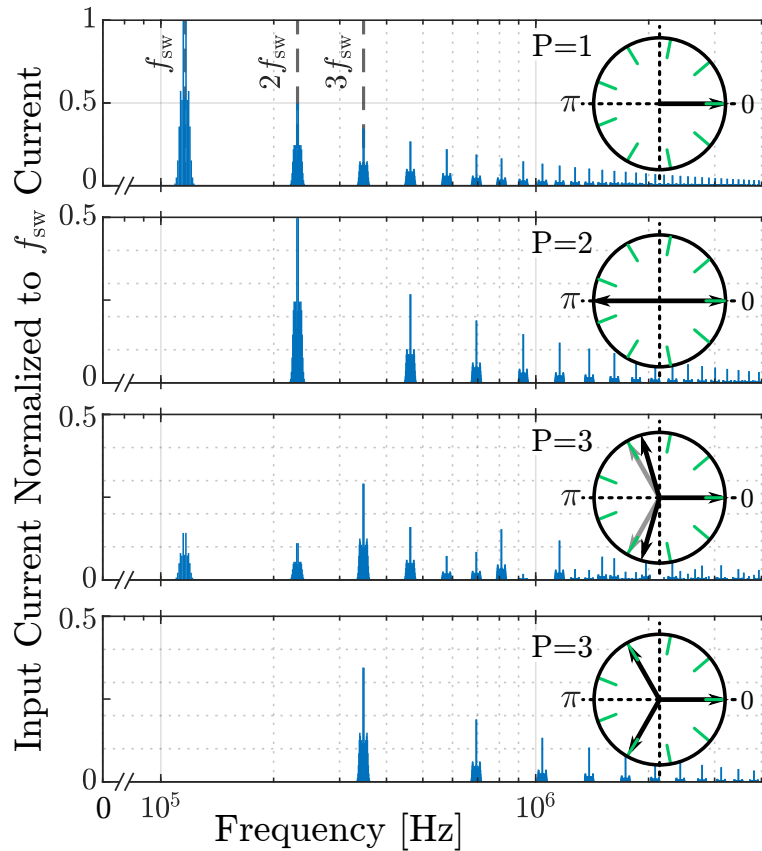


Figure 4.6: Calculated input current spectra in the region of interest (i.e., harmonics of f_{sw}) for $N = 10$.

($P = 1$), as well as the dual- ($P = 2$) and triple-interleaved ($P = 3$) cases, are shown in Fig. 4.5, with harmonic currents normalized to the total magnitude of the fundamental. Spectra shown are for the inductor current to emphasize harmonic cancellation; the load typically sees very little high-order switching harmonics since it is after this filter.

Note that for PSPWM, where the effective output frequency is an integer multiple of the switching frequency (which is seen at the input), P and $N - 1$ must be co-prime to simultaneously interleave the input and output. Otherwise, the phase shift that interleaves the input will wrap the phase at the output. This is illustrated with the unit circles representing the switching period in Fig. 4.5. The arrows show the phase of each interleaved converter, while the green ticks indicate the $N - 1$ full cycles of the output frequency each switching period. The middle two spectra show interleaving at the output, as reflected in each subsequent cancellation of effective switching frequency harmonic. However, with $P = 3$, this cannot occur at the phase shift required for interleaving the input (indicated with gray arrows). The last plot shows that the phase shift corresponding to interleaving the inputs for $P = 3$ leads to no reduction in load current harmonics. Section Section 4.4 will also show that a similar case arises when $P = 6$ and $N = 10$, though repeated dual-interleaving at the




Figure 4.7: Calculated input current harmonic content using double Fourier analysis and assuming sinusoidal output current.

output is possible.

To calculate the input current, the switching function for the DC side switch (i.e., S_{N-1}), containing the additional phase shift term, must be convolved with an expression for the output current. Given the sinusoidal output of the converters under study [1], this convolution can be simplified when the output current is assumed to be sinusoidal at the fundamental frequency [166]. Then, summing across the resulting spectra for interleaved converters $0, 1, \dots, P - 1$ again yields the input current harmonics of a given phase-leg, as shown in Fig. 4.6. Note that an appropriate choice of phase shift that interleaves the output can still significantly reduce input harmonics for $P = 3$.

The calculated input current spectra for a three-phase, segmented drive comprised of an array that is 6-fold interleaved in such a way is shown in Fig. 4.7. The simulated conducted EMI in $\text{dB}\mu\text{V}$ was derived by applying the transfer function for the line impedance stabilization network (LISN) used in Section 4.4 to the calculated input current harmonic amplitudes for a specified output current. Additionally, for comparison, the input current for a single 6-fold interleaved phase-leg, as well as that for a single, constituent FCML converter, are also plotted. Here, the single converter is shown to contain a broad spectrum of switching harmonics. However, as expected from the Fourier analysis, all but the $6m$ harmonics cancel in the interleaved phase-leg. Finally, although the full three-phase array spectra has greater peak magnitudes due to the higher total power output, it contains fewer harmonics in due to the cancellation of triplen sidebands—though this detail is not visible in the figure owing to the range of frequencies plotted.

4.3 Current Sharing & Parameter Sensitivity

While control loops on the local or global controller may be employed to enforce current sharing, the open-loop behavior provides a worst-case scenario to analyze. This work exam-

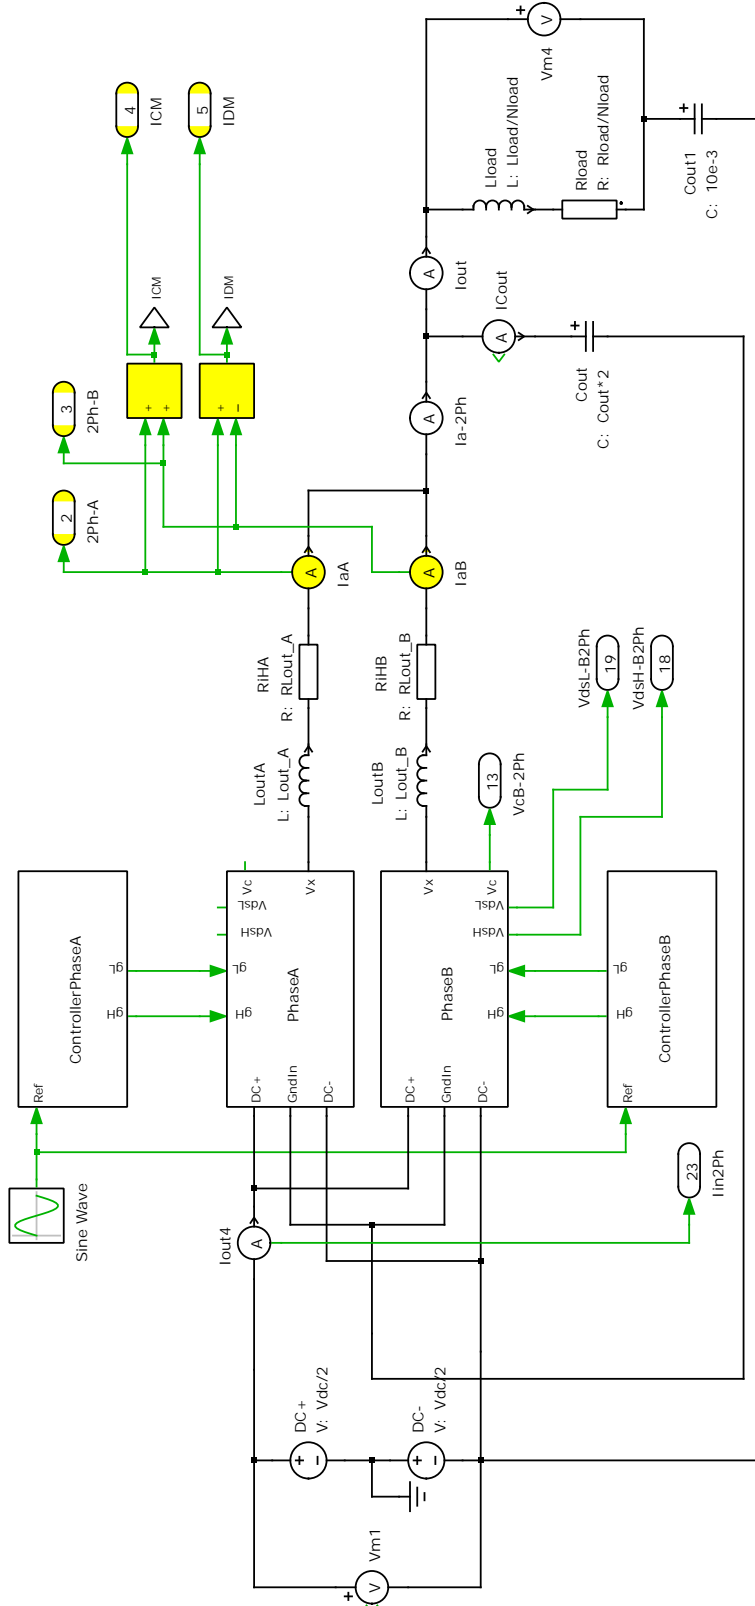


Figure 4.8: PLECS Blockset schematic of a dual-interleaved, 10-level FCML for Monte Carlo simulations. The two interleaved converter phases are the modules labeled **Phase A** and **Phase B** and contain the switches and flying capacitors whose parameters are varied throughout the simulation. Measurements of interest (pertaining to current sharing) are highlighted: the two respective phase currents, as well as their respective common- and differential-modes. The RMS values of these measurements, as well as the average of output power into R_{load} , are computed in a top-level Simulink simulation.

ines two sources of open-loop imbalance in a dual-interleaved, 10-level FCML module arising from parameter variation, either due to component tolerances or a bias consistent across an entire interleaved phase (e.g., from unequal cooling). The first case was examined through a Monte Carlo simulation to consider the variation in component values solely attributed to manufacturer stock. In this way, these parameters are treated as random variables. The key component parameters are the flying capacitances, device on-resistances, and the inductance and resistance of the output inductors in each phase.

For some components or product lines, manufacturers will release histograms detailing the parameter distributions. However, such data were unavailable for the components in the proposed converter of this thesis. Nonetheless, passive component part numbers often specify a tolerance, tol , typically used to describe the six-sigma spread in standard deviation, σ , around the nominal value, i.e. $\pm 3\sigma$; thus, $\sigma = tol/3$. It is also reasonable to assume that the parameters of many component types, especially passives, will have a Gaussian distribution around the rated value, μ – especially when considering the large volumes of each type of component in production [167]. Therefore, for each Monte Carlo iteration, the value of each random variable is selected from a distribution parameterized by μ and σ ; in MATLAB, this is the command `normrnd(μ, σ)`. Note, for the paralleled components comprising the flying capacitance, σ is reduced from the manufacturer value by the square root of the number of individual components, i.e., $\sigma = tol/(3\sqrt{N_p \cdot N_s})$, where N_p and N_s are the number of parallel and series connected components, respectively.

Transistor devices do not typically have a tolerance like passives. Instead, there is a reported *typical* value, as well as a *maximum* value. In this case, an exponential distribution with offset was assumed as an approximation to the inferred, single-sided distribution provided by such datasheet representations. In MATLAB, this was generated using a Weibull distribution, pd , with 1 for both parameters: `pd = makedist('Weibull', 1, 1)`. Then given a nominal on resistance, R_{on} , a sample $R_{on}(i)$ could be chosen by scaling this value with a selection from the distribution: $R_{on}(i) = R_{on} \cdot (\text{random}(pd) \cdot 10 + 100) / 100$. Note, this framework does assume that components within a specific build have distributions representative of the entire product line; actual component lots used for converter production, especially small runs, will likely see slightly different and possibly tighter distributions associated with lot to lot component variation.

The exact implementation scripted the nominal model parameter definition from Chapter 2 and the distributions described above in a Simulink model, with Monte Carlo iterations executed in parallel (via `parsim`). The Simulink model incorporated a PLECS Blockset realization of the dual-interleaved inverter, shown in Fig. 4.8, to generate exact solutions. Measurements of circuit electrical performance, such as output power and phase currents, are recorded for 30 fundamental cycles – enough time to allow any startup transients and balancing perturbations due to adjusted component values to fully ring out. The Simulink model then computes the RMS values of phase currents and output values for post-processing.

Results of the Monte Carlo simulation run over 100 iterations is shown in the correlation matrix in Fig. 4.9. The key measurements include:

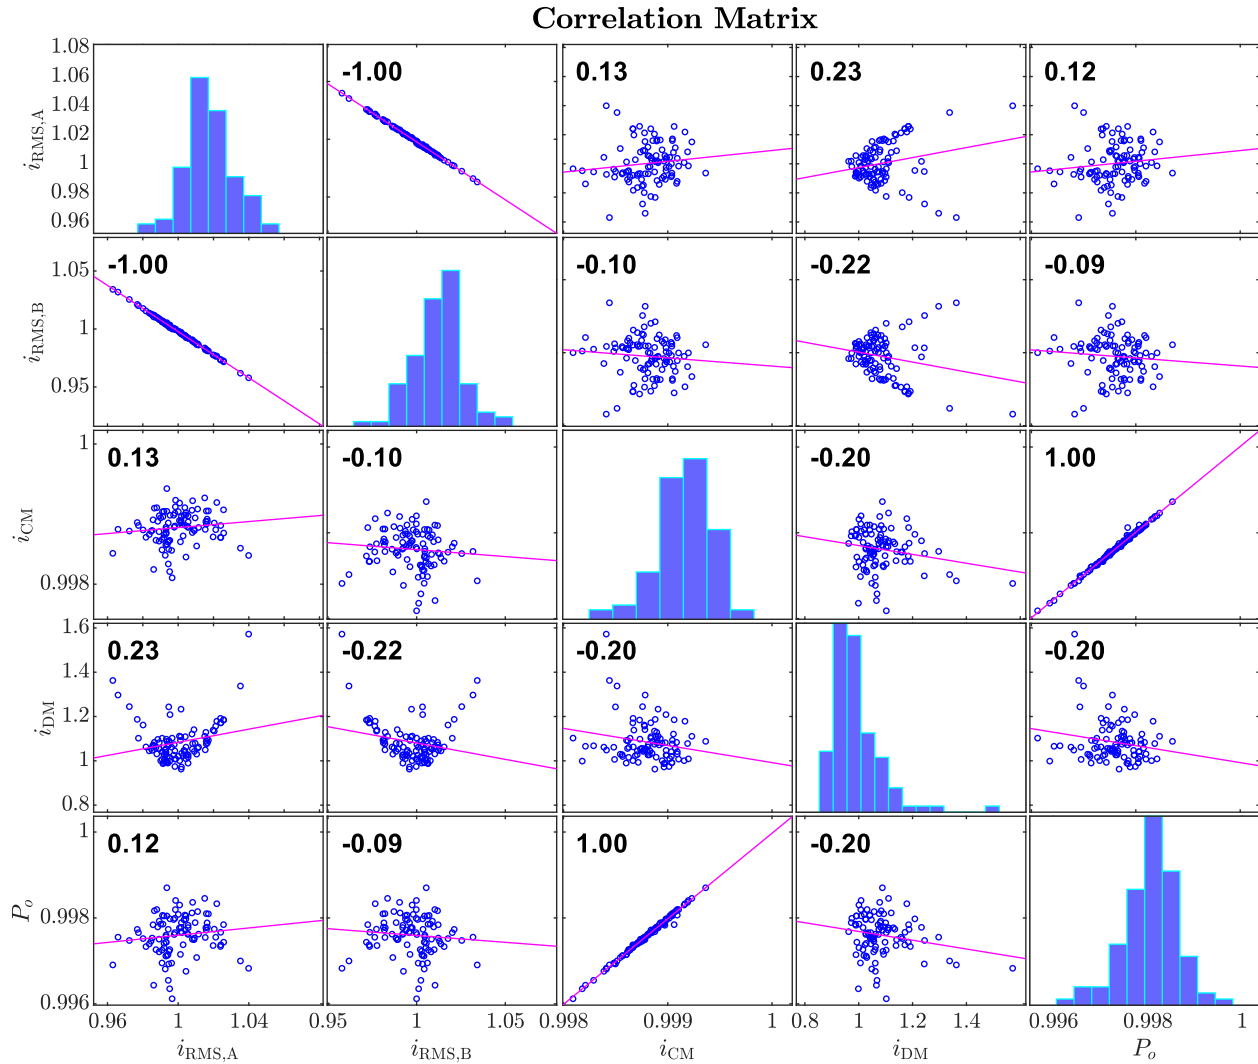


Figure 4.9: Correlation between key converter measurements from the Monte Carlo simulation.

- $i_{\text{RMS,A}}$ and $i_{\text{RMS,B}}$: the RMS values of the interleaved phases, normalized to the RMS values of the respective phase currents with converter parameters set to nominal.
- i_{CM} : the RMS of the common-mode component of the two phase currents, normalized to the RMS value with converter parameters set to nominal. Should be indicative of total output current and output power.
- i_{DM} : the RMS of the differential-mode component of the two phase currents, normalized to the RMS value with converter parameters set to nominal. Should be indicative of circulating current between the phases.
- P_o : total average output power, normalized to the output power under nominal converter parameters. Output power may be reduced, e.g., if a greater amount of phase

current is circulating (i.e., differential-mode), possibly due to parameter mismatch between phases.

There are three major takeaways from this matrix. First, the two phase currents are negatively correlated. This means that if parameter variations cause a decrease in output current from one interleaved phase, the relative current in the other phase increases. Additionally, this mismatch appears to have little effect on the circulating current (i_{DM}), instead only leading to diminished output current and power (i_{CM} and P_o , respectively). Although the correlation coefficients in some model outputs show marginal trends, it is clear from the spread of sample points that a linear trend is not actually present. Finally, where component variation does affect on current sharing, it is only on the order of 2-5%. This is highlighted in the left plot of Fig. 4.10 shows the variation in RMS output current between the interleaved phases, where only outliers (induced from simultaneous outliers in multiple parameters for a given scenario) produce this order of effect.

A second case was investigated to examine systemic parameter variation in a single phase. This situation may arise from manufacturing variations between individual modules, e.g. given lot to lot differences in component parameters. Additionally, variation in cooling across modules, or even interleaved phases on a single module, will drastically affect resistance of both passives and transistors alike. Using the framework in Chapter 3, a 5% to 25% reduction in mass flow of cooling air to one of the phases was simulated. The corresponding temperature elevation was used to adjust component resistances (increases up to 11%) on one of the interleaved phases accordingly. The results from this study are shown in the right plot of Fig. 4.10. Again, deviations from the nominal current stayed below 5%, though it is clear that current mismatch clearly increases directly with the temperature mismatch. Generally, increased losses in one phase will be offset by decreased losses in the other phase. However, the mismatch in current magnitude will have implications for harmonic cancellation; it can




Figure 4.10: Statistical results of simulations recording current sharing imbalance between phases of a dual-interleaved converter module given parameter variation.

be expected that, to the extent at which a mismatch exists, spectra that would otherwise cancel will instead remain at a magnitude equal to this mismatch.

4.4 Experimental Results

Several standards exist to characterize and qualify both radiated and conducted emissions from electronics. Radiated emissions include intentional transmitters, like WiFi or cellular communications, and unintentional radiators. In the latter, radio frequency interference is transmitted from circuit traces or conductors acting as antennae. Improper design of power converters, where switching creates high dv/dt edges with high harmonic frequencies that easily radiates from circuitry, can lead to significant radiated emissions. However, measuring emissions from such unintentional radiators requires complicated test equipment, including well characterized antennas, amplifiers and RF signal chain. Additionally, to ensure received emissions are only from the equipment under test, measurements must be made in an electrically quiet environment. This usually entails the use of a large, expensive shielded and anechoic chamber for the most robust measurements. Additionally, radiated emissions involve design insight beyond electrical parameters, up to and including physical component geometries, three-dimensional finite-element analysis, and enclosure construction. Therefore, this work will focus solely on conducted emissions, which are easier to measure and can still indicate overall EMC performance.

Most consumer devices and grid-connected equipment in the United States needs to adhere to Federal Communications Commission (FCC) standards, colloquially referred to as Part 15 [168]. Alternatively, CISPR 32 [169] (having absorbed the previous CISPR 22 standard [170]) represents an international variant of emissions limits and methods for consistent measurement. Indeed, most emissions standards are harmonized by the different regulating agencies or consortiums [171] to apply the same operating limits to the intended application. Conducted emissions measurements for any of these standards entails placing the equipment under test a specified distance from controlled ground planes, and connecting power cables to the ac mains using a LISN.

Electric vehicles and aircraft provide an entirely different electromagnetic environment. Aside from charging equipment, which would need to conform to the above CISPR or FCC requirements, the vehicle electric grid is self-contained. Additionally, the range of frequencies of interest also varies. For instance, frequency bands pertaining to shortwave radio, global positioning or avionics all have specific and stringent limits for unintentional radiators. Finally, the test setup for conducted emissions is fundamentally different. Whereas the standards above specifically indicate that the equipment under test be placed a specified height above the reference ground plane, in these tests the relevant standards dictate that the surface on which the equipment under test rests must be conductive and directly bonded to the reference ground plane.

To address the target application of electric vehicle drivetrains, the test setup of Fig. 4.11 was constructed to evaluate the interleaving strategy as applied to conducted emissions

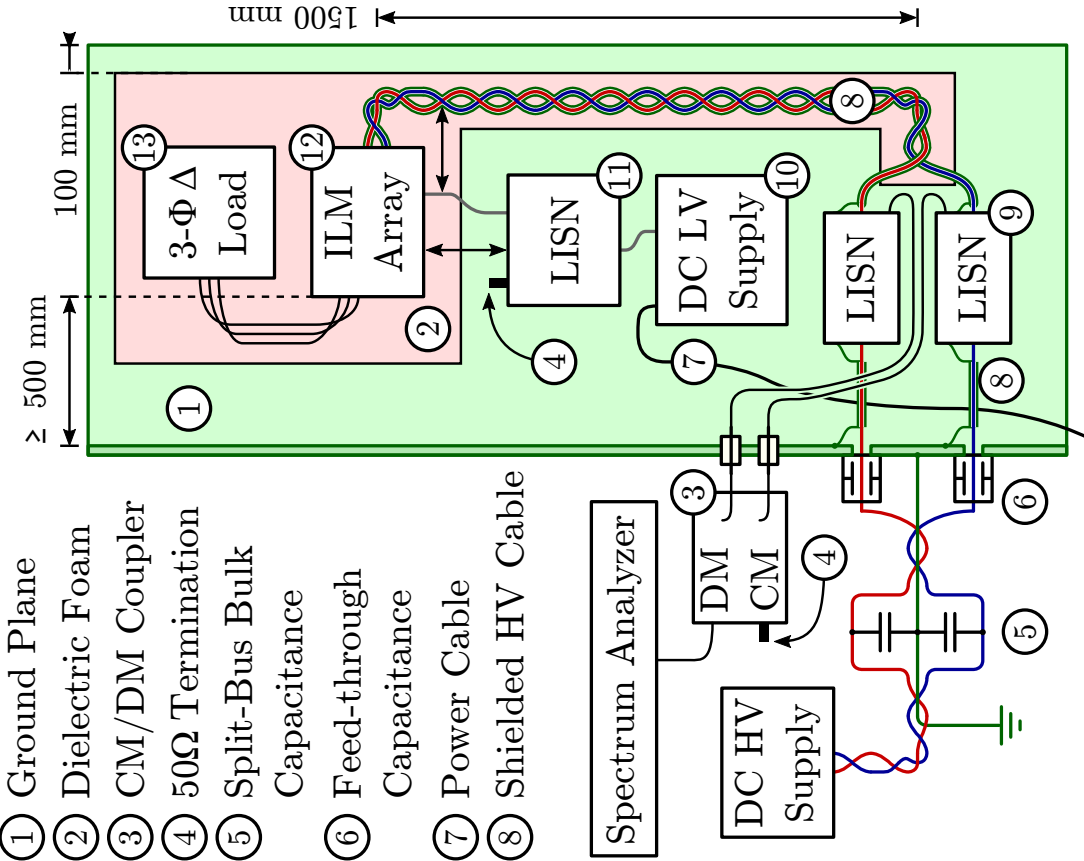



Figure 4.11: Electromagnetic compatibility pre-compliance experimental setup for measuring CISPR-25 conducted emissions.




Figure 4.12: Comparison of common-mode emissions measured at the HV LISNs with and without a common mode choke installed on the connection between the ILM array and the LV LISN.

defined in the CISPR-25 standard [172]. Two four foot by eight foot aluminum sheets were bonded to create a conductive surface for testing. Although the standard specifies measurements take place in a shielded enclosure, this configuration was considered sufficient for these initial “pre-compliance” tests to indicate relative performance improvements.

The two high-voltage (here, defined as 400 V to 1 kV) LISNs are 5 μ H 21702-5-TS-50-N from Solar Electronics Company, and the low-voltage (5 V) LISN is a 50 μ H LIN-115A from Com-Power Corporation. While this work focused solely on conducted emissions observed at the high-voltage (HV) connection, the low-voltage (LV) interconnection will also influence overall EMC performance. As such, it is important to draw attention to the impact of filtering on the logic power supply. Fig. 4.12 illustrates the impact of adding a broadband cable ferrite (with an impedance of approximately 200 Ω at 100 MHz) to the cable running between the LV LISN and the 5V terminal on the ILM array. Clearly, filtering of the LV terminals will have a direct impact on overall system EMC, especially for higher-frequency components that may more easily transfer across isolation barriers that nonetheless have non-negligible parasitic capacitive coupling.

While two LISNs are used to measure conducted emissions present on each of the high voltage lines, such a standard setup does not include a means to discern what emissions are differential-mode, and which are common-mode. Such a delineation is essential for informing an appropriate filtering strategy. To address this issue, a TBLM1 LISN Mate from Tekbox is installed between the two LISNs and the spectrum analyzer. Internal to this device, the signals from the two LISNs are mixed to either directly sum (for common-mode) or subtract (phase shift then sum, for differential-mode). The respective differential- or common-mode measurements were then made by selectively connecting the spectrum analyzer to the appropriate LISN Mate output.

Here, the device under test is a three-phase, segmented drive comprised of nine arrayed, dual-interleaved modules (ILMs), as shown in Fig. 4.13. This configuration corresponds to




Figure 4.13: Hardware photographs for the dual-interleaved module (ILM, on left) and three-phase, 18 converter array (right).

six paralleled ($P = 6$), 10-level ($N = 10$) FCML inverters (or three modules as described in Chapter 2) per phase-leg, or 18 arrayed converters in total. Correspondingly, each inverter is labeled by phase and converter number: A1...A6, B1...B6, C1...C6. Each converter module synchronizes phase-shifted fundamental and switching frequencies via optical communications from a global controller [173] (not shown) using the scheme from Section 4.1.3.

The inverter was supplied with a 400 V DC bus and commanded to drive a 950 Hz fundamental at 95% modulation into a 25Ω balanced, delta load. This frequency serves as a convenient test case corresponding to high-electrical-frequency machines while also remaining in a high-accuracy regime for the WT5000 meter used to record DC-AC efficiency. Measurements of the drain-source voltages of the low-side switches closest to the dc bus (S_{9B} in Fig. 4.1) were used to validate correct interleaving at the input. Fig. 4.14 shows these measured waveforms for the six converters within Phase A with the correct phase shift of $\pi/3$. Additional measurements also confirmed the correct phase shift of $2\pi/3$ between the array phase-legs. Further measurements at the switch node (v_a in Fig. 4.1) verified both three-phase operation, shown in Fig. 4.15a, as well as repeated dual-interleaving at the output, shown in Fig. 4.15b. This confirms the limitations imposed by the choice of P and $N - 1$ described in Section Section 4.2.

As the array assembly in this test setup was designed to facilitate ease of access to electrical nodes for debugging, both thermal and emissions performance were limited in this work; Future packaging improvements are expected to facilitate the full 100 kW output power target, while integration of shielding and filters will support better EMC. Nonetheless, the conducted EMI measurements for this configuration in Fig. 4.16 will still inform design of such filters, and highlights the benefits of interleaving. For instance, without interleaving between modules, Fig. 4.16a shows that significant harmonic content is present at many of the harmonics of the switching frequency. Note that, given each module is still dual-interleaved locally, content at odd harmonics is nonetheless reduced or eliminated compared


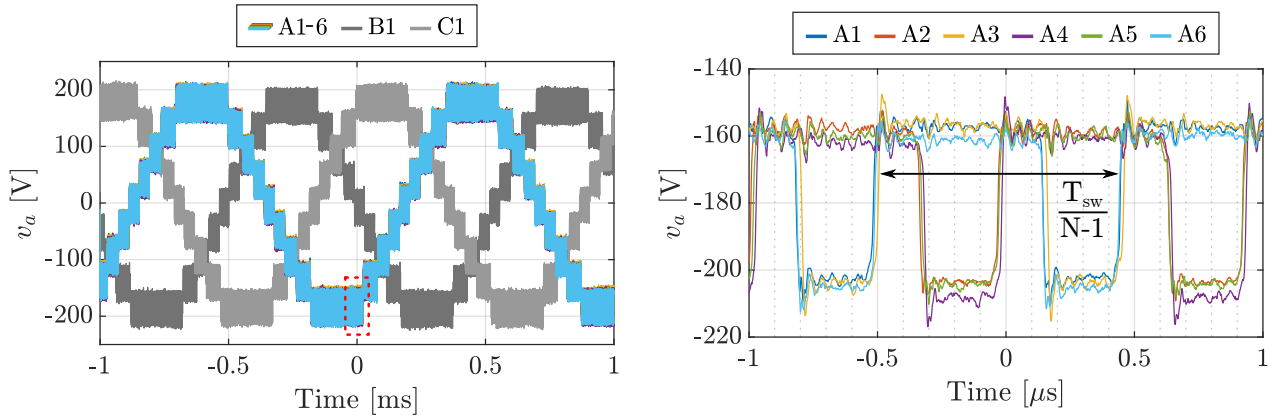


Figure 4.14: Six v_{ds} waveforms of the S_{9B} low-side switches (closest to the DC bus) measured for the converters of Phase A illustrate that 6-fold interleaving is achieved at the input.




(a) Switch node (v_a) waveforms for eight of the individual converter legs: all six interleaved modules within Phase A, as well as one measurement each from Phase B and from Phase C for reference.

(b) Zoomed view from the highlighted region above. The two sets of three overlapping waveforms indicate that only dual-interleaving is achieved at the output in this configuration.


Figure 4.15: Experimental waveforms at 400 V DC and 6 kW output power for the 3-phase array comprised of nine dual-interleaved FCML modules (16 arrayed converters).

to the spectra of the single FCML converter in Fig. 4.7. This is expected from the analysis of the dual-interleaved case in Fig. 4.6.

With 6-fold interleaving enabled across the arrayed modules, Fig. 4.16b shows that most non- $6m$ harmonics can be either eliminated or otherwise reduced by approximately 20 dBuV. As expected, the amplitude of the $6m$ harmonics increases slightly as the interleaving effectively shifts energy into these frequencies. The content at $2f_{\text{eff}}$ is particularly elevated as this is both a $6m$ harmonic of f_{sw} , and the repeated dual-interleaving at the output concentrates energy in this spectrum. For this converter setup, interleaving also more significantly attenuates CM harmonics compared to DM harmonics. This is possibly due to the use of a split-bus DC supply, and provides valuable insight into future filter design.



(a) Measured conducted emissions with no synchronization across arrayed, dual-interleaved modules.



(b) Measured conducted emissions with 6-fold interleaving at the input and associated synchronization of dual-interleaved module outputs.

Figure 4.16: Common-mode and differential-mode conducted emissions measurements of the three-phase segmented drive operating at a 400 V DC bus and 6 kW output power, measured at the HV terminals of the test setup.

4.5 Chapter Remarks

This section reported progress on a 100 kW-scale segmented drive comprised of nine arrayed, dual-interleaved, 18.9 kW FCML inverter modules. These modules were commanded and synchronized using a hierarchical control and communication scheme to demonstrate the benefits of interleaving across independent inverter modules. Experimental results were collected on a CISPR-25 pre-compliance test fixture, with CM and DM spectral performance obtained across several bus voltage and output power operating points. Future work should consider the inherent restrictions on interleaving at both the input and output of the array when evaluating choice of levels and number of paralleled modules for the full system design.

As the module and array design matures, future integration activities will further inform what filtering is necessary on the complete design. Additional shielding as part of the array enclosure will both help mitigate radiated emissions but also provide a new path for common-mode return currents. The magnitude of these interactions will need to be analyzed with respect to the corresponding enclosure geometries. However, the interleaving presented in this section is expected to continue to provide a system benefit – though the exact reduction of filter passives will be determined in future works.

Finally, it is worth noting again that the modulator used in this work was based on a single, trailing-edge ramp. However, prior work in [164] has shown that the up-down, triangular carrier results in a modulated output with better harmonic characteristics. Future work should be sure to use the this triangular carrier, as the ramp was only used in this work due to the ease of leveraging code for an existing modulator peripheral.

Chapter 5

System Reliability

Previous chapters have discussed module design, fabrication and testing, with a specific focus on scalability, high power density and high efficiency. However, for the approach to have any practical application, safe and cost-effective operation must be feasible over the operational lifetime of the hardware. Thus, having fulfilled many of the promises of the approach outlined in Section 1.3, this chapter now addresses the question of converter module and system reliability. In the aerospace industry, the complete aircraft design must be robust; fail-safe systems and multiple redundancies [174] are employed to keep catastrophic failure (i.e., loss of aircraft) rates to below 10^{-9} per flight hour [175]. For automotive applications, failures in the drive system are potentially less catastrophic, however industry expectations on failure rate are nonetheless stringent at 10^{-5} per 100,000 kilometers traveled [176].

As much of the research discussed thus far represents early development on a promising yet nascent approach towards solving industry challenges, it is necessary to provide context as to when the failure rates above apply. Fig. 5.1 illustrates a notional failure rate, $z(t)$, over the entire lifetime, t , of a generic component [177]. The *burn-in* period, with characteristically high initial failure rate, represents the early stages of development and operation. Here, undetected defects in individual components or systemic flaws in a design or manufacturing process are the predominant source of failures. These defective components, assemblies and systems would be largely eliminated before entry into service through qualification programs




Figure 5.1: The bathtub curve.




Figure 5.2: Panelized and contract assembled array of switching cell daughter boards.

or accelerated life testing. At the same time, as the product and production process evolves, this burn-in period can be reduced (or similarly, yield can be improved) when root causes of failures are traced back to specific design choices or manufacturing steps.

Given that the approach of this work is unconventional, as it were, many of the failures in extant hardware can likely be ascribed to this early burn-in period. Yet, many lessons and assembly improvements have already been introduced throughout the development culminating in this work. For instance, ceramic capacitors – lauded for their high-energy-density in Chapter 2 – are susceptible to mechanical stresses and cracking, a dominant root cause of failure [178, 179]. Hand soldering of these components, the procedure for early concepts in [66, 67, 73], led to high thermal stress, inconsistent solder application, and incomplete cleaning – all potentially contributing to the apparent increased failure rate, especially at high voltage.

These assembly risks were significantly mitigated by contracting the switching cell daughter board (groups of transistors, gate drivers and the ceramic flying capacitors) assembly to a manufacturer with an automated placement, reflow soldering and X-ray verification process for the converter in [80]. A photo of the contract assembled, panelized switching cells is shown in Fig. 5.2. This change improved the consistency of the manufactured cells, though ignored common design guidance on minimum clearances between the capacitors and circuit board edges. As such, the boards needed to be laser-depaneled to minimize reduce mechanical stress. Even with this added step, the close proximity of the capacitors the edge of the daughter boards nonetheless led to mechanical stress during the additional reflow stage when the daughter boards were installed onto the converter motherboard. The charring from the laser-depaneling also appeared to introduce short-circuits, likely from carbonized residue that remained even after light cleaning – introducing a new failure mode.

This work took steps to transition the entire converter build to an automated assembly process. As described in Chapter 2, switching cell daughter boards were eliminated altogether. Where ceramic capacitors still needed to be placed at the converter board edge,

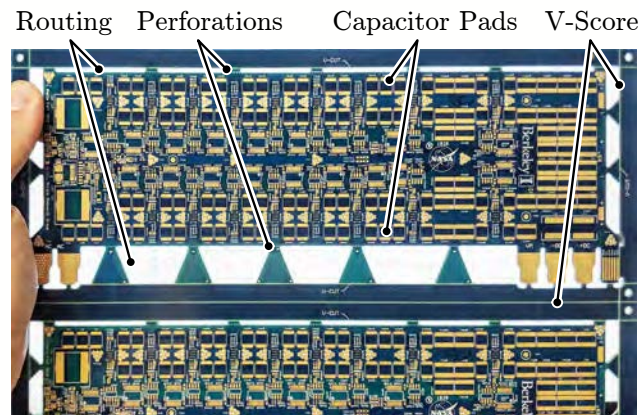


Figure 5.3: Panelized converter module circuit boards with features (routing and breakaway tabs) to mitigate mechanical stress from post-assembly depaneling.

additional features were added to the design to improve assembly yield, shown in Fig. 5.3. These included routed tabs to hold the nominal board shape in place when fixtured to assembly rails and breakaway perforations located away from the sensitive ceramic components. Furthermore, the assembled boards were subjected to a consistent bring-up process and validation checklist before fully committing them to operational testing. While subsequent testing eventually required manual rework and hand soldering, the commitment to a manufacturing process and handling guidelines significantly reduced the uncertainty surrounding early converter failures and enabled the range of experiments presented in both this work and in [173].

On the opposite end of the timeline in Fig. 5.1 lies the *wear-out* period. Although the converter hardware described in this work has not been tested to the extent to reasonably assume failures have been due to wear-out, there is an understanding of how wear-out manifests for the constituent components. Most mechanisms are treated as temperature driven [180]. For instance, repeated cycling of operating temperatures can cause thermo-mechanical fatigue via repeated thermal expansion and contraction. This is particularly relevant for solder joints and power modules with bond wires [181–183]. Higher temperatures will also accelerate many chemical processes, including degradation – notably in electrolytic capacitors [184], as well as ceramic capacitors [185], film capacitors and silicon semiconductors [183]. This strong correlation between thermal performance, peak and total energy dissipated, and failure rates drives many “design for reliability” studies in the literature, which then often entail minimizing losses [186–188] or hot-spots [181, 189].

However, not all wear-out mechanisms are of thermal or thermo-mechanical origins. Thus, broader *physics of failure* analyses, which consider origins of degradation, have gained popularity for more accurately describing certain failure modalities [190, 191]. For instance, while Arrhenius relationships can capture dominant modes of semiconductor degradation [192], a survey of research on gallium nitride power semiconductors indicated a disparate range of activation energies used in calculating the Arrhenius term [193]. Indeed, a relia-

bility model based on temperature alone will likely underestimate the failure rate of these devices [194]. Instead, transistor wear-out, evidenced by e.g., increased gate current or on resistance, appears to be largely driven by an inverse piezoelectric effect [195] and a voltage- and current-dependent trapping effect [196]. As such, models that can incorporate several interdependent stress parameters, such as the Boltzmann-Arrhenius-Zurkov model [197], may prove to be the most accurate approach to model wear-out. Nonetheless, practical performance in a power converter will ultimately depend on the exact design, assembly process, application and use case. Thus, physics of failure methods will need to be complemented with long-term empirical testing and statistical analysis [198].

In the two extremes of converter lifetime described above, process control and qualification can be used to improve burn-in performance, while lifetime measurement and maintenance can manage wear-out near end-of-life. Although these strategies can also improve outcomes during the *useful life* period of a converter, the failure rates cited at the beginning of the chapter largely pertain to “chance” occurrences in this period. In other words, failures arising outside of the scope of qualification or maintenance plans (e.g., cosmic rays) must either be statistically unlikely (e.g., outliers in the qualified population) or otherwise adequately covered (e.g., through system redundancy). In the remaining sections, the term *reliability* will describe whether either a converter module or the entire converter array can continue to provide rated operation during a fault condition or after such a failure of some kind. In the following sections, this will be presented through three specific examples:

1. the ability of a converter module to withstand exogenous faults, e.g., impedances, currents or voltages outside of nominal ranges seen at the load,
2. the specific failure mechanisms that arise within a converter module and possible mitigation schemes at the component level, and
3. the appropriate selection of levels or redundant modules to yield an acceptable system failure rate for the intended operation.

Although the proposed topology, as well as many of the constituent electronic components, are still early in the development life cycle, these examples will nonetheless illustrate progress towards addressing the demanding failure rate requirements for vehicle applications. Furthermore, the critical analysis can be incorporated into future design studies. This way, the methodology presented thus far can be augmented to add an aspect of “design for reliability” into the performance metrics.

5.1 Short-Circuit Fault Mitigation

Portions of this section are adapted in part or in whole from [199] and [200].

As discussed, robust converter behavior during transient and fault conditions must also be demonstrated to ensure reliable operation [111].




Figure 5.4: Single-leg representation of an N -level flying-capacitor multilevel converter with a split dc bus. The output filter capacitor is connected back to the dc-link midpoint. A short-circuit occurs with a sudden connection of fault impedance, Z_F .

Converter start-up is one such condition, because the total dc bus voltage often exceeds the rating for individual switches in the most efficient, power-dense designs [56]. As such, switching operation must ensure that flying capacitors charge in a controlled sequence and, in the case of a grid-connected inverter, that line currents do not overcharge capacitors through device reverse conduction. To this end, several solutions have been analyzed and experimentally demonstrated [1, 63, 201–203] at bus voltages up to 1 kV to overcome this concern.

Other research has focused on dynamic response to dc link transients and load steps. Several analytical models describing flying capacitor dynamics under various modulation schemes and loads have been developed and corroborated in hardware demonstrations [64, 116, 165, 201, 204, 205]. Additionally, converter response to modulation and load steps has been investigated on a testbed used to automate such analysis [121].

Fault-tolerant operation under internal component failures has also received significant attention. For instance, in the event of a switch short-circuit fault, if adjacent devices can survive the energy from capacitor charge redistribution [206] and all surviving devices can withstand the higher drain-source voltage associated with lower level-count operation, a converter can be reconfigured continue operating in a degraded state [71, 207, 208]. Conversely, it is well understood that an open-circuit switch failure cannot be recovered via switch reconfiguration [206] without additional bypass [209], disconnect [210], or dc bus mid-point connection circuitry [211, 212].

Recent work has also sought to elucidate FCML behavior in the event of an ac-side short circuit fault, e.g., as a result of a motor winding failure, accumulation of dust and moisture or an electrically-close grid fault[200]. While the fault-tolerant studies above may provide a path to remedial action under these conditions, the proposed strategies are only applicable once such a fault had already resulted in a component failure within the converter. Therefore, this section seeks to complement this existing literature by proposing a method to *prevent* such failures resulting from ac-side short circuit conditions. The section proceeds as follows: Section 5.1.1 reviews the model presented in [200] and applies it to several recently proposed FCML inverters to establish key operating limits and critical response times during the fault. Next, Section 5.1.2 proposes strategies to rapidly detect the short-circuit fault, while Section 5.1.3 describes a control mode to recover from and ride through a fault once detected. Then,

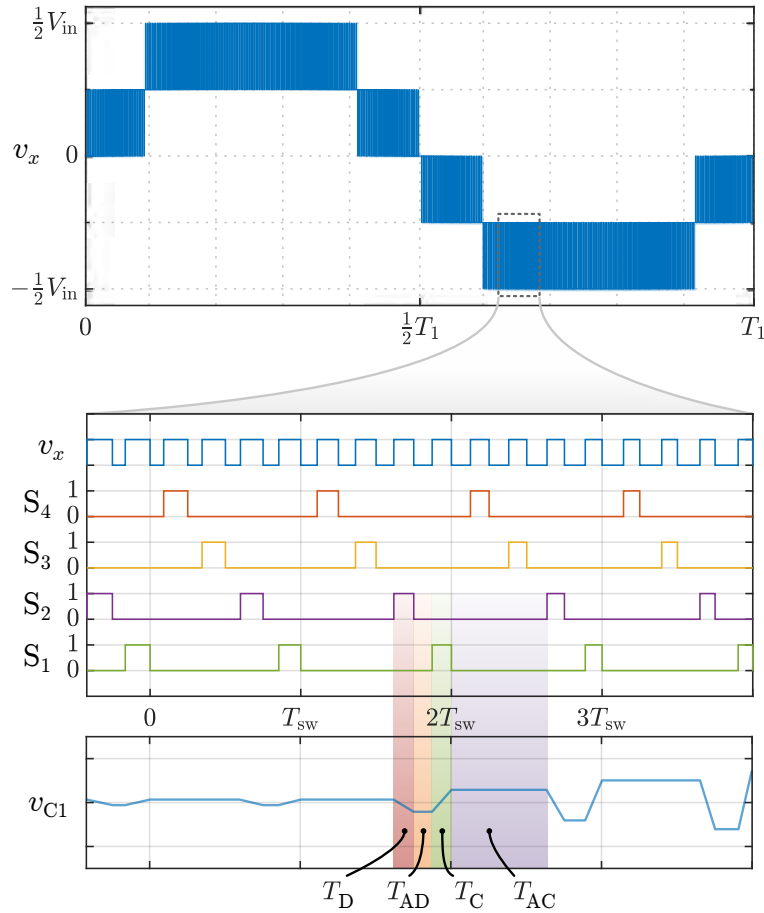


Figure 5.5: Switching waveforms of a 5-level ($N=5$) inverter. v_x is plotted over a fundamental period T_1 , while the zoomed view illustrates the high effective frequency and gate signals obtained using PSPWM. The bottom plot illustrates the charge and discharge of a flying capacitor over each switching period, with a short-circuit fault introduced between between T_{sw} and $2T_{sw}$.

Section 5.1.4 demonstrates the proposed strategy on a configurable platform with component ratings characteristic of high-power-density designs, and thus requiring high-speed response to short-circuit faults.

5.1.1 Short-Circuit Dynamics and Critical Timing

The N -level FCML inverter phase-leg in Fig. 5.4 consists of $N - 1$ complementary switch pairs and $N - 2$ flying capacitors. The top plot of Fig. 5.5 shows the multilevel output waveform at the switching node, v_x . Here, the duty cycle of a 5-level converter is modulated over a full fundamental period, T_1 . While several modulation strategies have been proposed for the FCML[64, 205, 213–215], phase-shift pulse width modulation (PSPWM) is often employed, since it provides balanced charging of the flying capacitors and a high effective

output frequency ($N - 1$ times the switching frequency), and can be adapted to allow space-vector techniques in three-phase systems [216]. The zoomed plots in the figure show how this effective switching frequency emerges from the phase-shifted switching actions. Given that a flying capacitor C_f charges or discharges based on the state of the adjacent switch pairs, a switching period can then be divided into the four intervals, as illustrated for C_1 : a charging period T_C , a discharging period T_D , and two disconnected states T_{AC} and T_{AD} that occur after charging and discharging, respectively. These times are functions of N , the switching period T_{sw} , and the duty cycle of a switch D where $D \in [0, 1]$. In the last plot of Fig. 5.5, the charging and discharging actions are balanced until an ac-side short circuit occurs between T_{sw} and $2T_{sw}$.

During a short-circuit fault, two factors must be considered: how quickly the current through L increases, and the rate of the mismatch between the charging and discharging currents a flying capacitor experiences during the times T_C and T_D , respectively. The consequence of this mismatch is illustrated in Fig. 5.5 after the fault occurs: as the short circuit current through the output inductance increases over the fault duration, the charging and discharging of the flying capacitors will no longer balance and voltage excursions will increase until component voltage ratings are exceeded. Thus, it will be useful for the designer to understand how choice of operating parameters and passive sizes affect the time until overcurrent or overvoltage failures can occur. A full derivation of current and voltage dynamics during a short-circuit fault is available in [200]; the following will describe the practical application of these models to calculate these critical mitigation times.

The overcurrent condition is directly tied to the minimum time until either component thermal or saturation current limits are exceeded. Though the controller is nominally modulating D such that $v_x = (2D - 1)V_{in}/2$, D can be safely treated as constant over the short time-scale of the fault. The evolution of inductor current as a function of duty cycle can then be expressed as

$$i_L(t) \approx \frac{(2D - 1)V_{in}}{2L} \cdot t + I_{PF}, \quad (5.1)$$

where I_{PF} is the pre-fault current. This linear approximation can also be easily adapted to include a current-dependent inductance when using soft-saturating inductors. It can readily be determined that, assuming a stiff dc bus, the worst-case short-circuit current slew will occur at $D = 1$ (or symmetrically in the split-bus system, at $D = 0$). Considering any component peak current limit, I_{max} , time until this current limit is reached is then

$$t_{crit-i} = \frac{2L}{V_{in}}(I_{max} - I_{PF}). \quad (5.2)$$

The overvoltage condition is somewhat more complicated. Unlike the case with overcurrent, the worst-case duty cycle, D_{wc} cannot occur at $D = 1$ as no flying capacitors are connected at this duty cycle. However, parametric sweeps of the duty cycle in short-circuit simulations do indicate the worst-case conditions will occur in $1 - \frac{1}{N-1} < D < 1$, i.e., the

Table 5.1: Table of key parameters and critical fault clearing times for recently published converters.

Ref.	N	f_{sw}	V_{in}	P_{out}	$I_{out,pk}^a$	D_{wc}	ave. C_f	L	V_{max}	R_{Dson}	t_{crit-i}	t_{crit-v}
[60]	5	60 kHz	1100 V 1200 V	5.0 kW	20.4 A	0.762 0.764	1.2 μ F	60 μ H	650 V	25 m Ω	10.9 μ s 10.0 μ s	19.0 μ s 15.8 μ s
[118] ^b	6	144 kHz	400 V	7 kW	21.1 A	0.800	2.6 μ F	22 μ H	150 V	7 m Ω	26.3 μ s	20.8 μ s
[62] ^c	9	115 kHz	1000 V	9.7 kW	20.2 A	0.884	1.6 μ F	5 μ H	160 V	10 m Ω	1.8 μ s	0.2 μ s
[2]	10	85 kHz 115 kHz	1000 V	16.8 kW 18.9 kW	36.3 A 40.9 A	0.889	1.7 μ F	7.5 μ H	200 V	8 m Ω	2.7 μ s 2.6 μ s	1.0 μ s 1.9 μ s
Sec. 5.1.4 ^c	5	115 kHz	370 V	1.1 kW	13.3 A	0.771	2.8 μ F	7.5 μ H	160 V	10 m Ω	7.6 μ s	5.3 μ s
	6					0.800					7.6 μ s	7.8 μ s

^aPer-phase for interleaved designs.

^bThe unipolar dc supply and unifier used in this reference would require a modified version of the approach presented in this work (e.g., a different definition of D throughout the fundamental cycle). To lend use of the parameters and operating conditions directly to this analysis, the converter is instead assumed to also have a split dc bus to the effect that the output voltage rating is correspondingly halved (i.e., 120 V_{ac} instead of 240 V_{ac}).

^cThis reference lists an EPC2034 voltage rating of 200 V; however, calculations in this work use 160 V per the manufacturer errata.

uppermost level¹. The capacitor charging dynamics for this range of D are found by integrating the linear estimate of inductor current from (5.1) into a flying capacitance, C_f , for the charge intervals illustrated in Fig. 5.5 over several switching periods. For the uppermost level where D_{wc} is found, these intervals are and defined as:

$$T_C = T_D = T_{sw}(1 - D) \quad (5.3)$$

$$T_{AC} = T_{sw} \left(\frac{1}{N-1} - D \right) \quad (5.4)$$

$$T_{AD} = T_{sw} \left(D + \frac{1}{N-1} - 1 \right) \quad (5.5)$$

With T_{sw} typically much shorter than the modulation period, D can again be safely treated as constant in this analysis.

$$t_{crit-v} = \left(\frac{T_C}{T_C + \tau_f} \right) \left(\frac{T_C}{2} - \frac{T_{sw}}{N-1} - \frac{2L}{(2D-1)V_{in}} \left[I_{PF} \cdot \left(1 + \frac{\tau_f}{T_C} \right) + \frac{C_f}{T_C} \left(\frac{V_{in}}{N-1} - V_{max} \right) \right] \right) \quad (5.6)$$

After deriving the charging and discharging of adjacent flying capacitors, the corresponding voltage stress applied to the series-connected, off-state switch can be calculated. A key result of this analysis – how quickly a fault must be cleared before the switch stress exceeds the switch rating or other limit, V_{max} – is provided in (5.6) for unity power factor, with $\tau_f = R_{DSon}C_f$. As the focus of Section 5.1 is on the protection scheme and fault ride-through, the full derivations are omitted here. However, both critical times in (5.2) and (5.6) were verified in simulation, and a broader discussion of the modeling, as well as a detailed model for use with non-unity power factor, is available in [200].

These models can serve to provide critical assessment of extant designs and how quickly a controller must respond to adequately protect them. For example, the FCML has been proposed for high-speed motor drives [62], bi-directional EV chargers [118] and grid-tied photovoltaic inverters [60], with hardware demonstrations showing high power densities and efficiencies possible with the topology. Given the descriptions in the cited works, Table 5.1 summarizes results of the model above with respective critical short-circuit detection and mitigation timing requirements based on the peak current and drain-source voltage ratings of the switches in each design. For each converter type, D was swept to identify the worst-case duty cycle D_{wc} requiring the shortest fault clearing time for each respective V_{max} .

Note that, by seeking to increase component utilization through operating devices closer to their ratings, the minimum time to detect and mitigate a short circuit fault becomes quite short – as evident for the design of [62]. Generally, the designs with higher level count, output inductance, flying capacitance and switching frequency will have greater overhead in critical mitigation times. These trade-offs, and the speed of available fault detection and mitigation

¹The duty cycles corresponding to negative currents (i.e., $0 < D < \frac{1}{N-1}$) will produce equivalent worst-case results. However, for simplicity, discussion will only be in the context of positive currents in this work.


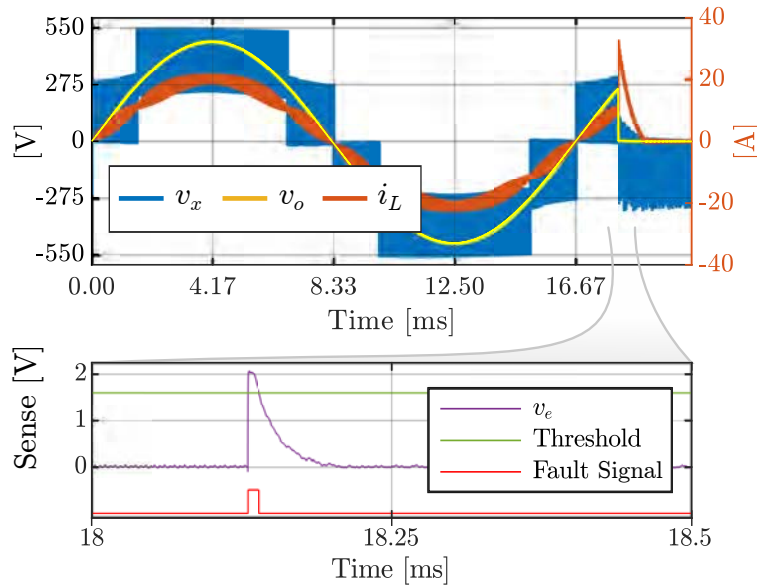
(a) Analog circuit for short-circuit detection based on v_x and v_o .(b) PLECS simulation of the circuit above for the 5-level converter from [60]. The detection time, t_{FD} , is on the order of microseconds, after which mitigating action (described in Section 5.1.3) is taken.

Figure 5.6: Analog detection of short-circuit fault from the mismatch between estimated output voltage and actual output voltage.

strategies, can be included in future design efforts to ensure that both performance and resiliency is considered.

5.1.2 Short-Circuit Detection

Methods of fault-detection for multilevel converters span a large body of work, including sliding-mode observers [217], machine learning [218, 219], detection from switch-node measurements [220], principle component analysis [221, 222] and others [206, 210, 211, 223]. While this research helps to address the complex challenge of detecting faults in these converters, many of these approaches target individual component failures or may not provide the rapid detection necessary for an ac-side short-circuit fault in an FCML inverter. Therefore, two short-circuit fault detection methods are proposed below.

5.1.2.1 Analog Detection

The first method relies on a comparator or logic circuit to detect an excursion in voltage or current beyond a fault limit. For instance, many current sensors have a programmable over-current limit that triggers a fault flag in a relatively short detection time. This scheme can be supplemented by applying additional criteria on current and voltage d/dt (if known) to reduce false trips during expected load transients [224].

Alternatively, significant deviations or dips in the difference between the output voltage (after the filter inductor) and the average switch-node voltage can also indicate a fault. In normal operation, the output voltage v_o , will approximately track the average value of the switch-node voltage v_x . An example circuit for this approach is illustrated in Fig. 5.6a. Here, v_x is low-pass filtered to create the output voltage estimate. Then, a rectifier circuit is used to compute the absolute error between the estimated and measured output voltage. The fault is detected when this difference exceeds a threshold, which is set to be greater than the mismatch during normal operation. This circuit is simulated under unity power factor for the converter and conditions in the first row of Table 5.1 in Fig. 5.6b.

If implemented entirely in analog circuitry (i.e., no analog-to-digital conversion stage), the detection time, t_{FD} , can be near instantaneous. Though filtering requirements and finite amplifier bandwidth will increase the detection time in a practical implementation, the circuit is can quickly detect severe faults. However, cases where the fault occurs at low modulation or near a zero-crossing lead to longer detection times, while the lack of a digital interface limits versatility.

5.1.2.2 Model-Predictive Detection

A second method that is more robust to these issues is the *model-predictive* approach [225]. This method requires digital sampling and control, and leverages measurements and predictions of the output filter state variables, i_L and v_o . The state-space model for the filter used in Fig. 5.4 is shown in (5.7).

$$\frac{d}{dt} \begin{bmatrix} i_L \\ v_o \end{bmatrix} = \begin{bmatrix} -\frac{R_s}{L} & -\frac{1}{L} \\ \frac{1}{C_o} & 0 \end{bmatrix} \begin{bmatrix} i_L \\ v_o \end{bmatrix} + \begin{bmatrix} -\frac{1}{L} & 0 \\ 0 & -\frac{1}{C_o} \end{bmatrix} \begin{bmatrix} v_x \\ i_o \end{bmatrix} \quad (5.7)$$

This model can be used in digital controller with a sample period, T_s , by discretizing the system with a zero-order hold:

$$\begin{bmatrix} i_{L,\text{pred}}(k+1) \\ v_{o,\text{pred}}(k+1) \end{bmatrix} = \mathbf{A}_{\text{ZOH}} \begin{bmatrix} i_L(k) \\ v_o(k) \end{bmatrix} + \mathbf{B}_{\text{ZOH}} \begin{bmatrix} v_x(k) \\ i_o(k) \end{bmatrix} \quad (5.8)$$

$$\text{where } \mathbf{A}_{\text{ZOH}} = e^{\mathbf{A}T_s} \text{ and } \mathbf{B}_{\text{ZOH}} = \int_0^{T_s} e^{\mathbf{A}\tau} d\tau. \quad (5.9)$$

Fault detection is then based on error between the output voltage predicted from the model and the next measurement:

$$v_{o,\text{error}}(k+1) = |v_o(k+1) - v_{o,\text{pred}}(k+1)|. \quad (5.10)$$

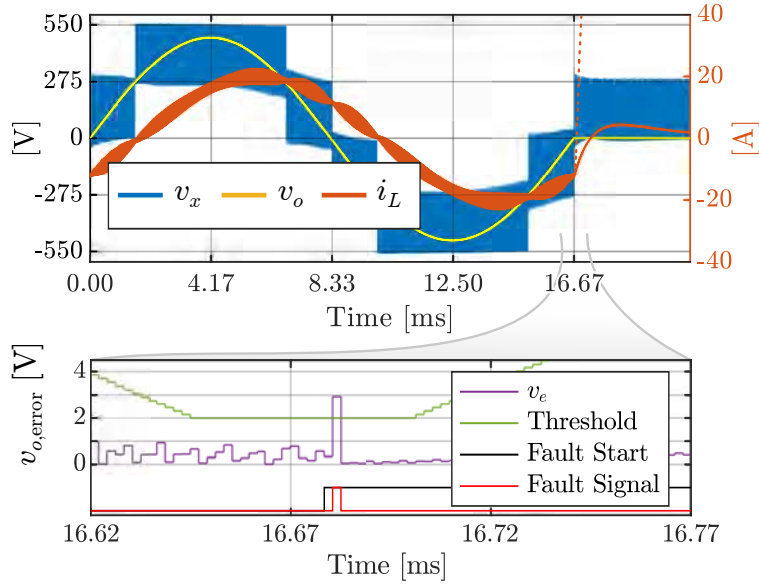


Figure 5.7: PLECS simulation of the model-predictive approach for the 5-level converter from [60] at power factor of 0.8. The detection time, t_{FD} , is the model time-step, T_s – after which mitigating, action (described in Section 5.1.3) is taken. For comparison, the dashed line shows the trajectory of the inductor current, i_L , without mitigation.

To reduce the requirements on sensing circuitry, the inputs $v_x(k)$ and $i_o(k)$ may be approximated as:

$$v_x(k) \approx V_{in} \cdot (D(k) - 0.5) \quad (5.11)$$

$$i_o(k) \approx i_L(k) - \frac{C_o \cdot (v_o(k) - v_o(k-1))}{T_s} \quad (5.12)$$

Though these approximations may reduce the accuracy necessary for closed-loop control of the state variables, they are sufficient for detection of unexpected transients.

This detection scheme is simulated in Fig. 5.7, where T_s is reasonably chosen to be twice the effective output frequency, i.e. $T_s = 2(N-1)T_{sw}$. Here, the converter from the first row of Table 5.1 is again operating at the same VA rating, but with a power factor of 0.8. This time, the fault is triggered near the zero crossing of the voltage (i.e., with the modulator reference close to $D = 0.5$), which had posed issues for the analog detection scheme of Fig. 5.6a. In this case, the fault is detected after one sample period – after which the model evolve without error as the new measurements of v_o are under the faulted condition. Note, in this implementation the threshold is adaptive, derived from the duty ratio. This allows for decreased sensitivity to expected load steps.




Figure 5.8: Representative time-reversed, complementary switching strategy to reset inductor fault current for $N=5$.

5.1.3 Mitigation and Ride-Through

Once a fault is detected through the means described in Section 5.1.3, the converter can enter a mitigation phase that drives the output current to zero while still maintaining capacitor voltage balance. While one approach could simply power the converter down, additional “internal balance booster” circuitry would be required to maintain voltage balance[111], implying additional components and power dissipation. However, if the converter remains online, the controller can proactively mitigate the fault by fixing $D = 0.5$ (i.e., the zero-output-voltage command). This general concept of driving the output voltage to zero during a fault in an FCML converter is analogous to closing the ground-referenced switch and opening the high-side switch in a two-level design. The added challenge in an FCML converter is that, compared to a two-level design, the converter must operate continuously to maintain this zero voltage – and capacitor voltage balance. A general description of such an operating mode is also suggested in [111], though performance is not experimentally verified.

In addition to the high-speed detection described in Section 5.1.2, other delays in control actions should also be considered. For instance, if the modulator employs regular sampling,

the peak current before mitigation, i_{peak} , is approximately

$$i_{\text{peak}} \approx I_{\text{PF}} + \frac{2(D_{\text{PF}} - 0.5)V_{\text{in}}}{2L} \cdot t_{\text{max}}, \quad (5.13)$$

$$\text{where } t_{\text{max}} = t_{\text{FD}} + T_{\text{PWM}}. \quad (5.14)$$

Here, D_{PF} is the pre-fault duty cycle and T_{PWM} is the PWM delay that depends on how the signals are generated. For regularly sampled modulation, T_{PWM} is generally T_{sw} for single-edge sampling and $\frac{1}{2}T_{\text{sw}}$ for double-edge sampling. In either case, T_{PWM} may dominate the response. Therefore, the switching frequency should be increased as much as possible at the start of zero-command mitigation. As FCML implementations in recent literature leverage FPGA control or high-speed PWM peripherals, and fast-switching, wide-bandgap devices, headroom for this increased switching frequency is often available. Indeed, the controller implemented on a MAX10 FPGA in this work approximates natural sampling; gate signals are updated instantaneously with D . With the high-speed clock set to 125 MHz, $T_{\text{PWM}} = 8$ ns.

Instead of simply commanding the reference to zero during the fault, the PWM operation may also be modified to reverse the effects of the capacitor unbalancing during the fault. The intuition is to apply switching actions that are time-reversed and complementary to the actions that occurred during the fault. From the perspective of the PWM generation, this may entail both inverting the carriers and running them in reverse while commanding the reference to the opposite of its pre-fault value. Practically, this may mean a digital counter in the PWM module changes from incrementing to decrementing, while from a modulation perspective, leading modulation momentarily becomes lagging modulation and vice versa. The gate signals for such a strategy implemented for a 5-level converter is given in Fig. 5.8. A PLECS simulation of this approach is shown in figure Fig. 5.9, with a subset of the simulation time highlighting the fault, protection and ride-through dynamics. This scenario used a sufficiently large output inductor (50 μH) to extend the fault duration and better illustrate the capacitor charge and discharge actions, though the same behavior can be observed for the values used in the experimental section.

Note that this simulation models few reactive parasitics and interactions between these and additional high-frequency load/short-circuit dynamics may introduce challenges in practical implementations. Furthermore, though the feedforward that drives the reference to the opposite of the pre-fault value is an intuitive approach, shaping the inductor current closed-loop may be a more appropriate solution. That said, given the time-scales discussed, the controller bandwidth to do so would need to be upwards of a megahertz. A final potential drawback for either protection approach is the increased dv/dt at the switch node necessary to drive the current back to zero. Nominal FCML operation limits the voltage levels to $\frac{V_{\text{in}}}{N-1}$, whereas the instant change in references discussed above would lead to a voltage step of $\frac{V_{\text{in}}}{2}$ to V_{in} at the switching node v_x . While this slew may be rate-limited via control, or compensated for in the output filter design, this trade-off between atypically high dv/dt and faster fault protection time should be considered when applying these techniques.

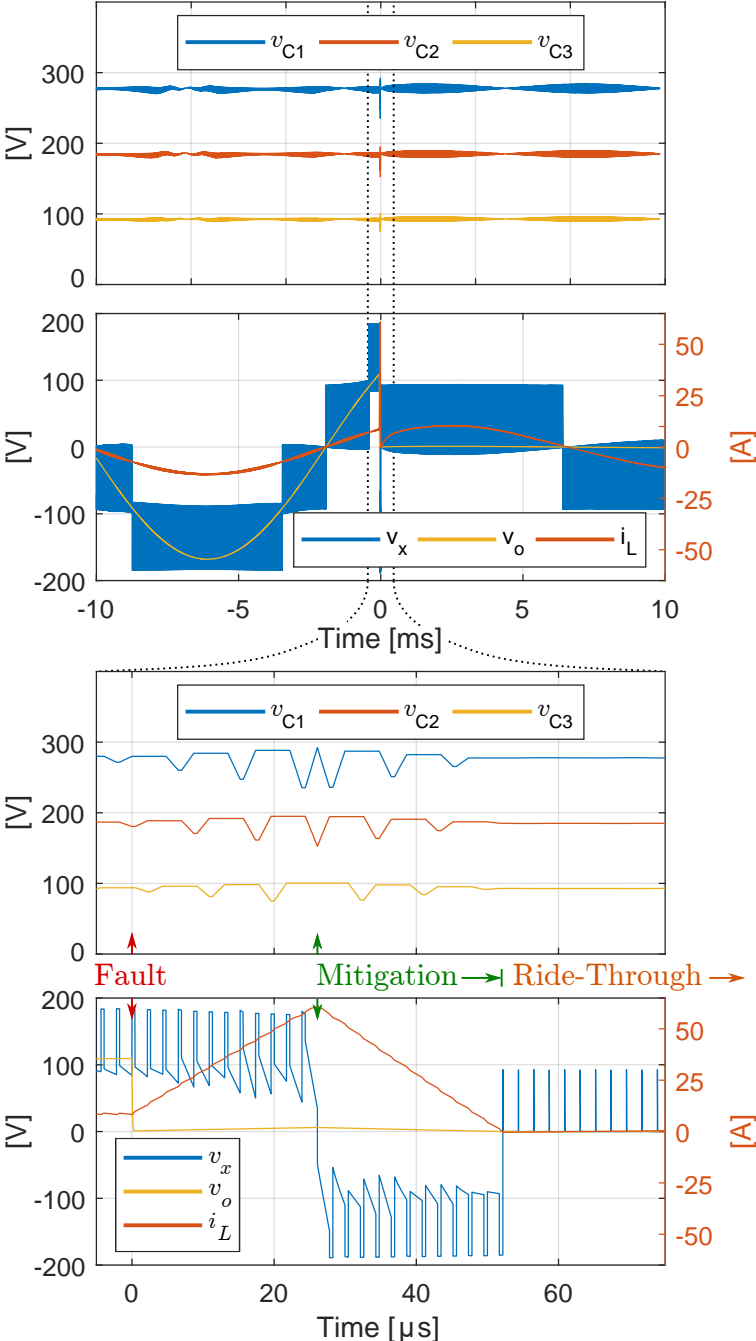


Figure 5.9: Simulation of a mitigation strategy that resets the inductor current via reverse capacitor charge/discharge actions.

After the mitigation action has taken place, i.e., the inductor current has been commanded to zero, the converter enters the ride-through phase. Depending on the nature of the load, the converter may wait until a line voltage is again present at the output (e.g., in the grid-following case) or actively determine whether the short-circuit has cleared (e.g., in a grid-forming or standalone case). This work assumes the latter, so the post-fault converter applies a low magnitude modulation during ride-through to detect when the load (or output current) is again indicative of the nominal impedance. In the absence of large startup current requirements, this nominal load impedance at the rated output voltage and power, Z_{\max} , can be specified. Then, if the converter can operate at sufficiently low modulation during a short so as not to exceed any output current limits (e.g., 3%), a series of output current and voltage measurements satisfying $i_{o,\text{meas}} \cdot Z_{\max} < v_{o,\text{meas}}$ indicates the fault has cleared and normal operation can resume.

5.1.4 Experimental Verification

The analog comparator detection scheme and the proposed mitigation strategy were experimentally tested for a dc-link voltage of 370 V – greater than twice the voltage rating of the transistors. The hardware used to demonstrate this fault, shown in Fig. 5.10, is a variation of the converter used in [1] and is configurable from 2 to 10 levels. While this converter is nominally dual-interleaved, only one phase was tested in this work, with the other left disconnected. Furthermore, while the converter has demonstrated high-efficiency and output powers in excess of 10 kW in [1], operation without heat sinking (for easy access to flying capacitor voltages) limited the demonstrations in this section to 1.1 kW.

An ACS733 current sensor is used to measure the inductor current, i_L . Fault detection is implemented using the analog overcurrent trip circuitry built into this component. Once the programmable current limit is exceeded, a sensor flag alerts the FPGA generating the converter gate signals to toggle the mode from normal operation to fault mitigation. The FPGA also reads digitized current measurements from the ACS733 at the effective output frequency. While these measurements are not used for model-predictive detection in this work, they are used to determine the start of the ride-through period (once the current has been reset) and the eventual restart of output power (once the fault has been cleared). For instance, the changeover between the mitigation action and ride-through is triggered when either a measurement between ± 0.5 A or a sign change (i.e., the zero-crossing) occurs.

To demonstrate the detection and mitigation strategy, the converter is first commanded to deliver power to a 12.5 Ω load at 950 Hz. The fundamental is chosen to be near the operating point of high-specific-power machines for drive applications [226, 227], but also illustrates fault recovery well on the plotted timescale. Then, an isolated switch toggles to momentarily short (with a nominal impedance of 80 m Ω) the converter output for a duration of 400 μ s. This ac switch is based on the design from [121], but instead uses two back-to-back NVHL040N65S3F MOSFETs connected in common-source. The resulting output current and voltages during this transient are shown in Fig. 5.11 for both a 5-level and 6-level FCML. Note that the 5-level FCML nominally has infinite balancing time-constants when

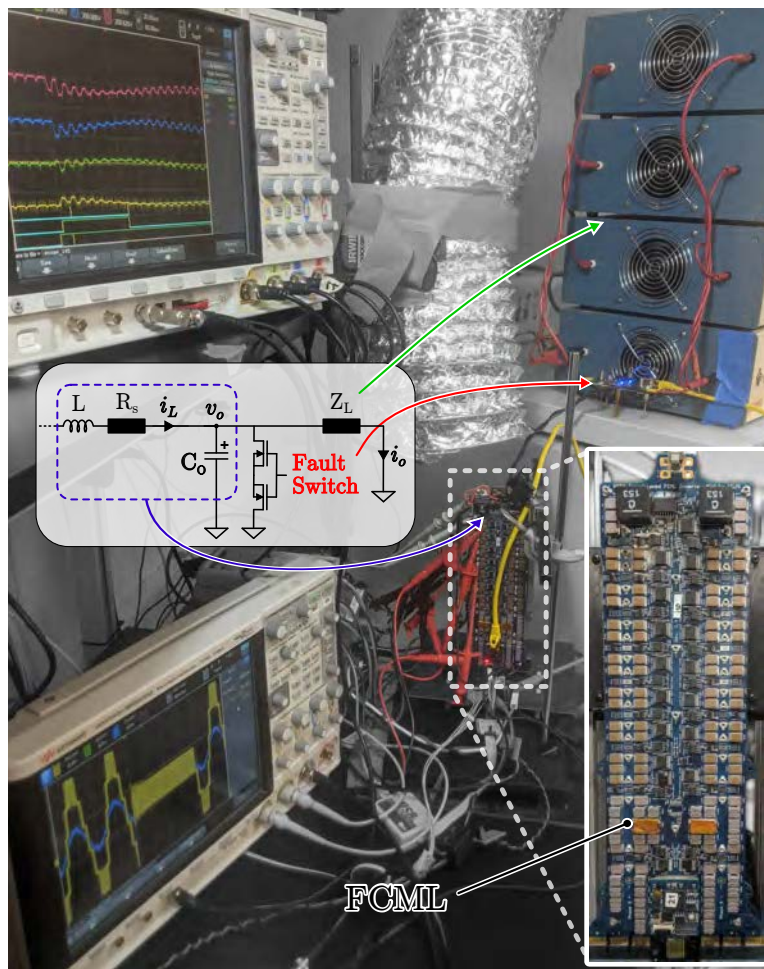



Figure 5.10: Experimental setup, including the load, the switch to simulate faults and an air cooling system. The inset highlights the converter hardware from [1] used in the tests.

$D = 0.5$ [76, 201], as is the case for all odd-level converters. This means that when idle-mode operation is activated, the flying capacitors will not re-balance to their nominal values. Nevertheless, if fast fault detection is provided the change in the steady-state flying-capacitor voltage is negligible – as seen in Fig. 5.11a.

After the fault and rapid mitigation action completes in both cases, a second, negative current peak can be observed. This is not a part of the system response, but rather the outcome of the reduced modulation ride-through discussed in Section 5.1.3. Here, modulation is set to 3% of the pre-fault value – sufficient to induce a measurable output current (and detect whether a short is still present) but not exceed the converter rating. This current can be seen to continue sinusoidally until the fault is cleared after 400 μs . The converter monitors the output current over this time and resumes full power once the current returns to the expected range for the nominal $12\ \Omega$ load and 3% modulation over a conservative number of cycles. This restart occurs at around 2 ms after the start of the fault.

(a) Measured waveforms during operation, short-circuit fault, ride-through and restart of the converter in a 5-level configuration.



(b) Measured waveforms during operation, short-circuit fault, ride-through and restart of the converter in a 6-level configuration.




Figure 5.11: Experimental results for the hardware demonstration of fault detection and ride-through at 1.1 kW output power.

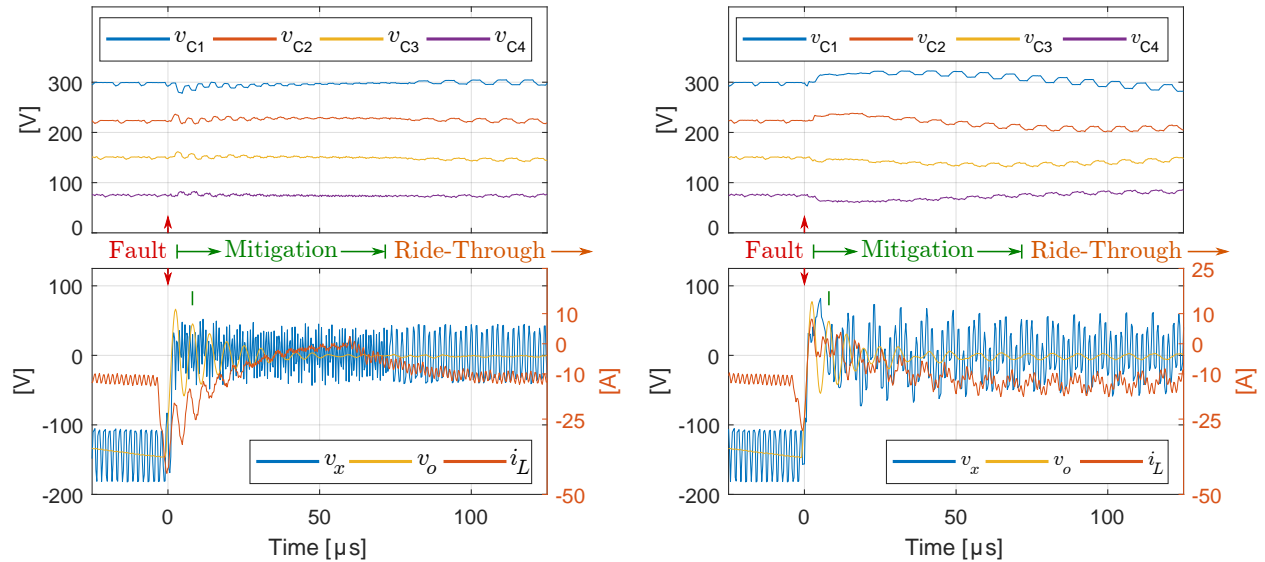
(a) Strategy that drives the switch node, v_x , to zero. (b) Strategy that reverses the effects of the fault.

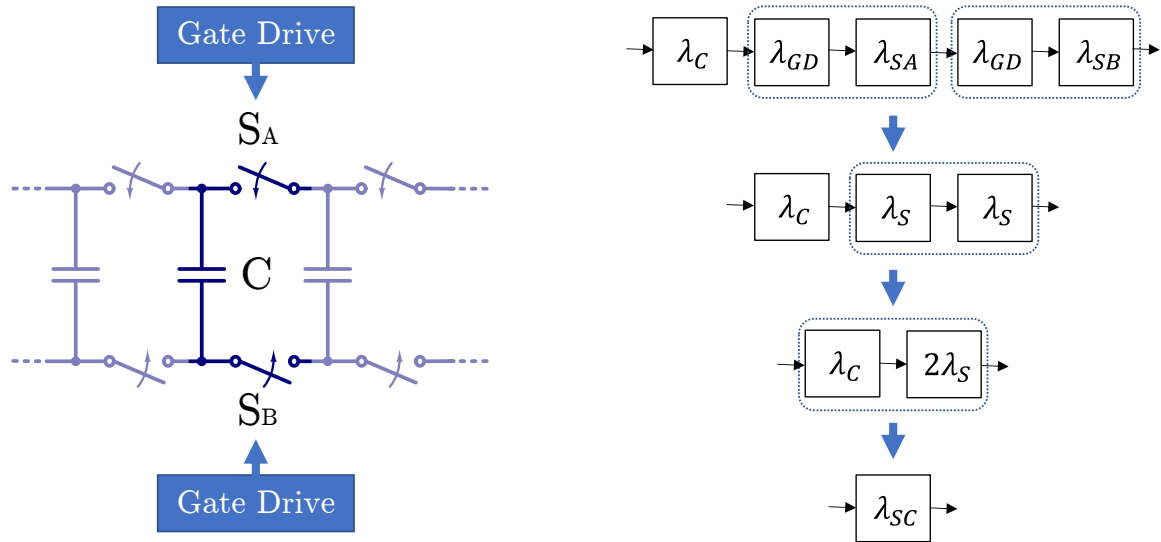
Figure 5.12: Zoomed experimental waveforms showing short-circuit transient for the two mitigation actions for a 6-level FCML.

To provide better insight into the fault transient and mitigating actions, Fig. 5.12 plots a detailed view of the time interval around the fault. Here, the fault begins just before $t = 0$, when the output voltage, v_o , decreases and the inductor current, i_L , rapidly increases in magnitude. In less than $5 \mu\text{s}$, the controller detects an overcurrent fault and initiates the mitigating action. For the response of Fig. 5.12b, this entails commanding the reference to zero and doubling the switching frequency (while also triggering the oscilloscopes capturing these waveforms). The current then decays to zero in less than $50 \mu\text{s}$. During this time, it can be seen that the flying capacitors still experience voltage excursions due to the momentarily increased current, though these excursions diminish as the inductor current resets. Once the inductor current reaches zero, the ride-through phase starts and the current visibly increases once again as from the low modulation applied at the output.

The fault-reversal strategy (time-reversed and inverted carriers, with inverted reference) has also been demonstrated, and results are plotted in Fig. 5.12b. Here, the detection time is unchanged, but the duration of the peak fault current is drastically reduced. Given that this approach is primarily feedforward and the inductor current feedback has a finite sampling rate, the inductor current does exhibit overshoot after resetting owing to the sampling delay in detecting the reset. Nonetheless, the fault is reset nearly instantaneously. Further improvements to these mitigation strategies should recognize that the output inductance of the experimental hardware is nearly ten times smaller than that in the simulation from Section 5.1.3; precisely reversing the induction current action is much more challenging. Furthermore oscillations clearly occur between the load, short-circuit and output filter (as visible

by ringing on v_o) and indicate the response of the output filter should be considered in the complete design review.

5.2 Cell-Level Component Failures and Redundancy



(a) Switching cell components: flying capacitor, two switches and corresponding gate drive circuits.

(b) The series reliability structure of each switching cell simplified as a single failure rate.

Figure 5.13: Switching cell reliability structure for analyzing cell-level failures and redundancy.

Prior work in [228] has shown that redundancy through spare switching cells can improve system-level reliability more rapidly than module redundancy as cell (and component count) increases. Some multilevel topologies provide inherent cell-level redundancy; in the case of a faulty component, the affected cell may be bypassed – allowing continued operation of the inverter phase-leg, albeit in a degraded state (e.g, reduced level count, decreased output power, etc.). For example, given a failure in one switch of the cascaded H-bridge inverter submodule, the cell may be bypassed by closing both of the top or bottom switches of the H-bridge. The half-bridge variant of the modular multilevel converter also supports this functionality with the addition of a single extra bypass switch [229]. Other variants of submodule for modular multilevel converters support can add this functionality with further additions to the component count [230].

However, the flying capacitor multilevel converter does not generally have this cell-level redundancy, except under certain faults and with appropriate component voltage and current ratings [231]. Therefore, this section discusses if and how failures of a specific component in a given switching cell can be mitigated. Referring to Fig. 5.13a, the following component failures are considered:

- Switch S_k : short-circuit failure

- Switch S_k : open-circuit failure
- Switch S_k : loss of controllability
- Gate drive: loss of controllability
- Flying capacitor C_k : short-circuit failure
- Flying capacitor S_k : open-circuit failure

Note that a failure in any of these components essentially leads to a failure of the switching cell, as indicated by the series structure in Fig. 5.13b.

5.2.1 Switch S_k : short-circuit failure

Although a short-circuit fault is not guaranteed to manifest as an ideal, low-resistance path, consideration of such a condition provides a useful starting point to analyze both the fault dynamics as well as the steady-state, post-fault condition. The notable result from this fault, where the mitigating action is to close the surviving complementary switch, is the loss of a discrete voltage level. This is apparent for faults in intermediate cells, where switches S_k and \bar{S}_k closed and adjacent flying capacitors C_k and C_{k+1} (or the DC capacitance) are in parallel. In the case of a fault in cell 1 (at the ac outlet), C_1 simply becomes shorted.

For an incipient fault, transient stresses may be avoided through a graceful transition from N to $N - 1$ levels [71]. However, in an acute fault, the stresses associated with the discharge of the adjacent capacitors must be tolerated by the failed and surviving devices for the converter to remain operational. Prior work in [231] presented an energy stress associated with the re-balancing of charge that silicon IGBT dies must tolerate. However, for discrete MOSFET devices, the pulsed current limit is the most appropriate constraint. While desaturation circuits have also been demonstrated to protect individual GaN devices from an overcurrent condition [232], they do not represent a path to mitigate the device once it has failed short. Nonetheless, it may be possible to leverage the same type of circuit to rapidly detect a short circuit failure and implement the dynamic level change described above.

5.2.2 Switch S_k : open-circuit failure

A switch failing as an open-circuit is most likely caused by the mechanical disconnection of the device terminals, or from failure of the gate drive circuitry. In the case of the former, cell redundancy may be achieved through the use of parallel switches; these devices would continue normal operation in parallel with the open circuit switch. Indeed, paralleling devices already provides an appealing strategy to reduce the overall on-resistance of the switching cells or increase the overall power handling capability of the inverter. The surviving device then either must carry the full current (potentially increasing likelihood of subsequent failures due to increased current stress), or the overall inverter power must be decreased.

If the failure is not an ideal open circuit but instead an intermediate state (i.e., a loss of controllability), a redundant switch could still allow the circuit to be reconfigured for operation at $N - 1$ levels. Here, voltages on C_k and C_{k-1} would be equalized and both the redundant switch in S_k and \bar{S}_k would be closed. The converter would then continue to operate in an $N - 1$ state, removing the failed switch from operation. Otherwise, in the absence of a redundant switch or in the case where the gate drive circuitry failed in an off or uncontrollable state, cell functionality cannot be recovered. As shown in [231], charge balance on the capacitors is no longer possible and continued operation would overcharge capacitors while simultaneously failing to deliver a sinusoidal output.

5.2.3 Gate drive failure

If the gate drive circuitry is compromised, there is little recourse in terms of circuit redundancy – especially if all redundant switches share the same gate driver. However, if an independent gate driver is provided per device, it is possible that the redundant driver could provide the same recourse as the open-circuit case above: bypass the affected gate driver and switch. If redundant gate drivers is not an option (as this would be a significant component increase), a similar recourse could be to add a fail-bypass, where the affected switch is forced into a bypass state if the gate driver can no longer provide the appropriate switching signal. This would be similar to the way in which redundant actuators on aircraft are designed to overpower a failed or stuck actuator for a given control surface [174]. The containment action is then identical to the open-circuit switch failure above that addressed loss of controllability.

5.2.4 Flying capacitor C_k : short-circuit failure

In order to survive an acute short-circuit fault of the flying capacitance C_k , all switches S_x and \bar{S}_x for $1 \leq x \leq k$ need to withstand energy from the lower voltage flying capacitors discharging through switch reverse conduction. Furthermore, the switches S_{k+1} and \bar{S}_{k+1} need to withstand the full, pre-fault voltage across C_{k+1} . Post fault operation at $N - k$ levels is then dependent on whether the capacitors and switches $k + 1$ and above can tolerate the new, elevated voltages. Indeed, failure of higher voltage flying capacitors (closer to the dc side) requires switches to withstand greater reverse conduction from lower voltage capacitors and surviving switching devices to be significantly overrated in voltage than a failure in a lower voltage switch; it may be possible to provide partial coverage from appropriate component derating to allow for some level of fault tolerance.

5.2.5 Flying capacitor C_k : open-circuit failure

An open-circuit failure of a flying capacitor C_k requires no additional components to allow for $N - 1$ operation. However surviving capacitors will need to be rebalanced, while all surviving capacitors and switches will need to be rated to withstand the new, elevated

voltages. Unlike the tandem switches adjacent to a bypassed level in [71], either S_k/\bar{S}_k or S_{k+1}/\bar{S}_{k+1} may remain closed to reduce switching losses.

5.3 Parallel Module Redundancy




Figure 5.14: Simple Markov model illustrating three phases (A,B and C) of six paralleled modules ($P = 6$).

One of the major benefits of the array described in Chapter 4 is that, in addition to interleaving opportunities, redundant modules can improve system reliability. To examine this benefit, Markov modeling was used to consider how individual module failures are covered in different redundant configurations. A simple Markov model describing a three-phase inverter array with six paralleled modules in each phase leg ($P = 6$) is shown in Fig. 5.14. When a module fails, the system enters a degraded state. If the modules can be assumed to have independent failure rates, at least from the healthy state, then the failure rate is the product of the module failure, λ_m , and the total number of modules in the system: $3P\lambda_m$ in Fig. 5.14. A degraded system might then be repaired, for example during scheduled maintenance or, in the case of more electric aircraft, as needed – e.g., between flights before the next departure. Additionally, if the degraded state requires surviving modules to carry a greater proportion of the load, as is the case with *active redundancy*, the failure rate of modules operating in the degraded system, λ_d , may be greater than the nominal λ_m . This consideration will be discussed further in Section 5.4. Finally, if F is then the level of redundancy built into the system, i.e. the number of modules failures that can be tolerated, then when $F + 1$ modules have failed, the system has failed.

In systems with a more nuanced level of redundancy, Markov reliability models are a useful analysis tool. Consider the case where the inverter array is paired with a high-power motor in an open-winding configuration [174]. The loss of a module in any of the phase-legs could then be compensated by the remaining modules within that phase-leg, by modifying the current vectors supplied by the other two phase-legs, or a combination of both that minimizes the overall increase in losses. In such a system, the failure rates of each phase are coupled, and the full system must be derived. This type of configuration is shown in Fig. 5.15 for $P = 6$ and $F = 3$, where analysis has already become quite complicated – especially since failure rates are adjusted according to each degraded state ($\lambda_{d1}, \lambda_{d2}, \dots$). Here, the stochastic matrix, Γ , can still be derived by hand as equivalent states have been reduced to the form as shown. However, for the 1 MW-scale targets for more electric aircraft, many more paralleled modules may be required; an array constructed from the modules in Chapter 2 would require

$$P_{\min} = \left\lceil \frac{1 \text{ MW}}{3 \cdot 18.9 \text{ kW}} \right\rceil = 18, \quad (5.15)$$

where P_{\min} is the minimum paralleled modules needed in a phase-leg. As the number of redundant module, F , increases, so does the number of possible degraded system states.

While a closed form expression for Γ may still be possible, and brute force methods present a path forward for these advanced systems, the following study instead focuses on a subset of configurations pertinent for extant development. In particular, a standard machine configuration is assumed: a single set of three phase terminals with the motor in a delta or wye configuration. In this way, power to each phase is only supplied by modules within a single phase-leg. This section also assumes fuses or other trip mechanisms can isolate individual failed modules, and do so with perfect coverage. Finally, module failure rates are assumed to be independent across phases, allowing the system reliability to be inferred from the per-phase reliability. In other words, the failure rate of a phase-leg is first calculated via a simplified Markov reliability model, and the system failure rate (i.e., a failure in any of the three phases) will thus be a factor of three greater.

Four such configurations are considered: *active redundancy* and *cold standby*, both *with* and *without maintenance*. These latter cases incorporating maintenance are a touchstone of the modular design in earlier chapters, and the overarching approach of this work: instead of designing a single, high-power converter through the use of fewer, high-current power modules like those in Chapter 1 – the failure of which may be catastrophic or otherwise hard to service and repair – swappable yet high-performance modules proposed in this work provides an alternative paradigm. Indeed, sufficient reliability for future vehicles might be attained through occasional service, just like replacing a spark plug in a car or a pitot tube in an aircraft.

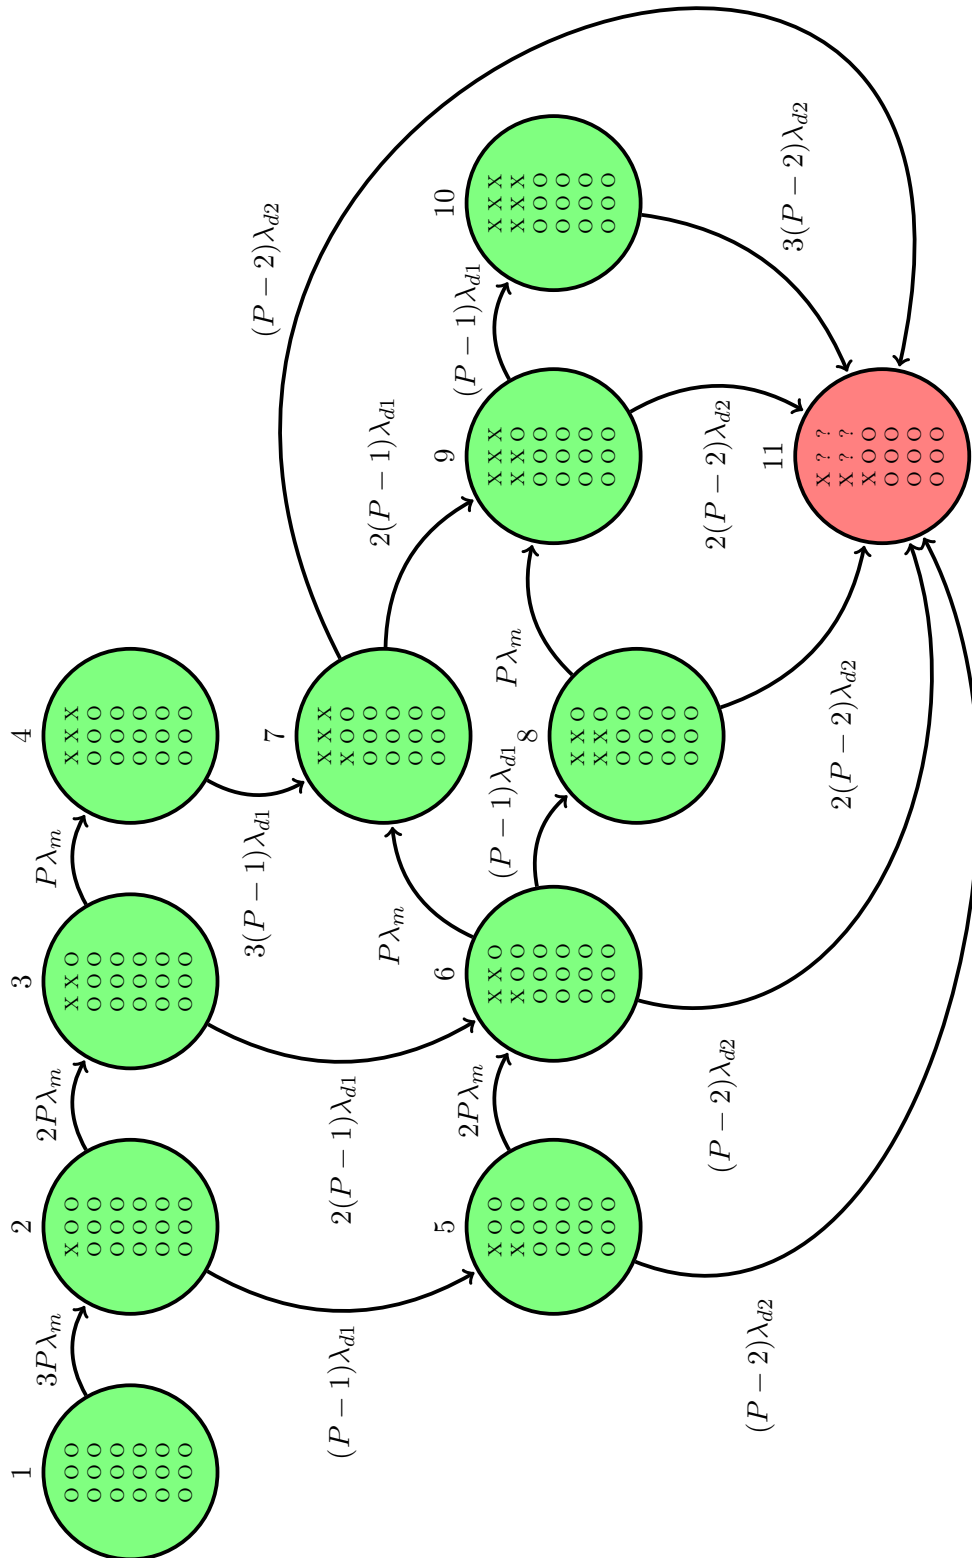


Figure 5.15: Markov reliability model for a three-phase array for $P = 6$ and $F = 2$, where failures in a single phase-leg have complex relationship with the overall system reliability. State 1 is the *healthy* state and state 11 is the *failed* state, with the remaining states indicating some degree of system degradation.

In each configuration, the stochastic matrix Γ takes the form

$$\Gamma = \left[\begin{array}{c|c} A & \begin{matrix} 0 \\ \vdots \\ 0 \end{matrix} \\ \hline 0 \dots \dots 0 & 0 \end{array} \right], \quad (5.16)$$

where the block matrix A is specific to each configuration and λ_{dF} is the failure rate of the degraded state just before the failed state. Each of the following configurations will describe a three-phase array comprised of $P_{\min} + F$ paralleled modules.

5.3.1 Active Redundancy without Maintenance




Figure 5.16: Markov reliability model for a phase-leg implementing active redundancy.

Active redundancy is the configuration where all modules within a phase-leg are energized and sharing the total load. This means that the failure rate of the healthy phase-leg is a product of the module failure rate, λ_m , and the total number of modules ($P_{\min} + F$). While the efficiency might be improved through reducing power (and possibly loss) in each module, this increased failure rate is a caveat and should be considered as part of the system optimization. A model of an irreparable phase-leg with active redundancy is shown in Fig. 5.16. Here, and in all of the following models, 0 represents the absorbing, failed state. This system is described by a diagonal transition matrix, where the only non-zero terms lie on the diagonal and superdiagonal:

$$A = \left[\begin{array}{cccc} -(P_{\min} + F)\lambda_m & (P_{\min} + F)\lambda_m & & \\ & -(P_{\min} + F - 1)\lambda_{d1} & (P_{\min} + F - 1)\lambda_{d1} & \\ & & \ddots & \\ & & & -(P_{\min} + 1)\lambda_{d(F-1)} & (P_{\min} + 1)\lambda_{d(F+1)} \\ & & & & -P_{\min}\lambda_{dF} \end{array} \right] \quad (5.17)$$

5.3.2 Cold Standby without Maintenance

Cold standby is a configuration where redundant modules are only energized once a failure in the phase-leg is detected. Note, for an inverter array comprised of FCML modules, this




Figure 5.17: Markov reliability model for a phase-leg implementing cold standby.

likely requires an additional breaker or switch mechanism to connect each spare as needed – though the exact implementations of this device are not considered here. A model of an irreparable phase-leg with cold standby is shown in Fig. 5.17. Again, there is a single path to the absorbing state that passes through the same number of intermediate, degraded states as the active redundant configuration. However, it is apparent that failure rates remain a constant product of P_{\min} and the module failure rate λ_m . While this is, at first glance, an advantage over the active redundant configuration, the nominal module failure rates may be different in the two configurations as the individual load per module is reduced in the active redundant case. Like the active redundant case, this system is also described by a diagonal transition matrix, where the only non-zero terms lie on the diagonal and superdiagonal; here, $\lambda_{dF} = \lambda_m$ to remain compatible with (5.16):

$$A = \begin{bmatrix} -P_{\min} \lambda_m & P_{\min} \lambda_m & & & \\ & & \ddots & & \\ & & & -P_{\min} \lambda_m & P_{\min} \lambda_m \\ & & & & -P_{\min} \lambda_{dF} \end{bmatrix} \quad (5.18)$$

5.3.3 Active Redundancy with Maintenance




Figure 5.18: Markov reliability model for a repairable phase-leg implementing active redundancy.

In the active redundant configuration proposed above, each module failure would bring the system closer towards a terminal failure. This may be an accurate model if the array were installed on a vehicle in a location that was difficult to access, or if the modules were so tightly integrated that replacement is impossible. But this is a potential caveat to any designs using power modules, where assembly and integration is a one-way process. However, if the inverter array is designed using the methodology in this work, repairs can be carried

out as needed or when otherwise convenient to provide a significant benefit to long-term system reliability.

A repairable, active redundant system is shown in Fig. 5.18. Here, the repair rate μ is used to represent the frequency at which the system can be repaired – for instance, every 12,000 miles (converted into hours) or at an annual tune up. This could also represent the mean flight length if, in the case of an aircraft, service is performed immediately once the flight has safely landed and passengers have deplaned at the destination airport. Note that the approach most likely to be taken would be to restore the aircraft to the healthy state; there is a transition path from all degraded states to the healthy state. This is also visible in the block matrix A below, which otherwise maintains much of the structure of (5.17). Additionally, the failed state remains an absorbing state; this is retained to represent the catastrophic failures (e.g., loss of vehicle) that result in irreparable system damage.

$$A = \begin{bmatrix} -(P_{\min} + F)\lambda_m & (P_{\min} + F)\lambda_m & & & & \\ \mu & -(P_{\min} + F - 1)\lambda_{d1} - \mu & (P_{\min} + F - 1)\lambda_{d1} & & & \\ \vdots & & \ddots & & & \\ \vdots & & & -(P_{\min} + 1)\lambda_{d(F-1)} - \mu & (P_{\min} + 1)\lambda_{d(F-1)} & \\ \mu & & & & -P_{\min}\lambda_{dF} - \mu & \end{bmatrix} \quad (5.19)$$

5.3.4 Cold standby with Maintenance





Figure 5.19: Markov reliability model for a repairable phase-leg implementing cold standby.

The repair structure introduced above can also be applied to the cold standby configuration, as illustrated in Fig. 5.19. The corresponding definition for A is likewise similar to (5.18), and $\lambda_{dF} = \lambda_m$ similarly asserted.

$$A = \begin{bmatrix} -P_{\min}\lambda_m & P_{\min}\lambda_m & & & & \\ \mu & -P_{\min}\lambda_m - \mu & P_{\min}\lambda_m & & & \\ \vdots & & \ddots & & & \\ \vdots & & & -P_{\min}\lambda_m - \mu & P_{\min}\lambda_m & \\ \mu & & & & -P_{\min}\lambda_{dF} - \mu & \end{bmatrix} \quad (5.20)$$

Figure 5.20: State probabilities over time using (5.21) and Γ from (5.17).

5.3.5 State Equations

To find the probability $\pi_k(t)$ that the inverter array is a given state $k \in [F + 1, F, \dots, 1, 0]$ at time t , the Kolmogorov forward equations are first written in matrix form:

$$\frac{\partial}{\partial t} \boldsymbol{\pi}(t) = \Gamma^T \boldsymbol{\pi}(t) \quad \text{where} \quad \boldsymbol{\pi}(t) = [\pi_{F+1}(t), \pi_F(t), \dots, \pi_1(t), \pi_0(t)]^T \quad (5.21)$$

From here, analytical solutions may be derived, or an ordinary differential may be used. For instance, `ode45` in MATLAB was used to generate Fig. 5.20. Here, $\pi_0(t)$ is the probability that the array has failed at time t . As the plots show, the state likelihood tends to progress through the degraded states over time, though a rapid progression to the failed state is possible – however unlikely. Thus it is apparent that increasing redundant modules from $F = 3$ in Fig. 5.20a to $F = 6$ in Fig. 5.20b reduces the likelihood of an early system failure.

Addition of further redundant states will asymptotically decrease the overall system failure rate, and it is useful to characterize the trade-off between the number of redundant modules and the total number of modules. While the plots in Fig. 5.20 are useful for developing an intuition of how the system reliability evolves over time, deriving the mean time to failure of the system (MTTF_S) as these parameters are varied provides greater insight into design trade-offs. The mean time to failure of the system is defined as:

$$\text{MTTF}_S = \int_0^\infty 1 - \pi_0 = \int_0^\infty \sum_{k=1}^{F+1} \pi_k \quad (5.22)$$

Although numerical integration of (5.21) will provide this asymptotic solution, the steady state value can also be determined by applying the final value theorem to the reduced system (less the absorbing states). The exact methodology is described in detail in Section 11.5 of [177], while the approach is only summarized here as follows:

$$A^T \boldsymbol{\pi}^T(t) = \frac{\partial}{\partial t} \boldsymbol{\pi}(t) \quad \text{Remove absorbing states from Kolmogorov equations. (5.23)}$$

$$A^T \boldsymbol{\Pi}^T(s) = s \boldsymbol{\Pi}^T(s) - \boldsymbol{\pi}(0) \quad \text{Take the laplace transform. (5.24)}$$

$$A^T \boldsymbol{\Pi}^T(0) = 0 \boldsymbol{\Pi}^T(0) - \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad \text{Apply the final value theorem } (s = 0). \quad (5.25)$$

$$\boldsymbol{\pi}^T(0) = A^{-T} \begin{bmatrix} -1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad \text{Invert square matrix } A \text{ and solve for } \boldsymbol{\Pi}^T(0). \quad (5.26)$$

These steps can be concisely written for the three-phase array under each configuration as:

$$\text{MTTF}_S = \frac{1}{3} \text{sum}\{A^{-T}[-1, 0, \dots, 0]^T\}, \quad (5.27)$$

where A is defined in Section 5.3.1 through Section 5.3.4. Recall that the division by three in this expression arises by treating failures in each phase-leg as an independent process and considering that the array reliability has a series structure consisting of three such processes. Note also that the -1 in (5.27) is due to the fact that all systems will start in the healthy state, i.e. $\pi_{F+1}(t = 0) = 1$; the negative sign is due to the application of the Laplace transform and final value theorem.

5.3.6 Power Density trade-off

The addition of any extra modules to improve system reliability will impose a power density penalty since each addition adds more mass and volume. A corresponding power density penalty, or PDP , can be defined based on the ratio of the minimum number of modules necessary to deliver the required power to the total number of modules in the system:

$$PDP = \frac{P_{\min}}{P} = \frac{P_{\min}}{P_{\min} + F} \quad (5.28)$$

While the addition of other hardware needed, e.g., for cold standby, might introduce further unaccounted for mass and volume into the system, this ratio provides a sufficient first-order scaling law for comparing trade-offs between power density and reliability.

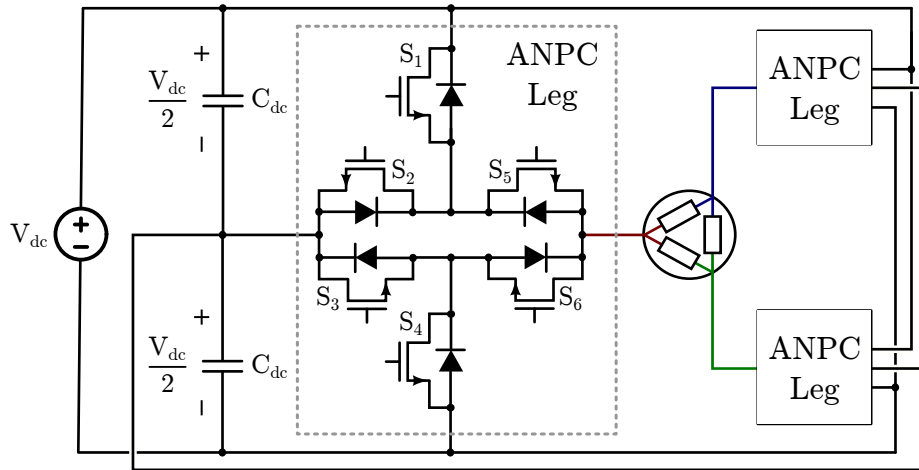


Figure 5.21: A phase-leg of the three-phase, 3-level ANPC for illustrating switch count per module.

5.3.7 Case Study: 1 MW Array for Electric Aircraft

To demonstrate the implications of this section, the trade-off between power density and system reliability is examined for a 1 MW inverter array with application in electric aircraft propulsion. Given the hypothesis that multiple, paralleled modules will be able to address power density and reliability simultaneously, this case provides an opportunity to examine how the module of Chapter 2 with a measured power density might scale to meet the needs; no additional assumptions or scaling laws are required, as would be the case for a lower-power application.

As discussed previously, arriving at reference individual failure rates is difficult because they depend tremendously on the specific design and indeed manufacturing process, quality control and so on. There is, however, reasonable reporting on engine shutdown rates, as well as other system architecture studies that try to peg the range of an inverters failure. Instead of arbitrarily choosing one of these numbers, or others out of a handbook, inverter failures ($\lambda = 8.6 \times 10^{-5}$) for high-power systems provide a reasonable range to assign for a full inverter system failure [233].

Additionally, instead of applying these failure rates directly to the array designs, this section pursues a comparative analysis. Here, the failure rate of 8.6×10^{-5} is applied to the 1 MW inverter design from [50], a three-level ANPC topology. A schematic for a single phase-leg for this converter is shown in Fig. 5.21, where it is seen to have six total switching devices. For simplicity of the comparison in this subsection, methods of failures involving the filters or dc bus capacitance are considered to be beyond the scope of this study. This way, the switch failure rate of the three-phase ANPC, λ_{sw} , can be approximated as

$$\lambda_{sw} \approx \frac{8.6 \times 10^{-5}}{3 \times 6}. \quad (5.29)$$

As discussed in Section 5.2, the dominant building block of the modules proposed in this

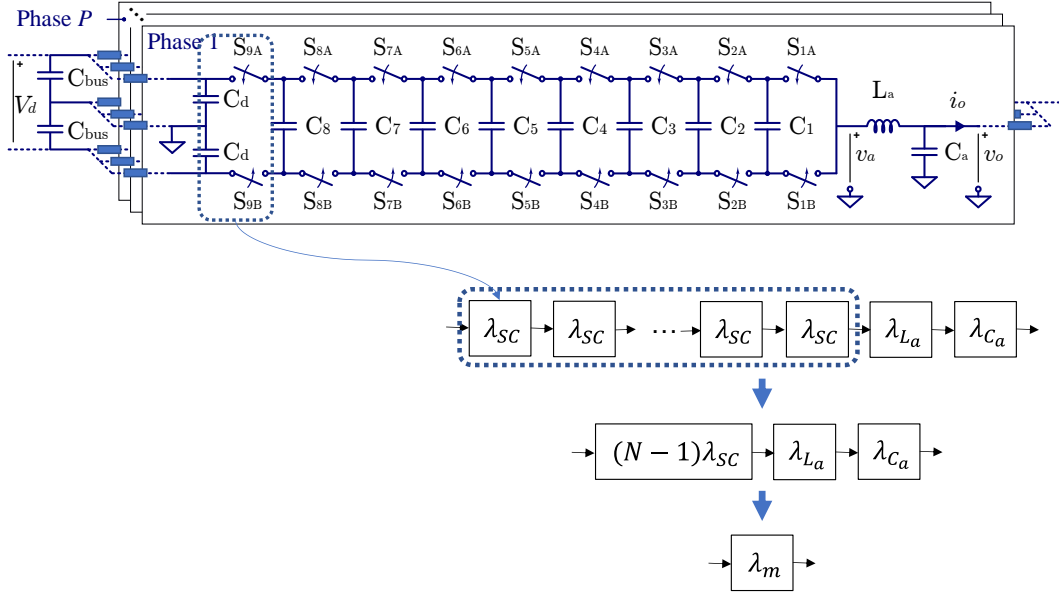


Figure 5.22: Module series reliability structure, dominated by the switching cell failure rate.

work are the switching cells. If design of these switching cells can become as robust as those used for the switches in the ANPC modules, a constraint on the switching cell failure rate, λ_{sc} , may then be asserted: $\lambda_{sc} = \lambda_{sw}$ for the purposes of this study. The switching cell is likewise the dominant contribution to the module failure rate, as illustrated in Fig. 5.22. In this way, a comparative failure rate for the modules, λ_m , can be defined as

$$\lambda_m = (N - 1)\lambda_{sc} = \frac{(N - 1) 8.6 \times 10^{-5}}{18}. \quad (5.30)$$

The results of this study are plotted in Fig. 5.23. To scale the design of Chapter 2 to 1MW, a minimum of 18 paralleled inverters is required per phase-leg – as shown derived in (5.15) – and up to twice that number was considered for phase-leg redundancy. The power density penalty, PDP , is plotted in blue against the left axis of Fig. 5.23. Here it can be seen to geometrically decay, and naturally falls to 0.5 for the case considering $2P_{min}$ redundancy. Model results for the $MTTF_s$ of each of the four cases from Section 5.3.1 through Section 5.3.4 are also plotted against the right axis, normalized to the three level ANPC benchmark; a higher value is better, where a value of 1 means the configuration is on par with the benchmark.

For the irreparable cases, the additional modules do improve reliability. However, there are diminishing returns for the active redundant case, as the increased failure rate from more modules counteracts gains due to extra modules. The cold standby case performs better, and will continue to trend linearly upward; however, the extra modules needed would make the system impractical from a power density point of view.




Figure 5.23: Reliability and power density of a 1 MW array with various levels of redundancy.

The repairable systems provide a much different outlook. Here, the repair rate was assumed to be every aircraft “A-Check.” Schedules vary with different airlines and manufacturers, but these regular maintenance checks occur roughly every 500 flight hours [234]. In both cases, only a moderate sacrifice in power density – about 80% of the nominal design – would be required to meet comparable reliability targets as the ANPC using this standard. However, referring back to figure Fig. 2.23a, for the current design, there is ample overhead in the power density of the proposed design to accommodate this trade-off, as well as penalties from other ancillary hardware necessary for this redundancy.

5.3.8 Aircraft Availability

While the examples above demonstrate how redundancy can drastically improve system reliability against catastrophic failures by incorporating a fixed repair rate, it is unlikely that airline operators will continue to fly an aircraft with a known failure between scheduled maintenance. It is more likely that the airline would service the aircraft as soon as the failure occurs. Indeed, a major benefit of the proposed modular design is that individual inverter modules might be easily swapped as needed. As such, aircraft *availability* may be the better metric. If a repair or replacement is made as soon as possible once a failure is detected, leading to a mean down time (MDT) between normal operation, availability may be approximated as:

$$\text{Availability} \approx \frac{\text{MTTF}_S}{\text{MTTF}_S + \text{MDT}} = \frac{1/(P\lambda_m)}{1/(P\lambda_m) + 1/\mu} \quad (5.31)$$

For example, considering regional flights of more electric aircraft, $\mu \approx 1/5 \text{ hr}^{-1}$ if flight durations of up to two hours comprise the bulk of segments, and a few additional hours are considered for service once the flight has landed. Using λ_m as defined in Section 5.3.7, and

$P = 22$ corresponding to the break-even MTTF from Fig. 5.23, a comparison may be made for designs of different level count. For instance, while a 10-level FCML using 200 V devices was considered in this work, a 5-level converter using 650 V devices may also be an appealing design candidate. Although several metrics will likely differ between the designs, the change in availability could be a significant consideration:

$$\text{Availability}_{10\text{-level}} = \frac{1057}{1057 + 5} = 99.5\% < \text{Availability}_{5\text{-level}} = \frac{1057^{\frac{9}{4}}}{1057^{\frac{9}{4}} + 5} = 99.8\% \quad (5.32)$$

5.4 Chapter Remarks

The chapter began by presenting methods for robust and rapid ac-side short-circuit fault detection and mitigation for the FCML inverter. An experimental platform was used to verify this fault mitigation on both even- and odd-level converters operating at high switching frequency. Additionally, models from [200] were applied to show that while some converter designs can provide adequate headroom to implement short-circuit protection, others that solely target extreme power density may push critical detection and mitigation times beyond the coverage of realizable implementations. However, this may be improved with modifications to design parameters, e.g., increasing the flying capacitance or output inductance, using higher voltage transistors or decreasing the rated current. Future work on practical applications of the FCML should consider these concerns as an integral part of the design process to balance performance with robustness.

Similarly, the second section also indicated that derating switching devices and capacitors can allow for coverage of certain component level failures. While additional detection and mitigation systems beyond what is presented in this work would be required, some aspects of cell-level redundancy – especially in the case of paralleled switches – could be implemented in future work.

Although the introduction to this chapter also made a point to emphasize that Arrhenius relationships cannot be the sole approach for reliability-based design, there is often a strong correlation between loss (and therefore temperature) in many power electronic systems [187]. For cases where a correlation between the operating point and resulting change in failure rates is known or informed by physics of failure, π factors or coefficients on the nominal failure rate may be used [235]:

$$\lambda = \lambda_0 \pi_U \pi_I, \pi_T \quad (5.33)$$

where λ_0 is the nominal failure rate, π_U is a voltage-dependent coefficient, π_I scales based on a current-dependence, and π_T represents the affect of operating temperature.

The Arrhenius model is typically used to describe π_T :

$$\pi_T = e^{\left(\frac{E_a}{k_B} \left(\frac{1}{T_0} - \frac{1}{T_d}\right)\right)} \quad (5.34)$$




Figure 5.24: Temperature-dependent failure rate coefficient, π_T , for the high-k ceramic capacitors considered in this work.

where k_B is the Boltzmann constant, T_0 is the comparison temperature, T_d is the operating temperature and E_a is an “activation energy” describing the dependence. For instance, for the ceramic capacitors used in this work, E_a can be as low as 1.19 eV [236] or as high as 1.96 [185], though it generally lies between 1-2 eV [237]. The corresponding π_T is plotted in Fig. 5.24 for a T_0 of 85°C. Here it can be seen that operating at a higher temperature (e.g., 100°C) could lead to a failure rate an order of magnitude higher than nominal. This has implications in the λ_d described above in the active-redundant case: if the surviving modules in a given phase need to increase output power (and therefore losses), but maintain the same cooling performance, reliability – at least of the capacitors – will be further degraded.

While the physics of failure for gallium nitride devices is still somewhat nascent, there is indication that a similar activation energy associated with off-state voltage may exist [196]. If such a correlation is significant, the associated π_T could also be a relevant metric at the design stage: choosing higher level counts would reduce the nominal off-state voltages of each device, while maintaining level count and choosing higher voltage rating devices might similarly provide a reliability boost.

As it were, the notion of reliability-driven design for these high performance applications is still evolving. For instance, while [238] developed a comprehensive reliability model of a more electric aircraft power system, the data was largely based off of MIL-HDBK-217F. Although this has been a standard reference of failure rate ranges for many years, a recent movement in the research community has denounced the accuracy of such methods and generally calls for more nuanced approaches [239, 240]. Indeed, the goal of the last section of this chapter was to do just that – though of course, some simplifications did need to be made to keep the model tractable for this work. For instance, the exact failure rate used is based on an assumption of electronics performance and notably carries with it a high level of

uncertainty. However, considering shutdown rates of extant engine technologies are on the order of 10^{-6} [17], the framework provides a reasonable starting point for these estimates.

As vehicle concepts continue to evolve, system level design choices will also provide a better understanding of the reliability needs of the drivetrain. For instance, designs with distributed electric propulsion in electric aircraft spread the overall vehicle thrust across many propulsors – compared to the typical two for extant commercial aircraft. If each propulsor has a separate inverter system, the reliability target for that inverter becomes a function of the aircraft configuration, e.g., how many propulsors must remain operational and the extent at which flight surfaces can be trimmed to compensate for those that are inoperable [241, 242].

Finally, depending on the exact motor implementation, both topology and array construction may be revisited to enhance reliability. For instance, multiple coil sets [227] or phases [157, 243] can provide an additional degree of freedom. If the coil sets of surviving phases can withstand temporarily greater output power, then what was lost from a module failure could be compensated by increasing output power across other windings to reduce the effects of the loss. Better yet, machines with independent or open-windings can provide a further degree of freedom in the event of a module or even phase failure [174, 244]. For example, an open-winding machine could be driven by an FCML on one side of the coil and a low-frequency bridge (i.e., an unfolder) on the other side. In this way, the same number of levels could be generated – but with fewer components, leading to a potential decrease in failure rates. Furthermore, the loss of a phase in this configuration could be compensated with compensation applied to the remaining, independent windings. As designs for future electric powertrains are still in early stages, consideration of nuances like these before integration activities take place can extend concepts presented in this chapter to produce a high-performance and high-reliability design.

Chapter 6

Future Work & Conclusion

This brief chapter concludes the thesis by reviewing some possible next steps and remaining open questions for further work. Although some of the following will be more mundane tasks or simple experiments only requiring more time and manpower, other ideas will need a bit more thought and analysis. It is the intent that this brief list may inform or motivate future work to improve the solutions in this work to progress beyond the results presented in the previous chapters.

6.1 Extended Component Characterization

Although noted in the concluding remarks of Chapter 2, the value of experimental component characterization in informing optimal converter design is worth further emphasis. This work was fortunate enough to benefit from prior work on characterization of several promising types of capacitors [78] and several gallium nitride transistors [77]. However, as many advances in converter performance are due to newly released components, and corresponding models that describe their behavior, an investment in a sustaining characterization effort can keep the loss models and design constraints at the forefront of the technology. Additionally, test fixtures for characterization are valued research projects within themselves, [77, 78, 95, 245]; any further improvements to repeatability and incorporation of new part numbers and operating regimes are likely to be as well received as the measurements themselves. Finally, this performance characterization effort can be coupled with a reliability characterization effort to gain insight on the long-term behavior of these new devices. Long-duration, burn-in tests across varied environmental conditions can be automated once the test stand is designed, and can provide valuable data and insights when designing for reliability.

6.2 Advanced Gate-Drive Techniques

As discussed in Chapter 2, this work employed a simple yet somewhat legacy method to supply isolated voltages and gate drive signals to each of the floating switches. This circuitry,




Figure 6.1: Highlight of the large area occupied by the gate drive circuitry in this work.

originally shown in Fig. 2.2, is highlighted in Fig. 6.1. From the figure, it is apparent that the gate drive occupies a large footprint. Fortunately, by using any of the techniques presented in [83], both the component count, cost, area and mass can be approximately halved by eliminating the isolated power supply chip in each switching cell of this work. This can lead to an improvement to improving volumetric power density by 10-20%. However, it is worth noting that the series string of bootstrap connections in these designs may have reliability implications: a failure of any cell in this chain will likely lead to module failure, and thus merits further study in critical operating regimes.

Implementation of newly developed gate drivers for wide-bandgap devices will also allow for enhanced performance across the operating space, as well as a new means for estimating converter health [246]. For instance, the programmable drive strength of these new drivers could be coupled with adaptive dead-time to tailor switching to varying operating conditions, even on the scale of a fundamental period. At light load, or anytime during the line cycle around the zero-crossing, current flowing through the converter is low and may not fully soft-charge in the way described in Chapter 2. In these low-current regimes, overlap losses would already be negligible, so a lower drive strength and increased dead-time could be used to slow turn-on and minimize risk of shoot-through and detrimental overshoot. Conversely, when heavily loaded or at the peaks of the line cycle when at least half of the converter is soft-switching (assuming an inductive load), drive strength could be increased and dead-times reduced to switch devices as fast as possible. Indeed, an initial implementation of adaptive dead-time over a line cycle alone showed an improvement of 0.05% for the converter of this work during low power tests. However, this modest improvement will scale with overlap losses – which Fig. ?? showed are significant at high power.

At the same time, these new drivers will be among the first allowing for integrated, bidirectional signal transmission; drain-source voltage measurements and fault states in the floating voltage domain can be read back to the controller. As the off-state drain-source voltage is defined by the adjacent flying capacitor voltages, the sensing in these new drivers

can enable robust estimation of flying capacitor balancing without the need for additional digitization and isolated sense circuitry. Furthermore, much like the fault detection methods in Section 5.1 could rapidly detect a fault at the ac output, these drivers could conclusively detect faults at the switching cell and alert the controller mitigate the failure instantaneously, without delays associated with the detection methods discussed.

6.3 Validation of Sense and Measurement Schemes

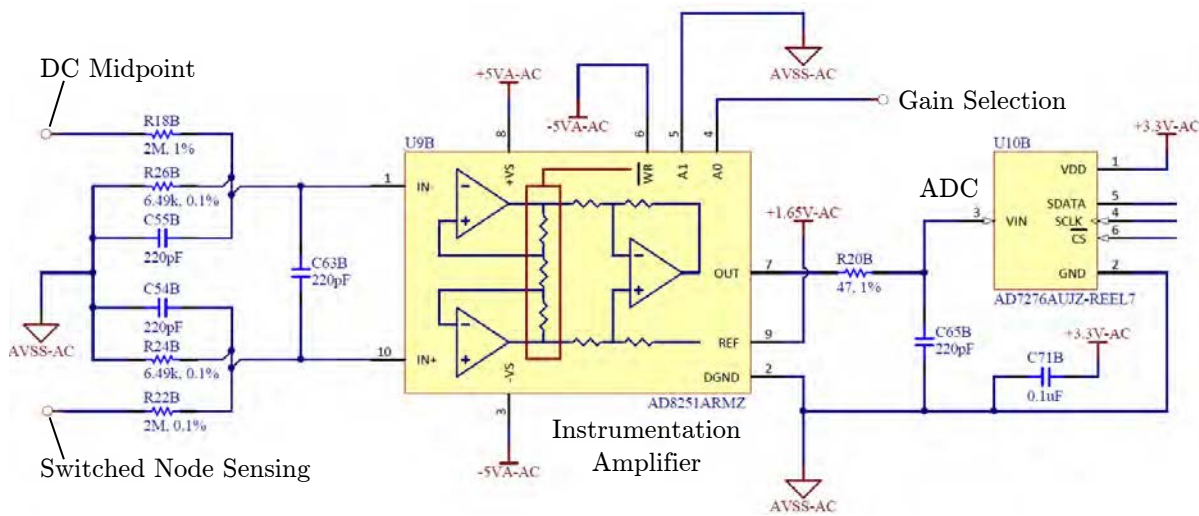


Figure 6.2: Instrumentation amplifier used to capture the high-frequency and high-dynamic-range waveform at the switched-node of the FCML. Gain can be toggled between 1 and 2 to enhance resolution near zero-crossings or at low modulation.

The high effective output frequency provides many advantages, but also introduces a significant measurement challenge at the output of the converter. This is further complicated at high dc bus voltages and output powers, where the output voltage and current both extend across a significant dynamic range. Measuring the switched-node is particularly useful, as it can be used to determine the state of the flying capacitors; since this voltage represents the systematic series connection of the various flying and dc bus capacitors, the respective voltages can be determined by measuring the output voltage and switch state over a series of switching actions [247, 248]. Yet, due to the interleaving effect of the PSPWM (see Section 1.3), there is no minimum pulse-width constraint on this waveform (a zero pulse-width occurs when $D/(N - 1)$ is an integer), while the maximum pulse-width is at most half of the effective switching period. For a 1 MHz effective output frequency, this means the pulse-width available for a measurement is at most only 500 ns. Therefore, more than having a high bandwidth, the analog front end for this measurement must have a very fast settling time such that the output of the sensing circuit is representative of the true switched

value by the time it is sampled. In the proposed voltage-sense circuit of Fig. 6.2, the settling time was simulated to be on the order of 300 ns. From initial experimentation with sensing this voltage, one approach to ensure a sample is valid was to first start a 300 ns timer after any gate transition. If another switch transition occurred before the time limit, the sampling process was reset. However, if the timer was able to complete, an analog-to-digital conversion is started. Finally, if the conversion finished before the next switch transition, the measurement and corresponding switch state are recorded for use in the estimation algorithm. Otherwise, the sample is discarded and the sampling process restarts.

For a 2 MSPS sample rate, this means two samples (one each at the high and low value of the switched node) are possible per period for effective duty ratios between 0.3 and 0.7. Future work may leverage this framework, but should closely study and empirically verify the implications of this irregular sampling strategy. Likewise, the noise-immunity of the circuit in Fig. 6.2 should be verified and possibly improved (e.g., with common-mode filtering or a different instrumentation amplifier implementation) when operating at high-voltage.

In addition to the switched-node, the output inductor current ripple is also at the effective switching frequency. While high-current sensors are available in bandwidths up to 1 MHz (e.g., the ACS733 current sensor discussed in Section 5.1), even higher bandwidth is needed if current-mode control is to be implemented at the switching frequency. Although [173] demonstrated closed-loop, three-phase array operation with current bandwidths below 100 kHz, high performance machines and control techniques may require output bandwidths in excess of this figure. If such high control bandwidth is indeed needed, a new design from [249] has demonstrated measurement capability of currents from dc to 35 MHz with high common-mode rejection. As this sensor required custom magnetics, such a measurement system should be considered in tandem with the magnetics design of the output filter to make the most efficient use of space.

6.4 Online Health-Monitoring

Switched-observer-based fault detection [250–252] was originally intended to provide online health-monitoring as part of this work; measuring incipient faults due to degraded component health is a hallmark of the switched-observer approach and would suit the applications targeted in this work. However, the issues with sampling described above extended to this aspect of the project as well. In particular, the switched state-space model used in the prior works does not extend well to the modulated FCML. Specifically, as the duty cycle is modulated across the fundamental line waveform, operating points near the zero minimum-pulse-width (when $D/(N - 1)$ is an integer) mean the converter spends an infinitesimal duration in the respective states – purely as a consequence of the PSPWM. If the observer (and relevant state measurements) are sampled on the time-scale of the switch carrier counter (i.e., every 5 ns for the 200 MHz system clock used in this work), the state error tracks correctly. However, if either the measurements or the observer are “undersampled,” i.e., to the degree at which the physical system may be realistically sampled (e.g., at 1 MHz – or




Figure 6.3: Current amplifier for the LMT01 temperature sensors on the modules in this work.

200 times slower than the physical system is stepped), the model error rapidly diverges – even in the absence of a fault. Further exploration of the switched-observer sampling may lead to a solution, but it is likely that other methods surveyed in Section 5.1 or [253] can provide more appropriate alternatives for the FCML.

Temperature estimation is also a useful means to evaluate component health [254]. As such, 16 discrete LMT01 temperature sensors were also included in the bill of materials for each module in Chapter 2. These two-wire sensors output a pulse-count modulated current between $35\mu\text{A}$ and $125\mu\text{A}$ depending on the temperature they sense. When one terminal is supplied by a voltage source, the other terminal can drive a voltage across a high-impedance sense resistor to communicate this temperature measurement to a microcontroller. Unfortunately, these current levels are particularly low, and the signal is highly susceptible to noise. Indeed, all measurement functionality was lost when operating the modules above 200 V. Therefore, the single transistor amplifier in Fig. 6.3 was designed to provide a current gain of 36 – resulting in a signal between 0.2 mA to 4.5 mA. Further testing on prototype hardware is required, but simulations show this to be a promising solution for incorporating temperature telemetry to the local controller feedback. Furthermore, if these temperature sensors can be made dependable, less time will be required to fixture converter thermocouples prior high power experiments.

6.5 Folded & Three-Dimensional Module Packaging

In addition to the vertical switching cell layout used in this thesis, and the lateral layouts it was contrasted to in Chapter 2, more elaborate switching cell designs have been recently proposed [61, 84]. These designs again leverage internal layers and blind or buried vias to create advantageous commutation loop routes that continue to reduce parasitic inductances. Additionally, they use lower-valued local bypass capacitors for the primary commutation loop, while electrically-thin routing is used to locate bulk flying capacitance in a more convenient location adjacent to the switching-cell, illustrated in Fig. 6.4a. Direct implementation of these new concepts can help clear out space around each transistor that is currently reserved for the large flying capacitors. This can simplify and improve thermal management; if the large capacitors can be moved away from the transistors, the post on the bottom of the heat sinks in Chapter 3 could be eliminated – allowing for a mechanically-robust and

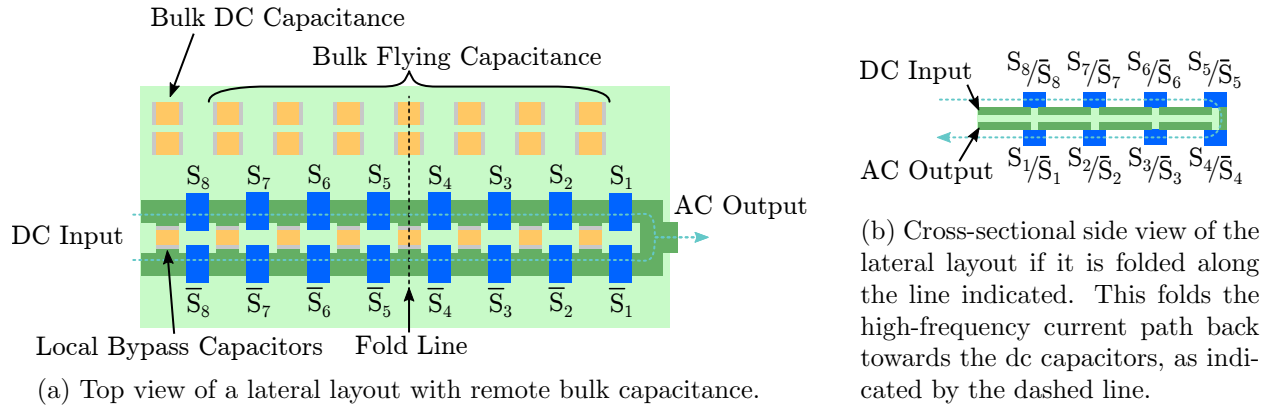


Figure 6.4: A conceptual illustration of how layouts that implement remote flying capacitance can improve the routing of switching currents in a 9-level FCML.

higher-performance thermal interface.

This new concept, where bulk capacitance can be located a relatively remote distance (several centimeters) from the primary switching cell commutation loop, permits other creative ideas. For instance, other high-frequency return paths could be optimally routed; in addition to the parasitic inductance in each commutation loop, there is also a parasitic inductance that extends the entire length from the dc terminals to the ac terminals, illustrated in Fig. 6.4a. While energy stored in this inductance contributes to the drain-source overshoot to a minor degree, it also represents the path in which switching currents filtered through the output capacitance must return; this loop length is over 30 cm in the current design! However, a folded design, where both the ac and dc terminals enter and exit at the same side of the converter, can realize the same benefits of anti-parallel current routing demonstrated in the electrically-thin switching cells. This is illustrated in the 9-level converter of Fig. 6.4b. Current flows from the dc terminals on the top side of the board down through devices from left to right, through the board, then right to left – ending at the ac terminals, also on the left of the converter. This path cancels the flux along the long string of series connected devices, while also creating a convenient connection of the output filter return, providing a very low impedance path for switching currents.

Another notable aspect of this folded configuration is that a high level-count can be achieved at half of the original converter length. This is more of a practical consideration than one dictated by performance; with a physically long converter PCB like the one in Chapter 2, it is easy to inadvertently flex the board while handling, reworking or installing. Unfortunately, this increases the likelihood of cracking the ceramic capacitors or flip-chip transistors which can lead to module failures. However, if the board aspect ratio and area can be reduced through an alternative layout, these concerns can be mitigated.

6.6 Improved Cooling and Integration

If liquid cooling does indeed represent the next iteration of thermal management for the inverter modules presented in this work, the work of [135] can perhaps inspire how the full array can be integrated into a drive system yet maintain the modularity and serviceability emphasized earlier in this work. More elaborate cooling systems are also possible – where Fig. 6.5 represents a concept in which a liquid cooling scheme for the inverters of this work is combined with the motor thermal management system. As rendered, this configuration might pose assembly and maintenance challenges – as the electrical and fluid ports are on both sides of the converter. However, if the folded geometry above is instead used, a seamless integration of both coolant and electrical connections could be devised.




Figure 6.5: Conceptual mock-up of how the converter array of this work might be scaled and integrated with the motor of [227]. In the rendering, a liquid cooling system and cold-plate is used to cool the module versus the individual, air-cooled heat sinks presented in Chapter 3.

6.7 Coupled Inductors for Interleaved Modules

Although Chapter 4 showed that mismatch in current sharing between interleaved inverter legs is expected to be minimal, a strategy that strictly enforces this policy will ensure a better loss distribution. While closed-loop control could actively regulate individual converter setpoints, integrating the magnetics of the output inductors across a single phase-leg could accomplish this passively. For instance, in [255], careful design of a coupled inductor at the output of an interleaved buck converter (topologically comparable to a 2-level inverter) enforced equal current sharing between the interleaved phases. Additionally, in [256] a similar outcome was achieved with interleaved voltage-source inverters. Besides improving current sharing, coupled inductors provide an opportunity to combine differential- and common-mode filters into a single magnetic core, improving core utilization and reducing overall mass. Indeed, [257] showed how a coupled inductor could be used to combine differential- and common-mode filter for interleaved 3-level ANPC inverters.

However, much like the cascaded bootstrap, it is worth considering what implications introducing this coupling among independent modules has on reliability. Specifically, if one module fails short- or open-circuit, it would be preferable that this failure not propagate




Figure 6.6: A root locus plot for the input-to-state transfer function for one of the flying capacitors of the 10-level converter of this work, as the switching frequency is swept from 20 kHz to 115 kHz. The arrows indicate the direction of the pole and zero movement as frequency increases, while the colors indicate pole-zero pairs.

to other healthy modules. Furthermore, while better current sharing is expected to lead to better current balancing, it is unclear how a coupled inductor between interleaved phases will affect passive voltage balancing of the flying capacitors.

6.8 Modeling of Converter Dynamics

Modeling FCML converters remains a formidable challenge. While Section 2.3 discussed an approach for evaluating behavior during converter startup, [116, 165, 204] all presented an insightful model of FCML dynamics using a sort of averaged model built from the Fourier synthesis of the switching functions for each cell. The states of this model are the flying capacitor voltages and, although the primary application of this approach was to model the input to state transfer function, it can likely provide other insights as well. For instance, root locus plots were generated for each of the input-to-state transfer functions, with Fig. 6.6 showing the movement of pole-zero pairs as switching frequency is increased from 20 kHz (for startup - see Chapter 2) to the nominal switching frequency of 115 kHz. This approach be used to indicate at which operating points potential instability may arise (as right half-plane

zeros approach the origin as frequency increases, indicated by the arrows). Additionally, the natural frequencies and damping associated with each pole-zero pair accurately describes the ripple present in flying capacitor voltage waveforms during simulation and operation.

6.9 Conclusion

As much of this chapter was spent highlighting the work yet to be done – and remarks were made at the conclusion of each chapter – it is worth recounting the extent of work that has been accomplished in this manuscript. At a basic level, this work produced the most power-dense, sine-output inverter module (that did not resort to liquid cooling) of extant architectures. Correspondingly, this work simultaneously achieved the high output power and high efficiencies required for applications in electric vehicle propulsion. This involved analysis and careful design of the thermal management system, as well as the careful test and measurement of both the thermal and electrical performance. Later in the work, a scheme to address scalability was proposed and demonstrated, validating the approach of initially independent module design followed by scaling through paralleling. Finally, significant reliability concerns pertaining to both the application and the chosen topology were addressed. Consideration was given to both prevention of critical faults through design for reliability, as well as active mitigation for protecting a converter when a fault has already occurred. Future applications of the unconventional approach in this work will thus be well suited for solving design challenges in electric vehicle propulsion – as well as any sector requiring high-performance inverter systems.

Bibliography

- [1] Nathan Pallo et al. “A 10-Level Flying Capacitor Multi-Level Dual-Interleaved Power Module for Scalable and Power-Dense Electric Drives”. In: *2020 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2020, pp. 893–898. DOI: 10.1109/APEC39645.2020.9124531.
- [2] Nathan Pallo, Roderick S. Bayliss, and R. C. N. Pilawa-Podgurski. “Multi-Phase Segmented Drive Comprising Arrayed Flying Capacitor Multi-Level Modules”. In: *2021 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2021.
- [3] Daniel Sperling. *Three Revolutions: Steering Automated, Shared, and Electric Vehicles to a Better Future*. Island Press, Mar. 1, 2018. ISBN: 9781610919067. URL: https://www.ebook.de/de/product/31591875/daniel_sperling_three_revolutions.html.
- [4] Chris Paine. *Who Killed the Electric Car?* Documentary. Sony Pictures Classics, 2006.
- [5] *GM Boosts Investment, Grows Electric Portfolio to Lead in EV Race*. Accessed: 07-08-2021. General Motors. Nov. 2020. URL: <https://media.gm.com/media/us/en/gm/home.detail.html/content/Pages/news/us/en/2020/nov/1119-electric-portfolio.html>.
- [6] Kevin Stark. “Lyft Has Lofty Climate Goals Experts Want Concrete Actions”. In: *KQED News* (July 2020). Print: July 3, 2020, Accessed: July 18, 2021. URL: <https://www.kqed.org/science/1966119/lyft-sets-target-for-electric-vehicles-ahead-of-a-clean-air-mandate>.
- [7] International Energy Agency. *Energy Technology Perspectives 2017*. English (US). France: Organisation for Economic Co-operation and Development (OECD), June 2017. ISBN: 978-9264275973. DOI: 10.1787/energy_tech-2017-en.
- [8] Masson-Delmotte, V., P. Zhai, H.-O. Pörtner, D. Roberts, J. Skea, P.R. Shukla, A. Pirani, W. Moufouma-Okia, C. Péan, R. Pidcock, S. Connors, J.B.R. Matthews, Y. Chen, X. Zhou, M.I. Gomis, E. Lonnoy, T. Maycock, M. Tignor, and T. Waterfield (eds.) “IPCC, 2018: Summary for Policymakers”. In: *Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strength-*

- ening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty*. Geneva, Switzerland, 2018.
- [9] International Energy Agency. *Energy Technology Perspectives 2020*. Organisation for Economic Co-operation and Development (OECD), Sept. 2020. DOI: 10.1787/d07136f0-en.
- [10] Benjamin J. Brelje and Joaquim R.R.A. Martins. “Electric, hybrid, and turboelectric fixed-wing aircraft: A review of concepts, models, and design approaches”. In: *Progress in Aerospace Sciences* 104 (Jan. 2019), pp. 1–19. DOI: 10.1016/j.paerosci.2018.06.004.
- [11] *Flightpath 2050 : Europe’s vision for aviation : maintaining global leadership and serving society’s needs*. Luxembourg: European Commission, 2011. ISBN: 9789279197246.
- [12] Cheryl L. Bowman, James L. Felder, and Ty.V Marien. “Turbo- and Hybrid-Electrified Aircraft Propulsion Concepts for Commercial Transport”. In: *2018 AIAA/IEEE Electric Aircraft Technologies Symposium (EATS)*. 2018, pp. 1–8.
- [13] David K. Hall et al. “Boundary Layer Ingestion Propulsion Benefit for Transport Aircraft”. In: *Journal of Propulsion and Power* 33.5 (Sept. 2017), pp. 1118–1129. DOI: 10.2514/1.B36321.
- [14] Luca Menegozzo and Ernesto Benini. “Boundary Layer Ingestion Propulsion: A Review on Numerical Modeling”. In: *Journal of Engineering for Gas Turbines and Power* 142.12 (Dec. 2020). DOI: 10.1115/1.4048174.
- [15] Kevin R. Moore and Andrew Ning. “Distributed Electric Propulsion Effects on Existing Aircraft Through Multidisciplinary Optimization”. In: *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. American Institute of Aeronautics and Astronautics, Jan. 2018. DOI: 10.2514/6.2018-1652.
- [16] Hyun D. Kim, Aaron T. Perry, and Phillip J. Ansell. “A Review of Distributed Electric Propulsion Concepts for Air Vehicle Technology”. In: *2018 AIAA/IEEE Electric Aircraft Technologies Symposium*. American Institute of Aeronautics and Astronautics, July 2018. DOI: 10.2514/6.2018-4998.
- [17] Alan H. Epstein and Steven M. O’Flarity. “Considerations for Reducing Aviation’s CO₂ with Aircraft Electric Propulsion”. In: *Journal of Propulsion and Power* 35.3 (2019), pp. 572–582. DOI: 10.2514/1.B37015. eprint: <https://doi.org/10.2514/1.B37015>. URL: <https://doi.org/10.2514/1.B37015>.
- [18] Ralph Jansen et al. “Overview of NASA Electrified Aircraft Propulsion (EAP) Research for Large Subsonic Transports”. In: *53rd AIAA/SAE/ASEE Joint Propulsion Conference*. American Institute of Aeronautics and Astronautics, July 2017. DOI: 10.2514/6.2017-4701.

- [19] Gerald Brown. “Weights and Efficiencies of Electric Components of a Turboelectric Aircraft Propulsion System”. In: *49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*. American Institute of Aeronautics and Astronautics, Jan. 2011. DOI: 10.2514/6.2011-225.
- [20] Ralph H. Jansen et al. “High Efficiency Megawatt Motor Preliminary Design”. In: *2019 AIAA/IEEE Electric Aircraft Technologies Symposium (EATS)*. 2019, pp. 1–13. DOI: 10.2514/6.2019-4513.
- [21] Ruirui Chen et al. “A Cryogenically Cooled MW Inverter for Electrified Aircraft Propulsion”. In: *AIAA Propulsion and Energy 2020 Forum*. American Institute of Aeronautics and Astronautics, Aug. 2020. DOI: 10.2514/6.2020-3598.
- [22] Christopher B. Barth et al. “Design, Operation, and Loss Characterization of a 1-kW GaN-Based Three-Level Converter at Cryogenic Temperatures”. In: *IEEE Transactions on Power Electronics* 35.11 (2020), pp. 12040–12052. DOI: 10.1109/TPEL.2020.2989310.
- [23] International Energy Agency. *Energy Technology Perspectives 2020: Special Report on Clean Energy Innovation*. Organisation for Economic Co-operation and Development (OECD), July 2020. DOI: 10.1787/ab43a9a5-en.
- [24] Ralph Jansen et al. “Turboelectric Aircraft Drive Key Performance Parameters and Functional Requirements”. In: *51st AIAA/SAE/ASEE Joint Propulsion Conference*. 2015. DOI: 10.2514/6.2015-3890. eprint: <https://arc.aiaa.org/doi/pdf/10.2514/6.2015-3890>. URL: <https://arc.aiaa.org/doi/abs/10.2514/6.2015-3890>.
- [25] Ralph Jansen, Kirsten P. Duffy, and Gerald Brown. “Partially Turboelectric Aircraft Drive Key Performance Parameters”. In: *53rd AIAA/SAE/ASEE Joint Propulsion Conference*. American Institute of Aeronautics and Astronautics, July 2017. DOI: 10.2514/6.2017-4702.
- [26] Kirsten P. Duffy and Ralph Jansen. “Partially Turboelectric and Hybrid Electric Aircraft Drive Key Performance Parameters”. In: *2018 AIAA/IEEE Electric Aircraft Technologies Symposium*. American Institute of Aeronautics and Astronautics, July 2018. DOI: 10.2514/6.2018-5023.
- [27] *Amendment No. 3 to the NASA Research Announcement (NRA) entitled ‘Research Opportunities in Aeronautics - 2015 (ROA-2015)’*. NASA Research Announcement NNH14ZEA001N. NASA, 2015. URL: <https://nspires.nasaprs.com/external/viewrepositorydocument/cmdocumentid=455490/ROA-2015%20NRA%20Amendment%20No.%203.pdf>.
- [28] Ralph Jansen, Cheryl Bowman, and Amy Jankovsky. “Sizing Power Components of an Electrically Driven Tail Cone Thruster and a Range Extender”. In: *16th AIAA Aviation Technology, Integration, and Operations Conference*. American Institute of Aeronautics and Astronautics, June 2016. DOI: 10.2514/6.2016-3766.

- [29] David J. Sadey et al. “A Generalized Power System Architecture Sizing and Analysis Framework”. In: *2018 Joint Propulsion Conference*. American Institute of Aeronautics and Astronautics, July 2018. DOI: 10.2514/6.2018-4616.
- [30] David C. Loder, Andrew Bollman, and Michael J. Armstrong. “Turbo-electric Distributed Aircraft Propulsion: Microgrid Architecture and Evaluation for ECO-150”. In: *2018 IEEE Transportation Electrification Conference and Expo (ITEC)*. 2018, pp. 550–557. DOI: 10.1109/ITEC.2018.8450180.
- [31] Advanced Research Projects Agency - Energy. *Aviation-class Synergistically Cooled Electric-motors with iNtegrated Drives*. Funding Opportunity No. DE-FOA-0002238. United States Department of Energy, 2019. URL: <https://arpa-e-foa.energy.gov/FileContent.aspx?FileID=6506584b-87b2-4d66-8c6c-2fefcfce754f>.
- [32] U.S. DRIVE Partnership. *Electrical and Electronics Technical Team Roadmap*. Tech. rep. U.S. Department of Energy, June 2017. URL: <https://www.energy.gov/sites/prod/files/2017/11/f39/EETT%20Roadmap%2010-27-17.pdf>.
- [33] U.S. DRIVE Partnership. *Partnership Plan*. Tech. rep. U.S. Department of Energy, Aug. 2020. URL: https://www.energy.gov/sites/default/files/2020/08/f77/USDRIVE_Partnership_Plan_2020_APPROVED_w_newVISION.pdf.
- [34] Ajay Kumar Morya et al. “Wide Bandgap Devices in AC Electric Drives: Opportunities and Challenges”. In: *IEEE Transactions on Transportation Electrification* 5.1 (2019), pp. 3–20. DOI: 10.1109/TTE.2019.2892807.
- [35] Iman Aghabali et al. “800V Electric Vehicle Powertrains: Review and Analysis of Benefits, Challenges, and Future Trends”. In: *IEEE Transactions on Transportation Electrification* (2020), pp. 1–1. DOI: 10.1109/TTE.2020.3044938.
- [36] X. Zhang et al. “Large electric machines for aircraft electric propulsion”. In: *IET Electric Power Applications* 12.6 (2018), pp. 767–779. ISSN: 1751-8660. DOI: 10.1049/iet-epa.2017.0639.
- [37] A. D. Anderson et al. “System Weight Comparison of Electric Machine Topologies for Electric Aircraft Propulsion”. In: *2018 AIAA/IEEE Electric Aircraft Technologies Symposium (EATS)*. July 2018, pp. 1–16.
- [38] Andy K. Yoon and Kirunba Haran. “High Frequency Electric Machines for Boundary Layer Ingestion Fan Propulsor”. In: *2018 AIAA/IEEE Electric Aircraft Technologies Symposium (EATS)*. 2018, pp. 1–8.
- [39] Shajjad Chowdhury et al. “Enabling Technologies for Compact Integrated Electric Drives for Automotive Traction Applications”. In: *2019 IEEE Transportation Electrification Conference and Expo (ITEC)*. 2019, pp. 1–8. DOI: 10.1109/ITEC.2019.8790594.
- [40] John Reimers et al. “Automotive Traction Inverters: Current Status and Future Trends”. In: *IEEE Transactions on Vehicular Technology* 68.4 (2019), pp. 3337–3350. DOI: 10.1109/TVT.2019.2897899.

- [41] Jianglin Zhu et al. “High efficiency SiC traction inverter for electric vehicle applications”. In: *2018 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2018, pp. 1428–1433. DOI: 10.1109/APEC.2018.8341204.
- [42] K. Olejniczak et al. “A compact 110 kVA, 140°C ambient, 105°C liquid cooled, all-SiC inverter for electric vehicle traction drives”. In: *2017 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2017, pp. 735–742. DOI: 10.1109/APEC.2017.7930776.
- [43] Yuheng Wu et al. “An Optimized Silicon Carbide based 2×250 kW Dual Inverter for Traction Applications”. In: *2020 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2020, pp. 911–916. DOI: 10.1109/APEC39645.2020.9124593.
- [44] Tyler Van Adamson et al. “An 800-V High-Density Traction Inverter – Electro-Thermal Characterization and Low-Inductance PCB Bussing Design”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* (2020), pp. 1–1. DOI: 10.1109/JESTPE.2020.3042211.
- [45] Subhadeep Bhattacharya et al. “A Dual Three-Level T-NPC Inverter for High-Power Traction Applications”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* 4.2 (2016), pp. 668–678. DOI: 10.1109/JESTPE.2016.2517819.
- [46] Zhongjing Wang et al. “Design and Validation of A 250-kW All-Silicon Carbide High-Density Three-Level T-Type Inverter”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* 8.1 (2020), pp. 578–588. DOI: 10.1109/JESTPE.2019.2951625.
- [47] Amol Deshpande et al. “Design of a High-Efficiency, High Specific-Power Three-Level T-Type Power Electronics Building Block for Aircraft Electric-Propulsion Drives”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* 8.1 (2020), pp. 407–416. DOI: 10.1109/JESTPE.2019.2952367.
- [48] Zhao Yuan et al. “Design and Evaluation of Laminated Busbar for Three-Level T-Type NPC Power Electronics Building Block With Enhanced Dynamic Current Sharing”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* 8.1 (2020), pp. 395–406. DOI: 10.1109/JESTPE.2019.2947488.
- [49] Li Zhang et al. “Evaluation of Different Si/SiC Hybrid Three-Level Active NPC Inverters for High Power Density”. In: *IEEE Transactions on Power Electronics* 35.8 (2020), pp. 8224–8236. DOI: 10.1109/TPEL.2019.2962907.
- [50] Di Zhang, Jiangbiao He, and Di Pan. “A Megawatt-Scale Medium-Voltage High-Efficiency High Power Density “SiC+Si” Hybrid Three-Level ANPC Inverter for Aircraft Hybrid-Electric Propulsion Systems”. In: *IEEE Transactions on Industry Applications* 55.6 (2019), pp. 5971–5980. DOI: 10.1109/TIA.2019.2933513.
- [51] Handong Gui et al. “Methodology of Low Inductance Busbar Design for Three-Level Converters”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* 9.3 (2021), pp. 3468–3478. DOI: 10.1109/JESTPE.2020.2999403.

- [52] Di Pan et al. “Control of MW-Scale High-Frequency SiC+Si Multilevel ANPC Inverter in Pump-Back Test for Aircraft Hybrid-Electric Propulsion Applications”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* 9.1 (2021), pp. 1002–1012. DOI: 10.1109/JESTPE.2020.2963890.
- [53] Anton Kersten et al. “Battery Loss and Stress Mitigation in a Cascaded H-Bridge Multilevel Inverter for Vehicle Traction Applications by Filter Capacitors”. In: *IEEE Transactions on Transportation Electrification* 5.3 (2019), pp. 659–671. DOI: 10.1109/TTE.2019.2921852.
- [54] Karun Potty et al. “Performance Improvement of Medium Voltage Modular Multilevel Converter based Motor Drive using SiC MOSFETs”. In: *2020 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2020, pp. 1189–1194. DOI: 10.1109/APEC39645.2020.9124415.
- [55] Tomas Modeer et al. “An analytical method for evaluating the power density of multilevel converters”. In: *2016 IEEE 17th Workshop on Control and Modeling for Power Electronics (COMPEL)*. 2016, pp. 1–5. DOI: 10.1109/COMPEL.2016.7556759.
- [56] Jon Azurza Anderson et al. “New Figure-of-Merit Combining Semiconductor and Multi-Level Converter Properties”. In: *IEEE Open Journal of Power Electronics* 1 (2020), pp. 322–338. DOI: 10.1109/OJPEL.2020.3018220.
- [57] T. A. Meynard and H. Foch. “Multi-level conversion: high voltage choppers and voltage-source inverters”. In: *Proc. 23rd Annu. IEEE Power Electron. Specialists Conf. (PESC)*. June 1992, 397–403 vol.1. DOI: 10.1109/PESC.1992.254717.
- [58] Yutian Lei, Wen-Chuen Liu, and Robert Carl Nikolai Pilawa-Podgurski. “An Analytical Method to Evaluate and Design Hybrid Switched-Capacitor and Multilevel Converters”. In: *IEEE Transactions on Power Electronics* 33.3 (2018), pp. 2227–2240. DOI: 10.1109/TPEL.2017.2690324.
- [59] Christopher B. Barth et al. “Experimental evaluation of capacitors for power buffering in single-phase power converters”. In: *2015 IEEE Energy Conversion Congress and Exposition (ECCE)*. 2015, pp. 6269–6276. DOI: 10.1109/ECCE.2015.7310539.
- [60] Andrew Stillwell and Robert C. N. Pilawa-Podgurski. “Design Optimization of a 1500 V GaN-Based Solar Inverter Using Flying Capacitor Multi-Level Converter Stages”. In: *2019 IEEE Energy Conversion Congress and Exposition (ECCE)*. 2019, pp. 4605–4612. DOI: 10.1109/ECCE.2019.8913007.
- [61] Samantha Coday, Ansel Barchowsky, and Robert C.N. Pilawa-Podgurski. “A 10-level GaN-based Flying Capacitor Multilevel Boost Converter for Radiation-Hardened Operation in Space Applications”. In: *2021 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2021, pp. 2798–2803. DOI: 10.1109/APEC42165.2021.9487143.

- [62] T. Modeer et al. “Design of a GaN-Based Interleaved Nine-Level Flying Capacitor Multilevel Inverter for Electric Aircraft Applications”. In: *IEEE Transactions on Power Electronics* 35.11 (2020), pp. 12153–12165. DOI: 10.1109/TPEL.2020.2989329.
- [63] Richardt H. Wilkinson, Thierry A. Meynard, and Hendrik du Toit Mouton. “Natural Balance of Multicell Converters: The General Case”. In: *IEEE Transactions on Power Electronics* 21.6 (2006), pp. 1658–1666. DOI: 10.1109/TPEL.2006.882951.
- [64] Brendan Peter McGrath et al. “Optimal Modulation of Flying Capacitor and Stacked Multicell Converters Using a State Machine Decoder”. In: *IEEE Transactions on Power Electronics* 22.2 (2007), pp. 508–516. DOI: 10.1109/TPEL.2006.889932.
- [65] Yutian Lei et al. “A 2 kW, single-phase, 7-level, GaN inverter with an active energy buffer achieving 216 W/in³ power density and 97.6% peak efficiency”. In: *2016 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2016, pp. 1512–1519. DOI: 10.1109/APEC.2016.7468068.
- [66] Christopher B. Barth et al. “Design and control of a GaN-based, 13-level, flying capacitor multilevel inverter”. In: *2016 IEEE 17th Workshop on Control and Modeling for Power Electronics (COMPEL)*. 2016, pp. 1–6. DOI: 10.1109/COMPEL.2016.7556770.
- [67] Tomas Modeer et al. “Design of a GaN-based, 9-level flying capacitor multilevel inverter with low inductance layout”. In: *2017 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2017, pp. 2582–2589. DOI: 10.1109/APEC.2017.7931062.
- [68] Yutian Lei et al. “A 2-kW Single-Phase Seven-Level Flying Capacitor Multilevel Inverter With an Active Energy Buffer”. In: *IEEE Transactions on Power Electronics* 32.11 (2017), pp. 8570–8581. DOI: 10.1109/TPEL.2017.2650140.
- [69] Andrew Stillwell, Yutian Lei, and Robert C. N. Pilawa-Podgurski. “A method to extract low-voltage auxiliary power from a flying capacitor multi-level converter”. In: *2016 IEEE 17th Workshop on Control and Modeling for Power Electronics (COMPEL)*. 2016, pp. 1–8. DOI: 10.1109/COMPEL.2016.7556747.
- [70] Andrew Stillwell, Margaret E. Blackwell, and Robert C. N. Pilawa-Podgurski. “Design of a 1 kV bidirectional DC-DC converter with 650 V GaN transistors”. In: *2018 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2018, pp. 1155–1162. DOI: 10.1109/APEC.2018.8341162.
- [71] M. E. Blackwell, A. Stillwell, and R. C. N. Pilawa-Podgurski. “Dynamic Level Selection for Full Range ZVS in Flying Capacitor Multi-Level Converters”. In: *2018 IEEE 19th Workshop on Control and Modeling for Power Electronics (COMPEL)*. 2018, pp. 1–8. DOI: 10.1109/COMPEL.2018.8460086.

- [72] Andrew Stillwell, Enver Candan, and Robert C. N. Pilawa-Podgurski. “Active Voltage Balancing in Flying Capacitor Multi-Level Converters With Valley Current Detection and Constant Effective Duty Cycle Control”. In: *IEEE Transactions on Power Electronics* 34.11 (2019), pp. 11429–11441. DOI: 10.1109/TPEL.2019.2899899.
- [73] Nathan Pallo et al. “Power-dense multilevel inverter module using interleaved GaN-based phases for electric aircraft propulsion”. In: *2018 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2018, pp. 1656–1661. DOI: 10.1109/APEC.2018.8341239.
- [74] G. Su and L. Tang. “A segmented traction drive system with a small dc bus capacitor”. In: *2012 IEEE Energy Conversion Congress and Exposition (ECCE)*. Sept. 2012, pp. 2847–2853. DOI: 10.1109/ECCE.2012.6342375.
- [75] E. Gurpinar and B. Ozpineci. “Loss Analysis and Mapping of a SiC MOSFET Based Segmented Two-Level Three-Phase Inverter for EV Traction Systems”. In: *2018 IEEE Transportation Electrification Conference and Expo (ITEC)*. June 2018, pp. 1046–1053. DOI: 10.1109/ITEC.2018.8450188.
- [76] Zichao Ye et al. “Investigation of capacitor voltage balancing in practical implementations of flying capacitor multilevel converters”. In: *2017 IEEE 18th Workshop on Control and Modeling for Power Electronics (COMPEL)*. 2017, pp. 1–7. DOI: 10.1109/COMPEL.2017.8013368.
- [77] T. Foulkes, T. Modeer, and R. C. N. Pilawa-Podgurski. “Quantifying Dynamic On-State Resistance of GaN HEMTs for Power Converter Design via a Survey of Low and High Voltage Devices”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* (2020), pp. 1–1. DOI: 10.1109/JESTPE.2020.3024930.
- [78] Samantha Coday and Robert C.N. Pilawa-Podgurski. “High Accuracy Calorimetric Measurements and Modeling of Ceramic Capacitor Losses Under Large Ripple Operation”. In: *2020 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2020, pp. 188–194. DOI: 10.1109/APEC39645.2020.9124490.
- [79] Christopher B. Barth et al. “Experimental Evaluation of Capacitors for Power Buffering in Single-Phase Power Converters”. In: *IEEE Transactions on Power Electronics* 34.8 (2019), pp. 7887–7899. DOI: 10.1109/TPEL.2018.2878825.
- [80] Nathan Pallo et al. “Modular heat sink for chip-scale GaN transistors in multilevel converters”. In: *2018 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2018, pp. 2798–2805. DOI: 10.1109/APEC.2018.8341414.
- [81] Zichao Ye and Robert C. N. Pilawa-Podgurski. “A power supply circuit for gate driver of GaN-based flying capacitor multi-level converters”. In: *2016 IEEE 4th Workshop on Wide Bandgap Power Devices and Applications (WiPDA)*. 2016, pp. 53–58. DOI: 10.1109/WiPDA.2016.7799909.

- [82] Zichao Ye et al. “Design and implementation of a low-cost and compact floating gate drive power circuit for GaN-based flying capacitor multi-level converters”. In: *2017 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2017, pp. 2925–2931. DOI: 10.1109/APEC.2017.7931112.
- [83] Zichao Ye et al. “Improved Bootstrap Methods for Powering Floating Gate Drivers of Flying Capacitor Multilevel Converters and Hybrid Switched-Capacitor Converters”. In: *IEEE Transactions on Power Electronics* 35.6 (2020), pp. 5965–5977. DOI: 10.1109/TPEL.2019.2951116.
- [84] Logan Horowitz et al. “A Method of Partial Inductances to Evaluate and Optimize Switching Cells”. In: *2021 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2021.
- [85] Amirnaser Yazdani and Reza Iravani. *Voltage-Sourced Converters in Power Systems: Modeling, Control, and Applications*. John Wiley and Sons, Jan. 26, 2010. 472 pp. ISBN: 0470521562. URL: https://www.ebook.de/de/product/9338772/yazdani_iravani_converters.html.
- [86] J. Lu, R. Hou, and D. Chen. “Loss Distribution among Paralleled GaN HEMTs”. In: *2018 IEEE Energy Conversion Congress and Exposition (ECCE)*. 2018, pp. 1914–1919. DOI: 10.1109/ECCE.2018.8557988.
- [87] J. Gareau, R. Hou, and A. Emadi. “Review of Loss Distribution, Analysis, and Measurement Techniques for GaN HEMTs”. In: *IEEE Transactions on Power Electronics* 35.7 (2020), pp. 7405–7418. DOI: 10.1109/TPEL.2019.2954819.
- [88] Z. Lin. “Study on the Intrinsic Origin of Output Capacitor Hysteresis in Advanced Superjunction MOSFETs”. In: *IEEE Electron Device Letters* 40.8 (2019), pp. 1297–1300. DOI: 10.1109/LED.2019.2924512.
- [89] M. Guacci et al. “On the Origin of the C_{oss} -Losses in Soft-Switching GaN-on-Si Power HEMTs”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* 7.2 (2019), pp. 679–694. DOI: 10.1109/JESTPE.2018.2885442.
- [90] Z. Tong et al. “Origins of Soft-switching Coss Losses in SiC Power MOSFETs and Diodes for Resonant Converter Applications”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* (2020), pp. 1–1. DOI: 10.1109/JESTPE.2020.3034345.
- [91] *Power MOSFET Basics: Understanding Gate Charge and Using it to Assess Switching Performance*. Application Note AN608A. Vishay Siliconix, Feb. 2016.
- [92] Julian Styles. “Common misconceptions about the body diode”. In: *Design World: Power and Energy Efficiency Handbook* (Oct. 2019), pp. 12–14. URL: <https://www.designworldonline.com/october-2019-special-edition-power-energy-efficiency-handbook/>.

- [93] T. Cappello, A. Santarelli, and C. Florian. “Dynamic RON Characterization Technique for the Evaluation of Thermal and Off-State Voltage Stress of GaN Switches”. In: *IEEE Transactions on Power Electronics* 33.4 (2018), pp. 3386–3398. DOI: 10.1109/TPEL.2017.2710281.
- [94] R. Hou et al. “Power Loss Characterization and Modeling for GaN-Based Hard-Switching Half-Bridges Considering Dynamic on-State Resistance”. In: *IEEE Transactions on Transportation Electrification* 6.2 (2020), pp. 540–553. DOI: 10.1109/TTE.2020.2989036.
- [95] K. Li et al. “Accurate Measurement of Dynamic on-State Resistances of GaN Devices Under Reverse and Forward Conduction in High Frequency Power Converter”. In: *IEEE Transactions on Power Electronics* 35.9 (2020), pp. 9650–9660. DOI: 10.1109/TPEL.2019.2961604.
- [96] G. Zulauf, M. Guacci, and J. W. Kolar. “Dynamic on-Resistance in GaN-on-Si HEMTs: Origins, Dependencies, and Future Characterization Frameworks”. In: *IEEE Transactions on Power Electronics* 35.6 (2020), pp. 5581–5588. DOI: 10.1109/TPEL.2019.2955656.
- [97] G. Zulauf et al. “The Impact of Multi-MHz Switching Frequencies on Dynamic On-Resistance in GaN-on-Si HEMTs”. In: *IEEE Open Journal of Power Electronics* 1 (2020), pp. 210–215. DOI: 10.1109/OJPEL.2020.3005879.
- [98] David Menzi, Morris Heller, and Johann W. Kolar. “iGSE-C_x – a New Normalized Steinmetz Model for Class II Multilayer Ceramic Capacitors”. In: *IEEE Open Journal of Power Electronics* 2 (2021), pp. 138–144. DOI: 10.1109/OJPEL.2021.3060874.
- [99] David Menzi et al. “Novel iGSE-C Loss Modeling of X7R Ceramic Capacitors”. In: *IEEE Transactions on Power Electronics* 35.12 (2020), pp. 13367–13383. DOI: 10.1109/TPEL.2020.2996010.
- [100] D. Neumayr et al. “Comprehensive large-signal performance analysis of ceramic capacitors for power pulsation buffers”. In: *2016 IEEE 17th Workshop on Control and Modeling for Power Electronics (COMPEL)*. 2016, pp. 1–8. DOI: 10.1109/COMPEL.2016.7556762.
- [101] W.-J. Gu and R. Liu. “A study of volume and weight vs. frequency for high-frequency transformers”. In: *Proceedings of IEEE Power Electronics Specialist Conference - PESC '93*. 1993, pp. 1123–1129. DOI: 10.1109/PESC.1993.472059.
- [102] C.R. Sullivan. “Optimal choice for number of strands in a litz-wire transformer winding”. In: *IEEE Transactions on Power Electronics* 14.2 (1999), pp. 283–291. DOI: 10.1109/63.750181.
- [103] W.G. Hurley, E. Gath, and J.G. Breslin. “Optimizing the AC resistance of multilayer transformer windings with arbitrary current waveforms”. In: *IEEE Transactions on Power Electronics* 15.2 (2000), pp. 369–376. DOI: 10.1109/63.838110.

- [104] Jieli Li, T. Abdallah, and C.R. Sullivan. “Improved calculation of core loss with non-sinusoidal waveforms”. In: *Conference Record of the 2001 IEEE Industry Applications Conference. 36th IAS Annual Meeting (Cat. No.01CH37248)*. Vol. 4. 2001, 2203–2210 vol.4. DOI: 10.1109/IAS.2001.955931.
- [105] K. Venkatachalam et al. “Accurate prediction of ferrite core loss with nonsinusoidal waveforms using only Steinmetz parameters”. In: *2002 IEEE Workshop on Computers in Power Electronics, 2002. Proceedings.* 2002, pp. 36–41. DOI: 10.1109/CIPE.2002.1196712.
- [106] Yehui Han et al. “Evaluation of Magnetic Materials for Very High Frequency Power Applications”. In: *IEEE Transactions on Power Electronics* 27.1 (2012), pp. 425–435. DOI: 10.1109/TPEL.2011.2159995.
- [107] Jonas Muhlethaler et al. “Core Losses Under the DC Bias Condition Based on Steinmetz Parameters”. In: *IEEE Transactions on Power Electronics* 27.2 (2012), pp. 953–963. DOI: 10.1109/TPEL.2011.2160971.
- [108] Jonas Muhlethaler et al. “Improved Core-Loss Calculation for Magnetic Components Employed in Power Electronic Systems”. In: *IEEE Transactions on Power Electronics* 27.2 (2012), pp. 964–973. DOI: 10.1109/TPEL.2011.2162252.
- [109] Ranjith Bramanpalli. *Accurate Inductor Loss Determination using Würth Elektronik’s REDEXPERT*. Tech. rep. ANP029. Würth Elektronik, June 2015. URL: https://www.we-online.com/web/en/electronic_components/produkte_pb/application_notes/exaktebestimmungvonspulenverlusten.php.
- [110] *Choosing Inductors for Energy Efficient Power Applications*. Tech. rep. 1400. Coilcraft, July 2020. URL: https://www.coilcraft.com/pdfs/Doc1400_Choosing_Inductors_for_Energy_Efficiency.pdf.
- [111] Panteleimon Papamanolis, Dominik Neumayr, and Johann W. Kolar. “Behavior of the flying capacitor converter under critical operating conditions”. In: *2017 IEEE 26th International Symposium on Industrial Electronics (ISIE)*. 2017, pp. 628–635. DOI: 10.1109/ISIE.2017.8001319.
- [112] A. Ruderman, B. Reznikov, and M. Margaliot. “Simple analysis of a flying capacitor converter voltage balance dynamics for DC modulation”. In: *2008 13th International Power Electronics and Motion Control Conference*. Sept. 2008, pp. 260–267. DOI: 10.1109/EPEPEMC.2008.4635276.
- [113] S. Thielemans, A. Ruderman, and J. Melkebeek. “Self-precharge in single-leg flying capacitor converters”. In: *2009 35th Annual Conference of IEEE Industrial Electronics*. Nov. 2009, pp. 812–817. DOI: 10.1109/IECON.2009.5415024.
- [114] A. Penczek et al. “Simple time-domain analysis of a multilevel DC–DC flying capacitor converter average aperiodic natural balancing dynamics”. In: *IET Power Electronics* 12.5 (2019), pp. 1179–1186. ISSN: 1755-4535. DOI: 10.1049/iet-pe1.2018.5070.

- [115] T.A. Meynard, M. Fadel, and N. Aouda. “Modeling of multilevel converters”. In: *IEEE Transactions on Industrial Electronics* 44.3 (1997), pp. 356–364. DOI: 10.1109/41.585833.
- [116] Brendan Peter McGrath and Donald Grahame Holmes. “Analytical Modelling of Voltage Balance Dynamics for a Flying Capacitor Multilevel Converter”. In: *IEEE Transactions on Power Electronics* 23.2 (2008), pp. 543–550. DOI: 10.1109/TPEL.2007.915175.
- [117] David Reusch and Johan Strydom. “Understanding the Effect of PCB Layout on Circuit Performance in a High-Frequency Gallium-Nitride-Based Point of Load Converter”. In: *IEEE Transactions on Power Electronics* 29.4 (2014), pp. 2008–2015. DOI: 10.1109/TPEL.2013.2266103.
- [118] Derek Chou, Kelly Fernandez, and Robert C. N. Pilawa-Podgurski. “An Interleaved 6-Level GaN Bidirectional Converter for Level II Electric Vehicle Charging”. In: *2019 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2019, pp. 594–600. DOI: 10.1109/APEC.2019.8721971.
- [119] Nathan Pallo, Tomas Modeer, and Robert C.N. Pilawa-Podgurski. “Electrically thin approach to switching cell design for flying capacitor multilevel converters”. In: *2017 IEEE 5th Workshop on Wide Bandgap Power Devices and Applications (WiPDA)*. 2017, pp. 411–416. DOI: 10.1109/WiPDA.2017.8170582.
- [120] *TDK Equivalent Circuit Model Library, C5750X6S2W225K250KA*. Tech. rep. 2015. URL: https://product.tdk.com/system/files/dam/technicalsupport/tvcl/pdf/mlcc_commercial_midvoltage_c5750_ecm.pdf.
- [121] Amanda Jackson et al. “A Modular Multi-Phase Actively Controlled Resistive Load Bank with Zero-Current Switching Capability and Integrated Snubbers”. In: *2021 IEEE Power and Energy Conference at Illinois (PECI)*. 2021, pp. 1–6. DOI: 10.1109/PECI51586.2021.9435213.
- [122] Yanjun Shi et al. “Design and implementation of a 100 kW SiC filter-less PV inverter with 5 kW/kg power density and 99.2% CEC efficiency”. In: *2018 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2018, pp. 393–398. DOI: 10.1109/APEC.2018.8341041.
- [123] Amol Deshpande et al. “A three-level, T-type, power electronics building block using Si-SiC hybrid switch for high-speed drives”. In: *2018 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2018, pp. 2609–2616. DOI: 10.1109/APEC.2018.8341385.
- [124] Eduardo Laloya et al. “Heat Management in Power Converters: From State of the Art to Future Ultrahigh Efficiency Systems”. In: *IEEE Transactions on Power Electronics* 31.11 (2016), pp. 7896–7908. DOI: 10.1109/TPEL.2015.2513433.

- [125] Sina Lohrasbi et al. “A Comprehensive Review on the Core Thermal Management Improvement Concepts in Power Electronics”. In: *IEEE Access* 8 (2020), pp. 166880–166906. DOI: 10.1109/ACCESS.2020.3021946.
- [126] I. Mudawar. “Assessment of high-heat-flux thermal management schemes”. In: *IEEE Transactions on Components and Packaging Technologies* 24.2 (2001), pp. 122–141. DOI: 10.1109/6144.926375.
- [127] Derek Chou et al. “An Interleaved 6-Level GaN Bidirectional Converter With an Active Energy Buffer for Level II Electric Vehicle Charging”. In: *2021 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2021.
- [128] Ki Wook Jung et al. “Embedded cooling with 3D manifold for vehicle power electronics application: Single-phase thermal-fluid performance”. In: *International Journal of Heat and Mass Transfer* 130 (2019), pp. 1108–1119. ISSN: 0017-9310. DOI: <https://doi.org/10.1016/j.ijheatmasstransfer.2018.10.108>. URL: <https://www.sciencedirect.com/science/article/pii/S0017931018326206>.
- [129] Remco van Erp, Georgios Kampitsis, and Elison Matioli. “Efficient Microchannel Cooling of Multiple Power Devices With Compact Flow Distribution for High Power-Density Converters”. In: *IEEE Transactions on Power Electronics* 35.7 (July 2020), pp. 7235–7245. DOI: 10.1109/TPEL.2019.2959736.
- [130] Angie Fan et al. “An innovative passive cooling method for high performance light-emitting diodes”. In: *2012 28th Annual IEEE Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM)*. 2012, pp. 319–324. DOI: 10.1109/STHERM.2012.6188867.
- [131] Yiyi Chen et al. “Direct Phase-Change Cooling of Vapor Chamber Integrated With IGBT Power Electronic Module for Automotive Application”. In: *IEEE Transactions on Power Electronics* 36.5 (2021), pp. 5736–5747. DOI: 10.1109/TPEL.2020.3031372.
- [132] Thomas Foulkes et al. “Fundamental limits of jumping droplet heat transfer”. In: *Applied Physics Letters* 116.9 (2020), p. 093701. DOI: 10.1063/1.5141744. eprint: <https://doi.org/10.1063/1.5141744>. URL: <https://doi.org/10.1063/1.5141744>.
- [133] Tiwei Wei et al. “High-Efficiency Polymer-Based Direct Multi-Jet Impingement Cooling Solution for High-Power Devices”. In: *IEEE Transactions on Power Electronics* 34.7 (2019), pp. 6601–6612. DOI: 10.1109/TPEL.2018.2872904.
- [134] Samantha Jones-Jackson, Romina Rodriguez, and Ali Emadi. “Jet Impingement Cooling in Power Electronics for Electrified Automotive Transportation: Current Status and Future Trends”. In: *IEEE Transactions on Power Electronics* 36.9 (2021), pp. 10420–10435. DOI: 10.1109/TPEL.2021.3059558.

- [135] Issam Mudawar, Peter E. Jimenez, and Robert E. Morgan. “Immersion-Cooled Standard Electronic Clamshell Module: A Building Block for Future High-Flux Avionic Systems”. In: *Journal of Electronic Packaging* 116.2 (June 1994), pp. 116–125. DOI: 10.1115/1.2905499.
- [136] Peter E. Jimenez and Issam Mudawar. “A Multi-Kilowatt Immersion-Cooled Standard Electronic Clamshell Module for Future Aircraft Avionics”. In: *Journal of Electronic Packaging* 116.3 (Sept. 1994), pp. 220–229. DOI: 10.1115/1.2905689.
- [137] Patrick Birbarah et al. “Water immersion cooling of high power density electronics”. In: *International Journal of Heat and Mass Transfer* 147 (2020), p. 118918. ISSN: 0017-9310. DOI: <https://doi.org/10.1016/j.ijheatmasstransfer.2019.118918>. URL: <https://www.sciencedirect.com/science/article/pii/S0017931019336002>.
- [138] Uwe Drofenik, Andrija Stupar, and Johann W. Kolar. “Analysis of Theoretical Limits of Forced-Air Cooling Using Advanced Composite Materials With High Thermal Conductivities”. In: *IEEE Transactions on Components, Packaging and Manufacturing Technology* 1.4 (2011), pp. 528–535. DOI: 10.1109/TCPMT.2010.2100730.
- [139] Andrija Stupar, Uwe Drofenik, and Johann W. Kolar. “Optimization of Phase Change Material Heat Sinks for Low Duty Cycle High Peak Load Power Supplies”. In: *IEEE Transactions on Components, Packaging and Manufacturing Technology* 2.1 (2012), pp. 102–115. DOI: 10.1109/TCPMT.2011.2168957.
- [140] Ercan M. Dede et al. “Optimal Design of Three-Dimensional Heat Flow Structures for Power Electronics Applications”. In: *Journal of Thermal Science and Engineering Applications* 11.2 (Dec. 2018). DOI: 10.1115/1.4041440.
- [141] Mohammad Anwar et al. “Traction Power Inverter Design for EV and HEV Applications at General Motors: A Review”. In: *2019 IEEE Energy Conversion Congress and Exposition (ECCE)*. 2019, pp. 6346–6351. DOI: 10.1109/ECCE.2019.8913098.
- [142] Cai Chen et al. “Flexible PCB-Based 3-D Integrated SiC Half-Bridge Power Module With Three-Sided Cooling Using Ultralow Inductive Hybrid Packaging Structure”. In: *IEEE Transactions on Power Electronics* 34.6 (2019), pp. 5579–5593. DOI: 10.1109/TPEL.2018.2866404.
- [143] Christopher Barth et al. “Experimental evaluation of a 1 kW, single-phase, 3-level gallium nitride inverter in extreme cold environment”. In: *2017 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2017, pp. 717–723. DOI: 10.1109/APEC.2017.7930773.
- [144] Nathan Andrew Pallo. “Overcoming electro-thermal barriers to achieve extreme performance power conversion for more electric aircraft”. MA thesis. University of Illinois at Urbana-Champaign, 2018. URL: <http://hdl.handle.net/2142/101087>.

- [145] Xuhui Feng et al. “Investigation of thermal interface materials using phase-sensitive transient thermoreflectance technique”. In: *Fourteenth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*. 2014, pp. 1296–1307. DOI: 10.1109/ITHERM.2014.6892430.
- [146] *Understanding Fan Curves*. Tech. rep. FE-2000. Twin City Fan, 2018.
- [147] Enver Candan et al. “A 6-level Flying Capacitor Multi-level Converter for Single Phase Buck-type Power Factor Correction”. In: *2019 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2019, pp. 1180–1187. DOI: 10.1109/APEC.2019.8722199.
- [148] M. Iyengar and A. Bar-Cohen. “Design for manufacturability of SISE parallel plate forced convection heat sinks”. In: *IEEE Transactions on Components and Packaging Technologies* 24.2 (2001), pp. 150–158. DOI: 10.1109/6144.926377.
- [149] Christoph Gammeter, Florian Krismer, and Johann W. Kolar. “Weight Optimization of a Cooling System Composed of Fan and Extruded-Fin Heat Sink”. In: *IEEE Transactions on Industry Applications* 51.1 (2015), pp. 509–520. DOI: 10.1109/TIA.2014.2336977.
- [150] Seri Lee. “Calculating Spreading Resistance in Heat Sinks”. In: *Electronics COOLING* (1998). URL: https://www.electronics-cooling.com/wp-content/uploads/2020/09/Electronics-Cooling_Fall-2020-Issue.pdf.
- [151] Roger C. Baker. *Flow Measurement Handbook: Industrial Designs, Operating Principles, Performance, and Applications*. Cambridge University Press, 2000. DOI: 10.1017/CB09780511471100.
- [152] Michael A. Crabtree. *The Concise Industrial Flow Measurement Handbook*. CRC Press, 2020.
- [153] *Measurement of fluid flow by means of pressure differential devices inserted in circular-cross section conduits running - Part 2: Orifice plates*. International Special Committee on Radio Interference, Mar. 2003.
- [154] François Corre. *HydrauCalc*. Version 2020b. 2020. URL: <https://www.hydraucalc.com/>.
- [155] Sreekant Narumanchi et al. “Thermal interface materials for power electronics applications”. In: *2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*. 2008, pp. 395–404. DOI: 10.1109/ITHERM.2008.4544297.
- [156] *PrioAir 8 EC Inline Duct Fan Technical Reference Performance Curve*. Fantech. URL: <https://shop.fantech.net/en-US/prioair--8--ec--inline--duct--fan/p99451>.

- [157] A. Shea and T. M. Jahns. “Hardware integration for an integrated modular motor drive including distributed control”. In: *2014 IEEE Energy Conversion Congress and Exposition (ECCE)*. 2014, pp. 4881–4887.
- [158] F. Savi et al. “Information Technologies for Distributed Machine Drives: An Overview”. In: *2019 IEEE International Electric Machines Drives Conference (IEMDC)*. 2019, pp. 1805–1809.
- [159] “Enhanced Pulse Width Modulator (ePWM) Module”. In: *Technical Reference Manual TMS320x2833x, TMS320x2823x*. Literature Number SPRUI07. Texas Instruments, Sept. 2019.
- [160] Roland E. Best. “Mixed-Signal PLL Building Blocks”. In: *Phase-locked Loops: Design, Simulation, and Applications*. McGraw-Hill professional engineering. McGraw-Hill, 2007. ISBN: 9780071499262. URL: <https://books.google.com/books?id=7YAn1QEACAAJ>.
- [161] V. Kratyuk et al. “A Design Procedure for All-Digital Phase-Locked Loops Based on a Charge-Pump Phase-Locked-Loop Analogy”. In: *IEEE Transactions on Circuits and Systems II: Express Briefs* 54.3 (2007), pp. 247–251. DOI: 10.1109/TCSII.2006.889443.
- [162] Curtis Barrett. *Fractional/Integer-N PLL Basics*. Tech. rep. SWRA029. Texas Instruments, 1999.
- [163] Luca Corradini et al. *Digital Control of High-Frequency Switched-Mode Power Converters*. WILEY, July 13, 2015. 368 pp. ISBN: 1118935101. URL: https://www.ebook.de/de/product/22512463/luca_corradini_dragan_maksimovic_paolo_mattavelli_regan_zane_digital_control_of_high_frequency_switched_mode_power_converters.html.
- [164] D. Grahame Holmes and Thomas A. Lipo. *Pulse Width Modulation for Power Converters: Principles and Practice*. IEEE Press Series on Power Engineering. Wiley, 2003. ISBN: 9780471208143.
- [165] Brendan Peter McGrath and Donald Grahame Holmes. “Analytical Determination of the Capacitor Voltage Balancing Dynamics for Three-Phase Flying Capacitor Converters”. In: *IEEE Transactions on Industry Applications* 45.4 (2009), pp. 1425–1433. DOI: 10.1109/TIA.2009.2023480.
- [166] B. P. McGrath and D. G. Holmes. “A General Analytical Method for Calculating Inverter DC-Link Current Harmonics”. In: *IEEE Transactions on Industry Applications* 45.5 (2009), pp. 1851–1859.
- [167] Dimitri P. Bertsekas and John N. Tsitsiklis. *Introduction to Probability*. 2nd. Athena Scientific, 2008. ISBN: 9781886529236. URL: <https://books.google.com/books?id=8LGi1AjSfpcC>.
- [168] *47 C. F. R. §15*. Electronic Code of Federal Regulations. URL: <https://www.ecfr.gov/cgi-bin/text-idx?node=pt47.1.15>.

- [169] *Electromagnetic compatibility of multimedia equipment - Emission requirements*. International Special Committee on Radio Interference, Oct. 2019.
- [170] *Information technology equipment - Radio disturbance characteristics - Limits and methods of measurement*. International Special Committee on Radio Interference, Sept. 2008.
- [171] Timothy Hegarty. *An overview of conducted EMI specifications for power supplies*. Tech. rep. SLYY136. Texas Instruments, 2018.
- [172] *Vehicles, boats and internal combustion engines – Radio disturbance characteristics – Limits and methods of measurement for the protection of on-board receivers*. International Special Committee on Radio Interference, Oct. 2016.
- [173] Christopher Brandon Barth. “Compact and Efficient Power Electronics System Design for Automotive, Solar, and Aerospace Applications”. Ph.D. dissertation. Urbana, IL: University of Illinois at Urbana-Champaign, Dec. 2019. URL: <http://hdl.handle.net/2142/106246>.
- [174] J. W. Bennett et al. “Fault-Tolerant Design Considerations and Control Strategies for Aerospace Drives”. In: *IEEE Transactions on Industrial Electronics* 59.5 (2012), pp. 2049–2058. DOI: 10.1109/TIE.2011.2159356.
- [175] Y.C. Yeh. “Triple-triple redundant 777 primary flight computer”. In: *1996 IEEE Aerospace Applications Conference. Proceedings*. Vol. 1. 1996, 293–307 vol.1. DOI: 10.1109/AERO.1996.495891.
- [176] Frede Blaabjerg et al. “Reliability of Power Electronic Systems for EV/HEV Applications”. In: *Proceedings of the IEEE* (2020), pp. 1–17. DOI: 10.1109/JPROC.2020.3031041.
- [177] Marvin Rausand, Anne Barros, and Arnljot Hoyland. *System Reliability Theory*. Wiley series in probability and statistics. Applied probability and statistics. Wiley John + Sons, Jan. 4, 2021. 864 pp. ISBN: 1119373522. URL: https://www.ebook.de/de/product/39286344/marvin_rausand_anne_barros_arnljot_hoyland_system_reliability_theory.html.
- [178] M. Cozzolino and G. Ewell. “A Fracture Mechanics Approach to Structural Reliability of Ceramic Capacitors”. In: *IEEE Transactions on Components, Hybrids, and Manufacturing Technology* 3.2 (1980), pp. 250–257. DOI: 10.1109/TCHMT.1980.1135611.
- [179] *Mechanical Cracking: The Major Cause for Multilayer Ceramic Capacitor Failures*. Tech. rep. AN0005. Knowles.
- [180] Y. Song and B. Wang. “Survey on Reliability of Power Electronic Systems”. In: *IEEE Transactions on Power Electronics* 28.1 (2013), pp. 591–604. DOI: 10.1109/TPEL.2012.2192503.

- [181] Vivek Raveendran, Markus Andresen, and Marco Liserre. “Improving Onboard Converter Reliability for More Electric Aircraft With Lifetime-Based Control”. In: *IEEE Transactions on Industrial Electronics* 66.7 (2019), pp. 5787–5796. DOI: 10.1109/TIE.2018.2889626.
- [182] M. Calabretta et al. “Power Semiconductor Devices and Packages: Solder Mechanical Characterization and Lifetime Prediction”. In: *IEEE Access* 9 (2021), pp. 22859–22867. DOI: 10.1109/ACCESS.2021.3056281.
- [183] Jayakrishnan Harikumar et al. “Failure Modes and Reliability Oriented System Design for Aerospace Power Electronic Converters”. In: *IEEE Open Journal of the Industrial Electronics Society* 2 (2021), pp. 53–64. DOI: 10.1109/OJIES.2020.3047201.
- [184] H. Wang and F. Blaabjerg. “Reliability of Capacitors for DC-Link Applications in Power Electronic Converters—An Overview”. In: *IEEE Transactions on Industry Applications* 50.5 (2014), pp. 3569–3578. DOI: 10.1109/TIA.2014.2308357.
- [185] B. S. Rawal and N. H. Chan. *Conduction and Failure Mechanisms in Barium Titanate Based Ceramics under Highly Accelerated Conditions*. Tech. rep. AVX Corporation, 1988.
- [186] Dirk Hirschmann et al. “Reliability Prediction for Inverters in Hybrid Electrical Vehicles”. In: *IEEE Transactions on Power Electronics* 22.6 (2007), pp. 2511–2517. DOI: 10.1109/TPEL.2007.909236.
- [187] E. Gurpinar et al. “Reliability-Driven Assessment of GaN HEMTs and Si IGBTs in 3L-ANPC PV Inverters”. In: *IEEE Journal of Emerging and Selected Topics in Power Electronics* 4.3 (2016), pp. 956–969. DOI: 10.1109/JESTPE.2016.2566259.
- [188] Adarsh Nagarajan et al. *Photovoltaic Inverter Reliability Assessment*. Tech. rep. NREL/TP-5D00-74462. Golden, CO: National Renewable Energy Laboratory, 2019.
- [189] Luca Concari et al. “Assessment of Efficiency and Reliability of Wide Band-Gap Based H8 Inverter in Electric Vehicle Applications”. In: *Energies* 12.10 (2019). ISSN: 1996-1073. DOI: 10.3390/en12101922. URL: <https://www.mdpi.com/1996-1073/12/10/1922>.
- [190] Jonathan Swingler. *Reliability Characterisation of Electrical and Electronic Systems*. English. Netherlands: Elsevier, Jan. 2015. ISBN: 9781782422259.
- [191] Saeed Peyghami, Zhongxu Wang, and Frede Blaabjerg. “A Guideline for Reliability Prediction in Power Electronic Converters”. In: *IEEE Transactions on Power Electronics* 35.10 (2020), pp. 10958–10968. DOI: 10.1109/TPEL.2020.2981933.
- [192] Henry Shu-hung Chung et al., eds. *Reliability of Power Electronic Converter Systems*. Power and Energy. Institution of Engineering and Technology, 2015. URL: <https://digital-library.theiet.org/content/books/po/pbpo080e>.

- [193] Enrico Zanoni et al. “AlGaN/GaN-Based HEMTs Failure Physics and Reliability: Mechanisms Affecting Gate Edge and Schottky Junction”. In: *IEEE Transactions on Electron Devices* 60.10 (2013), pp. 3119–3131. DOI: 10.1109/TED.2013.2271954.
- [194] Ronald A. Coutu et al. “Benefits of Considering More than Temperature Acceleration for GaN HEMT Life Testing”. In: *Electronics* 5.3 (2016). ISSN: 2079-9292. DOI: 10.3390/electronics5030032. URL: <https://www.mdpi.com/2079-9292/5/3/32>.
- [195] J.A. del Alamo and J. Joh. “GaN HEMT reliability”. In: *Microelectronics Reliability* 49.9 (2009). 20th European Symposium on the Reliability of Electron Devices, Failure Physics and Analysis, pp. 1200–1206. ISSN: 0026-2714. DOI: <https://doi.org/10.1016/j.microrel.2009.07.003>. URL: <https://www.sciencedirect.com/science/article/pii/S0026271409002364>.
- [196] Jesús A. del Alamo and Ethan S. Lee. “Stability and Reliability of Lateral GaN Power Field-Effect Transistors”. In: *IEEE Transactions on Electron Devices* 66.11 (2019), pp. 4578–4590. DOI: 10.1109/TED.2019.2931718.
- [197] A. Bensoussan. “How to quantify and predict long term multiple stress operation: Application to Normally-Off Power GaN transistor technologies”. In: *Microelectronics Reliability* 58 (2016). Reliability Issues in Power Electronics, pp. 103–112. ISSN: 0026-2714. DOI: <https://doi.org/10.1016/j.microrel.2015.12.020>. URL: <https://www.sciencedirect.com/science/article/pii/S0026271415302687>.
- [198] N. G. N. Irias et al. “Challenges in using the physics-of-failure approach in practical applications”. In: *2017 IEEE Industry Applications Society Annual Meeting*. 2017, pp. 1–8. DOI: 10.1109/IAS.2017.8101838.
- [199] Nathan Pallo et al. “Short-Circuit Fault Ride-Through of Flying-Capacitor Multilevel Converters through Rapid Fault Detection and Idle-mode Operation”. In: *2020 IEEE 21st Workshop on Control and Modeling for Power Electronics (COMPEL)*. 2020, pp. 1–8. DOI: 10.1109/COMPEL49091.2020.9265743.
- [200] Mads Graungaard Taul et al. “Theoretical Analysis and Experimental Validation of Flying-Capacitor Multilevel Converters under Short-Circuit Fault Conditions”. In: *IEEE Transactions on Power Electronics* (2021), pp. 1–1. DOI: 10.1109/TPEL.2021.3075447.
- [201] Steven Thielemans et al. “Improved Natural Balancing With Modified Phase-Shifted PWM for Single-Leg Five-Level Flying-Capacitor Converters”. In: *IEEE Transactions on Power Electronics* 27.4 (2012), pp. 1658–1667. DOI: 10.1109/TPEL.2011.2169993.
- [202] Hossein Sepahvand et al. “Start-up Procedure and Switching Loss Reduction for a Single-Phase Flying Capacitor Active Rectifier”. In: *IEEE Transactions on Industrial Electronics* 60.9 (2013), pp. 3699–3710. DOI: 10.1109/TIE.2012.2204709.
- [203] Derek Chou. “Compact and Efficient Power Electronics System Design for Automotive, Solar, and Aerospace Applications”. Ph.D. dissertation. Berkeley, CA: University of California, Berkeley, May 2021.

- [204] Brendan Peter McGrath and Donald Grahame Holmes. “Natural Capacitor Voltage Balancing for a Flying Capacitor Converter Induction Motor Drive”. In: *IEEE Transactions on Power Electronics* 24.6 (2009), pp. 1554–1561. DOI: 10.1109/TPEL.2009.2016567.
- [205] Brendan Peter McGrath and Donald Grahame Holmes. “Enhanced Voltage Balancing of a Flying Capacitor Multilevel Converter Using Phase Disposition (PD) Modulation”. In: *IEEE Transactions on Power Electronics* 26.7 (2011), pp. 1933–1942. DOI: 10.1109/TPEL.2010.2097279.
- [206] C. Turpin et al. “Fault management of multicell converters”. In: *IEEE Transactions on Industrial Electronics* 49.5 (2002), pp. 988–997. DOI: 10.1109/TIE.2002.803196.
- [207] Alian Chen et al. “Fault-tolerant design for flying capacitor multilevel inverters”. In: *2009 IEEE 6th International Power Electronics and Motion Control Conference*. 2009, pp. 1460–1464. DOI: 10.1109/IPEMC.2009.5157616.
- [208] Hao Wang et al. “A Short-Circuit Fault-Tolerant Strategy for Three-Phase Four-Wire Flying Capacitor Three-Level Inverters”. In: *2019 IEEE 10th International Symposium on Power Electronics for Distributed Generation Systems (PEDG)*. 2019, pp. 781–786. DOI: 10.1109/PEDG.2019.8807759.
- [209] Xiaomin Kou, K.A. Corzine, and Y.L. Familiant. “A unique fault-tolerant design for flying capacitor multilevel inverter”. In: *IEEE Transactions on Power Electronics* 19.4 (2004), pp. 979–987. DOI: 10.1109/TPEL.2004.830037.
- [210] Jalal Amini and Mehrdad Moallem. “A Fault-Diagnosis and Fault-Tolerant Control Scheme for Flying Capacitor Multilevel Inverters”. In: *IEEE Transactions on Industrial Electronics* 64.3 (2017), pp. 1818–1826. DOI: 10.1109/TIE.2016.2624722.
- [211] Joachim Druant et al. “Adding Inverter Fault Detection to Model-Based Predictive Control for Flying-Capacitor Inverters”. In: *IEEE Transactions on Industrial Electronics* 62.4 (2015), pp. 2054–2063. DOI: 10.1109/TIE.2014.2354591.
- [212] R.L. de Araujo Ribeiro et al. “Fault-tolerant voltage-fed PWM inverter AC motor drive systems”. In: *IEEE Transactions on Industrial Electronics* 51.2 (2004), pp. 439–446. DOI: 10.1109/TIE.2004.825284.
- [213] Jalal Amini. “An Effortless Space-Vector-Based Modulation for N -level Flying Capacitor Multilevel Inverter With Capacitor Voltage Balancing Capability”. In: *IEEE Transactions on Power Electronics* 29.11 (2014), pp. 6188–6195. DOI: 10.1109/TPEL.2014.2299716.
- [214] Amer M. Y. M. Ghias et al. “Voltage Balancing Method for a Flying Capacitor Multilevel Converter Using Phase Disposition PWM”. In: *IEEE Transactions on Industrial Electronics* 61.12 (2014), pp. 6538–6546. DOI: 10.1109/TIE.2014.2320216.
- [215] Amer M. Y. M. Ghias et al. “Single-Carrier Phase-Disposition PWM Implementation for Multilevel Flying Capacitor Converters”. In: *IEEE Transactions on Power Electronics* 30.10 (2015), pp. 5376–5380. DOI: 10.1109/TPEL.2015.2427201.

- [216] Yong-Long Syu et al. “Phase-Shifted PWM with Dynamic Phase Shift Control and Zero Sequence Injection to Minimize Inductor Current Ripple in Three-Phase Flying Capacitor Multilevel Converters”. In: *2020 IEEE 21st Workshop on Control and Modeling for Power Electronics (COMPEL)*. 2020, pp. 1–7. DOI: 10.1109/COMPEL49091.2020.9265783.
- [217] Shanshan Zhao et al. “Submodule Fault Diagnosis Based on Capacitor Voltage Sliding-Mode Observer and Fast Diagnosis Criterion for Modular Multilevel Converter”. In: *2018 IEEE Energy Conversion Congress and Exposition (ECCE)*. 2018, pp. 5725–5730. DOI: 10.1109/ECCE.2018.8558062.
- [218] Serkan Kiranyaz et al. “Real-Time Fault Detection and Identification for MMC Using 1-D Convolutional Neural Networks”. In: *IEEE Transactions on Industrial Electronics* 66.11 (2019), pp. 8760–8771. DOI: 10.1109/TIE.2018.2833045.
- [219] Shu Ye et al. “Fault Diagnosis and Tolerance Control of Five-Level Nested NPP Converter Using Wavelet Packet and LSTM”. In: *IEEE Transactions on Power Electronics* 35.2 (2020), pp. 1907–1921. DOI: 10.1109/TPEL.2019.2921677.
- [220] Mojtaba Ashourloo et al. “Fault Detection in a Hybrid Dickson DC–DC Converter for 48-V Automotive Applications”. In: *IEEE Transactions on Power Electronics* 36.4 (2021), pp. 4254–4268. DOI: 10.1109/TPEL.2020.3022764.
- [221] Tianzhen Wang et al. “Cascaded H-Bridge Multilevel Inverter System Fault Diagnosis Using a PCA and Multiclass Relevance Vector Machine Approach”. In: *IEEE Transactions on Power Electronics* 30.12 (2015), pp. 7006–7018. DOI: 10.1109/TPEL.2015.2393373.
- [222] Mahdi Houchati et al. “Fault detection in modular multilevel converter using principle component analysis”. In: *2018 IEEE 12th International Conference on Compatibility, Power Electronics and Power Engineering (CPE-POWERENG 2018)*. 2018, pp. 1–6. DOI: 10.1109/CPE.2018.8372596.
- [223] Majid T. Fard et al. “Fast online diagnosis of open-circuit switching faults in flying capacitor multilevel inverters”. In: *Chinese Journal of Electrical Engineering* 6.4 (2020), pp. 53–62. DOI: 10.23919/CJEE.2020.000030.
- [224] Rajiv K. Varma et al. “Technique for Fast Detection of Short Circuit Current in PV Distributed Generator”. In: *IEEE Power and Energy Technology Systems Journal* 3.4 (2016), pp. 155–165. DOI: 10.1109/JPETS.2016.2592465.
- [225] Tomislav Dragičević. “Model Predictive Control of Power Converters for Robust and Fast Operation of AC Microgrids”. In: *IEEE Transactions on Power Electronics* 33.7 (2018), pp. 6304–6317. DOI: 10.1109/TPEL.2017.2744986.
- [226] Yuanshan Chen et al. “Mechanical Design Considerations of an “Ironless,” High-Specific-Power Electric Machine”. In: *IEEE Transactions on Transportation Electrification* 3.4 (2017), pp. 855–863. DOI: 10.1109/TTE.2017.2733763.

- [227] Andy Yoon et al. “High-Frequency Electric Machines for Boundary Layer Ingestion Fan Propulsor”. In: *IEEE Transactions on Energy Conversion* 34.4 (2019), pp. 2189–2197. DOI: 10.1109/TEC.2019.2942775.
- [228] Mattia Guacci, Dominik Bortis, and Johann W. Kolar. “High-Efficiency Weight-Optimized Fault-Tolerant Modular Multi-Cell Three-Phase GaN Inverter for Next Generation Aerospace Applications”. In: *2018 IEEE Energy Conversion Congress and Exposition (ECCE)*. IEEE, Sept. 2018, pp. 1334–1341. DOI: 10.1109/ECCE.2018.8557754.
- [229] J. V. M. Farias et al. “Reliability-Oriented Design of Modular Multilevel Converters for Medium-Voltage STATCOM”. In: *IEEE Transactions on Industrial Electronics* 67.8 (2020), pp. 6206–6214. DOI: 10.1109/TIE.2019.2937050.
- [230] J. Xu, P. Zhao, and C. Zhao. “Reliability Analysis and Redundancy Configuration of MMC With Hybrid Submodule Topologies”. In: *IEEE Transactions on Power Electronics* 31.4 (2016), pp. 2720–2729. DOI: 10.1109/TPEL.2015.2444877.
- [231] P. Lezana et al. “Survey on Fault Operation on Multilevel Inverters”. In: *IEEE Transactions on Industrial Electronics* 57.7 (2010), pp. 2207–2218. DOI: 10.1109/TIE.2009.2032194.
- [232] R. Hou et al. “A Simple Desaturation-Based Protection Circuit for GaN HEMT With Ultrafast Response”. In: *IEEE Transactions on Power Electronics* 36.6 (2021), pp. 6978–6987. DOI: 10.1109/TPEL.2020.3040727.
- [233] Richard W. Jones, Christopher Grice, and Robert Clements. “Safety and Reliability Driven Design for the More Electric Aircraft”. In: *2019 14th IEEE Conference on Industrial Electronics and Applications (ICIEA)*. 2019, pp. 949–954. DOI: 10.1109/ICIEA.2019.8834038.
- [234] International Air Transport Association. *Maintenance Costs for Aging Aircraft*. First. 800 Place Victoria, Montreal, Quebec, 2018. ISBN: 978-92-9229-637-7.
- [235] *Electric components - Reliability - Reference conditions for failure rates and stress models for conversion*. International Electrotechnical Commission, Oct. 2019.
- [236] W. Minford. “Accelerated Life Testing and Reliability of High K Multilayer Ceramic Capacitors”. In: *IEEE Transactions on Components, Hybrids, and Manufacturing Technology* 5.3 (1982), pp. 297–300. DOI: 10.1109/TCHMT.1982.1135974.
- [237] David Donhuang Liu. “Highly Accelerated Life Stress Testing (HALST) of Base-Metal Electrode Multilayer Ceramic Capacitors”. In: *Capacitor and Resistors Technology Symposium (CARTS) International*. Electronics Components Industry Association. Mar. 2013, pp. 235–248.
- [238] Qianwen Xu et al. “Systematic Reliability Modeling and Evaluation for On-Board Power Systems of More Electric Aircrafts”. In: *IEEE Transactions on Power Systems* 34.4 (July 2019), pp. 3264–3273. DOI: 10.1109/tpwrs.2019.2896454.

- [239] A. Goel and R. J. Graves. “Electronic system reliability: collating prediction models”. In: *IEEE Transactions on Device and Materials Reliability* 6.2 (2006), pp. 258–265. DOI: 10.1109/TDMR.2006.876570.
- [240] C. Jais, B. Werner, and D. Das. “Reliability predictions - continued reliance on a misleading approach”. In: *2013 Proceedings Annual Reliability and Maintainability Symposium (RAMS)*. 2013, pp. 1–6. DOI: 10.1109/RAMS.2013.6517751.
- [241] Stefan Biser et al. “Design Space Exploration Study and Optimization of a Distributed Turbo-Electric Propulsion System for a Regional Passenger Aircraft”. In: *AIAA Propulsion and Energy 2020 Forum*. 2020. DOI: 10.2514/6.2020-3592. eprint: <https://arc.aiaa.org/doi/pdf/10.2514/6.2020-3592>. URL: <https://arc.aiaa.org/doi/abs/10.2514/6.2020-3592>.
- [242] Mayank V. Bendarkar et al. “Evaluation of Off-Nominal Performance and Reliability of a Distributed Electric Propulsion Aircraft during Early Design”. In: *AIAA Scitech 2021 Forum*. 2021. DOI: 10.2514/6.2021-1723. eprint: <https://arc.aiaa.org/doi/pdf/10.2514/6.2021-1723>. URL: <https://arc.aiaa.org/doi/abs/10.2514/6.2021-1723>.
- [243] Igor Bolvashenkov and Hans-Georg Herzog. “Degree of Fault Tolerance as a Comprehensive Parameter for Reliability Evaluation of Fault Tolerant Electric Traction Drives”. In: *Transactions on Environment and Electrical Engineering* 1.3 (2016), pp. 50–61. ISSN: 2450-5730. DOI: 10.22149/teee.v1i3.39. URL: <https://teee.eu/index.php/TEEE/article/view/39>.
- [244] Saakar Byahut and Alejandra Uranga. “Propulsion System Component Modeling for an All-Electric Commuter Aircraft Mission”. In: *AIAA Scitech 2020 Forum*. 2020. DOI: 10.2514/6.2020-0015. eprint: <https://arc.aiaa.org/doi/pdf/10.2514/6.2020-0015>. URL: <https://arc.aiaa.org/doi/abs/10.2514/6.2020-0015>.
- [245] Mattia Guacci, Dominik Bortis, and Johann W. Kolar. “On-state voltage measurement of fast switching power semiconductors”. In: *CPSS Transactions on Power Electronics and Applications* 3.2 (2018), pp. 163–176. DOI: 10.24295/CPSSSTPEA.2018.00016.
- [246] Robert Murphree et al. “A Flip-Chip Capable Low-Side and High-Side SOI Gate Driver with Variable Drive Strength for GaN Power FETs”. In: *2019 IEEE 7th Workshop on Wide Bandgap Power Devices and Applications (WiPDA)*. 2019, pp. 113–120. DOI: 10.1109/WiPDA46397.2019.8998884.
- [247] M A H Broadmeadow, G R Walker, and S M Bulmer. “An FPGA-based, single transducer, active balancing scheme for a 5-level flying capacitor converter with 1 MHz effective switching frequency”. In: *8th IET International Conference on Power Electronics, Machines and Drives (PEMD 2016)*. 2016, pp. 1–6. DOI: 10.1049/cp.2016.0188.

- [248] Ziyu Xia et al. “State Space Analysis of Flying Capacitor Multilevel DC-DC Converters for Capacitor Voltage Estimation”. In: *2019 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2019, pp. 50–57. DOI: 10.1109/APEC.2019.8722230.
- [249] Pascal S. Niklaus, Dominik Bortis, and Johann W. Kolar. “High-Bandwidth High-CMRR Current Measurement for a 4.8 MHz Multi-Level GaN Inverter AC Power Source”. In: *2021 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2021, pp. 200–207. DOI: 10.1109/APEC42165.2021.9487044.
- [250] Kieran T. Levin, Eric M. Hope, and Alejandro D. Domínguez-García. “Observer-based fault diagnosis of power electronics systems”. In: *2010 IEEE Energy Conversion Congress and Exposition*. 2010, pp. 4434–4440. DOI: 10.1109/ECCE.2010.5618444.
- [251] Jason Poon et al. “A linear-switched observer for large-signal state estimation in power electronics”. In: *2012 15th International Power Electronics and Motion Control Conference (EPE/PEMC)*. 2012, LS3b.3-1-LS3b.3-5. DOI: 10.1109/EPEPEMC.2012.6397429.
- [252] Xiangyu Ding et al. “Fault Detection and Isolation Filters for Three-Phase AC-DC Power Electronics Systems”. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 60.4 (2013), pp. 1038–1051. DOI: 10.1109/TCSI.2012.2221222.
- [253] Jason Poon et al. “Model-Based Fault Detection and Identification for Switching Power Converters”. In: *IEEE Transactions on Power Electronics* 32.2 (2017), pp. 1419–1430. DOI: 10.1109/TPEL.2016.2541342.
- [254] Pamela J. Tannous et al. “Dynamic temperature estimation of power electronics systems”. In: *2017 American Control Conference (ACC)*. 2017, pp. 3463–3469. DOI: 10.23919/ACC.2017.7963482.
- [255] Yi Dou, Ziwei Ouyang, and Michael A. E. Andersen. “Integrated Coupled Inductors With Functionality of Current Balancing Transformer for Two-Phase Synchronous DC-DC Converters”. In: *IEEE Transactions on Power Electronics* 35.5 (May 2020), pp. 4472–4476. DOI: 10.1109/TPEL.2019.2948882.
- [256] Gabriel J. Capella et al. “Current-Balancing Technique for Interleaved Voltage Source Inverters With Magnetically Coupled Legs Connected in Parallel”. In: *IEEE Transactions on Industrial Electronics* 62.3 (2015), pp. 1335–1344. DOI: 10.1109/TIE.2014.2345345.
- [257] Ruirui Chen et al. “Coupled Inductor Design for Interleaved Three-level Active Neutral Point Clamped Inverters Considering EMI Noise Reduction”. In: *2019 IEEE Applied Power Electronics Conference and Exposition (APEC)*. 2019, pp. 257–264. DOI: 10.1109/APEC.2019.8721817.