CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

WATERMARK

PHD THESIS



Keeping batteries “just right”

Do you remember the childhood story of Goldilocks and the Three Bears, where Goldilocks was always looking for something that was neither too hot nor too cold, but “just right”?

A lithium-ion battery is not so different. It also has a preferred temperature range in which it performs well and ages slowly. If it becomes too cold, less ‘juice’ can be extracted, and charging can accelerate ageing. If it becomes too hot, it degrades faster and, in extreme cases, can become a safety risk.

Keeping the average temperature within a suitable range is not enough. Different regions inside a battery can reach different temperatures, creating hot and cold spots that can affect performance and ageing. The challenge is thus not only to keep the battery at the right temperature but also to maintain a reasonably uniform temperature. This thesis investigates how that can be achieved by combining cooling through the battery surface and its electrical terminals. These two types of cooling have different strengths, and instead of choosing only one, a controller decides how the available coolant should be shared between them to harness their individual strengths as the battery operating conditions change. Fast mathematical models are developed to estimate the temperature distribution inside the battery, while smart control is used to adjust the cooling in real time. The work also considers battery ageing directly.

The results show that making a battery as cold and thermally uniform as possible does not always minimise degradation. What matters is when, where, and how much cooling is applied. In short, the goal is to give the battery what Goldilocks wanted: the right conditions at the right time, helping it perform well, remain safe, and last longer.

WATERMARK

CHALMERS
UNIVERSITY OF TECHNOLOGY

GODWIN KWADWO PEPRAH • Modelling and real-time predictive control for hybrid cooling of lithium-ion batteries • 2026

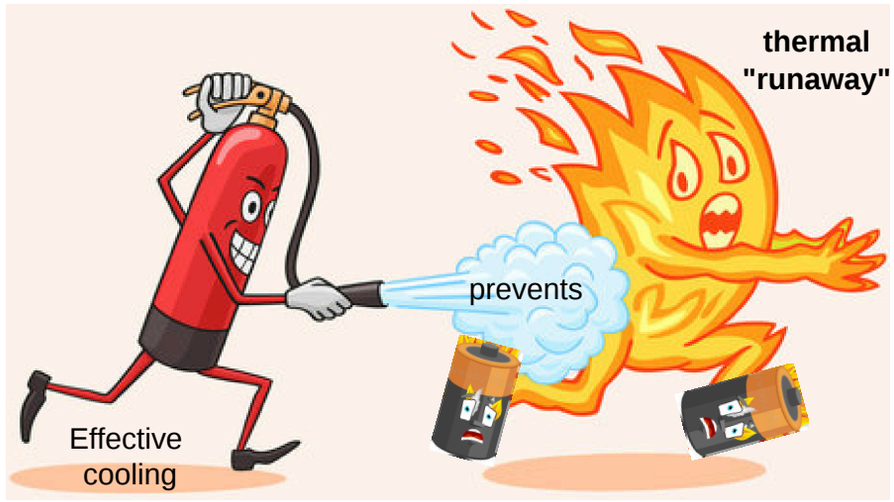

Modelling and real-time predictive control for hybrid cooling of lithium-ion batteries

For enhanced thermal safety, performance, and lifetime

GODWIN KWADWO PEPRAH

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Modelling and real-time predictive control for hybrid cooling of lithium-ion batteries

For enhanced thermal safety, performance, and lifetime

Godwin Kwadwo Peprah

Department of Electrical Engineering
Chalmers University of Technology
Gothenburg, Sweden, 2026

Modelling and real-time predictive control for hybrid cooling of lithium-ion batteries

For enhanced thermal safety, performance, and lifetime

Godwin Kwadwo Peprah
ISBN 978-91-8103-510-0

Acknowledgements, dedications, and similar personal statements in this thesis reflect the author's own views.

© Godwin Kwadwo Peprah 2026 except where otherwise stated.

Selected material from the author's licentiate thesis: Godwin K. Peprah, "Thermal modelling and control of lithium-ion batteries: For enhanced thermal safety and lifetime," Licentiate thesis, Chalmers University of Technology, Gothenburg, Sweden, 2024., is republished in this Ph.D. thesis.

Doktorsavhandlingar vid Chalmers tekniska högskola
Ny serie nr 5967
ISSN 0346-718X
doi:10.63959/chalmers.dt/5967

Department of Electrical Engineering
Chalmers University of Technology
SE-412 96 Gothenburg, Sweden
Phone: +46 (0)31 772 1000

Cover:
An illustration of battery cooling mitigating thermal runaway risk, reflecting the central theme of this thesis: modelling and predictive cooling control for enhanced battery safety, performance, and lifetime.
Adapted from "Fire extinguisher" by strels, Adobe Stock, asset no. 82202940.

Printed by Chalmers Digital Printing
Gothenburg, Sweden, November 2026

Modelling and real-time predictive control for hybrid cooling of lithium-ion batteries

For enhanced thermal safety, performance, and lifetime

Godwin Kwadwo Peprah
Department of Electrical Engineering
Chalmers University of Technology

Abstract

Lateral surface cooling can limit temperature rise in lithium-ion batteries through its large heat-transfer area, whereas tab cooling can reduce internal thermal gradients by exploiting the high thermal conductivity of the current collectors. This thesis investigates how these complementary characteristics can be integrated through modelling and predictive control.

A control-oriented thermal modelling framework is developed using the Chebyshev spectral-Galerkin method and model component decomposition. The resulting reduced-order two-dimensional models enable independent manipulation of thermal boundary conditions at the tabs and lateral surface while predicting spatial temperature distributions at low computational cost. Building on this framework, tab and surface cooling are integrated via a constrained optimal coolant allocation problem in which coolant is dynamically distributed among the cooling interfaces. A first-principles coolant model and valve-actuation model are coupled with the electrothermal dynamics. The resulting nonlinear problem is solved using real-time iteration model predictive control (MPC), which closely reproduces nonlinear MPC while reducing mean computation time from several seconds to approximately 19 ms. The framework is extended to degradation-aware control through a coupled distributed electrothermal–degradation model that captures spatial variations in temperature, current density, state of charge, and capacity loss. Model evaluations reveal increasingly pronounced electrothermal and degradation non-uniformities in larger cells.

Across the control studies, integrated tab-surface cooling provides the best overall trade-off between thermal regulation, uniformity, degradation, and coolant utilisation compared to conventional cooling configurations. Together, these contributions establish a basis for hybrid battery cooling combining spatially resolved modelling with real-time predictive control in battery systems.

Keywords: Lithium-ion batteries, battery thermal management, integrated tab and surface cooling, control-oriented thermal modelling, distributed electrothermal modelling, battery degradation, coolant modelling and allocation, model predictive control, real-time control.

To my family.

List of Publications

This thesis is based on the following appended publications:

[A] **Godwin K. Peprah**, Yicun Huang, Torsten Wik, Faisal Altaf, Changfu Zou, "Thermal Modeling of Battery Cells for Optimal Tab and Surface Cooling Control". Published in *IEEE Transactions on Control Systems Technology*, vol. 34, no. 4, pp. 1831–1844, July 2026, doi: 10.1109/TCST.2026.3675726.

[B] **Godwin K. Peprah**, Torsten Wik, Masood Tamadondar, Changfu Zou, "Model Predictive Coolant Allocation for Integrated Tab–Surface Cooling of Battery Cells". Submitted to *IEEE Transactions on Control Systems Technology* for review, 2026. Available as *arXiv preprint: 2607.10872*.

[C] **Godwin K. Peprah**, Phil Kollmeyer, Torsten Wik, Changfu Zou, "Degradation-Aware Model Predictive Coolant Allocation for Integrated Tab-Surface Cooling of Battery Cells". Manuscript, To be submitted to *IEEE Transactions on Control Systems Technology*.

Other publications related to the thesis, not appended, are:

[D] **Godwin K. Peprah**, Yicun Huang, Torsten Wik, Faisal Altaf, Changfu Zou, "Control-Oriented 2D Thermal Modelling of Cylindrical Battery Cells for Optimal Tab and Surface Cooling". Published in *2024 American Control Conference (ACC)*, Toronto, ON, Canada, 2024, pp. 4651-4656, doi: 10.23919/ACC60939.2024.10644585.

[E] **Godwin K. Peprah**, Yicun Huang, Torsten Wik, Faisal Altaf, Changfu Zou, "Model Predictive Cooling Control of Cylindrical Battery Cells Through Tab and Surface Channels". Published in *2025 American Control Conference (ACC)*, Denver, CO, USA, 2025, pp. 4071-4076, doi: 10.23919/ACC63710.2025.11107872.

[F] **Godwin K. Peprah**, Francesco Liberati, Faisal Altaf, Gilbert Osei-Dadzie, Alessandro Di Giorgio, Antonio Pietrabissa, "Optimal load sharing in reconfigurable battery systems using an improved model predictive control method". Published in *29th Mediterranean Conf. on Control and Automation (MED)*, Puglia, Italy, 2021, pp. 979-984, doi: 10.1109/MED51440.2021.9480237.

Contents

Abstract iii

List of Publications vii

Acknowledgements xi

Acronyms xix

I Overview 1

1 Introduction 3
- 1.1 Motivation . . . 3
- 1.2 Research questions . . . 7
- 1.3 Thesis contributions . . . 8
- 1.4 Thesis outline . . . 10

2 Battery thermal management 11
- 2.1 Lithium-ion batteries and temperature effects . . . 11
- 2.2 Battery thermal management systems . . . 13
  - Integrated tab-surface cooling . . . 13

3 Battery and cooling system modelling **17**
3.1 Battery model 17
Electrical modelling 17
Thermal modelling 21
Degradation modelling 25
3.2 Cooling system model 30
Coolant modelling 31
Valve actuation modelling 32
3.3 Coupled battery–cooling system 33

4 Predictive battery cooling control **37**
4.1 Optimal coolant allocation 37
4.2 Model predictive control 40
Real-time iteration model predictive control 41

5 Summary of included papers **45**
5.1 Paper A 45
5.2 Paper B 47
5.3 Paper C 48

6 Concluding remarks and future work **51**
6.1 Conclusions 51
6.2 Synthesis and broader implications 55
6.3 Future work 55

References **59**

II Papers **65**

A Thermal modelling **A1**

B Real-time coolant allocation **B1**

C Real-time degradation-aware coolant allocation **C1**

Acknowledgements

A ship is safe in harbour, but that is not what ships are built for.

Those words gave me courage as I sat aboard a Turkish Airlines flight bound for Rome, leaving behind the familiarity of home in pursuit of higher education. I remember the excitement, the uncertainty, and the many questions running through my mind. I was venturing into unfamiliar waters, with no guarantee of what awaited me. All I knew was that I wanted to learn, grow, and discover how far my ambitions could take me. Looking back, that decision set in motion a journey far more rewarding than I could ever have imagined.

Italy was incredibly good to me. I left not only with a Master's degree but also with friendships and an extended family who have shaped my life in countless ways. It would be a disservice to my PhD story not to acknowledge Francesco Liberati, who supervised my Master's thesis and guided me towards my very first scientific publication. I have also cherished our coffee catch-ups at Gatsby near Vittorio Emanuele whenever my travels have brought me back to the Eternal City. That first research experience opened my eyes to the world of academia. Combined with the inspiration I drew from close friends such as James Adu-Amankwah and Eric Owusu-Obeng, who had already pursued PhDs, it became increasingly difficult to imagine taking any path other than a PhD to deepen my understanding in battery technology. There is something powerful about surrounding yourself with people whose ambitions challenge you to dream bigger. I was fortunate to find such giants early in my journey.

Speaking of giants, I fear the ink in my pen would run dry before I could adequately describe the influence Faisal Altaf has had on my academic and professional development. From my Master's thesis in Rome through to the completion of this PhD, Faisal has been a mentor and a source of inspiration. Thank you for seeing possibilities in me, for opening doors, and for continuing to invest in my growth. I am eternally grateful, and I hope to extend the same generosity to others as my career progresses.

It is often said that a PhD supervisor can make or break a doctoral student's experience. I was blessed with two who helped me stand tall. To my main supervisor, Professor Changfu Zou, thank you for your unwavering commitment to excellence and for pushing me to become a better researcher. You were always ready, armed with your Adobe PDF annotation tools, to ensure that our journal manuscripts were nothing short of outstanding. I still remember the bittersweet feeling of opening yet another file ending in `_cz_comments.pdf`

during the preparation of my first journal paper. After more than 25 rounds of revisions, I occasionally wondered whether we would ever reach the finish line. But hindsight is indeed 20/20. We produced a paper we could be proud of, and I emerged from that experience a far better researcher and scientific writer. I will never take for granted the many occasions you went above and beyond, including late evenings and weekends spent helping us prepare manuscripts for submission. Thank you for your patience, high standards, encouragement, and belief in my potential. And lest I forget, thank you for introducing me to that Chinese restaurant in Hisingen. That recommendation has been a game-changer!

To Professor Torsten Wik, thank you for bringing warmth, wisdom, and a distinctly Swedish *lagom* approach to supervision. Your calm, easy-going personality was one of the deciding factors in my choice of Chalmers over TU Wien for my doctoral studies. Although your supervision style was often wonderfully relaxed, you never missed an opportunity to decorate my manuscripts with red ink whenever the occasion demanded it. I learned to appreciate both the freedom you gave me and the precision you expected in return. Changfu and Torsten, your open-door policy meant more to me than you may realise. Being able to walk into your offices without an appointment, ask questions, debate ideas, or simply seek advice created an environment in which I felt genuinely supported. Thank you both for the opportunities to lecture, attend international conferences, and build professional networks. I could not have asked for a better supervisory team.

I am equally grateful to my co-authors and collaborators, particularly Faisal Altaf, Yicun Huang, Masood Tamadondar, and Phil Kollmeyer, for the ideas, discussions, and contributions that helped bring this work to life. My special thanks go to Dr Phil Kollmeyer, who welcomed me to McMaster University in Canada with open arms. Your hospitality made Canada feel less like a temporary research destination and more like a second home. I also extend my appreciation to the professors who contributed to my doctoral development, especially Professor Anders Grauers, my examiner, for his encouragement and support throughout this journey.

I gratefully acknowledge the financial support of the Swedish Electromobility Centre, the Swedish Energy Agency, the Swedish Research Council, and the European Union's Horizon Europe programme through the Marie Skłodowska-Curie Actions. Their support provided the resources and oppor-

tunities that enabled this research, including international collaboration and research exchange.

Of course, a PhD is never only about research. It is also about the people who make the years memorable, the friendships that carry you through difficult moments, and the laughter that reminds you there is life beyond MATLAB, manuscripts, and deadlines.

To my *Gothenburg fellas*, haaa where do I even begin? Your Highness DJ Andyy, thank you for keeping our feet moving on the dance floors of Third Place, Bivac, and Draken. I have always enjoyed our battery conversations, although I am still not sure whether they qualify as work or leisure. I hope to make you proud by finally putting my own DJ set to good use, rather than waiting forever to hear you play my favourite, *Midnight Prayer*. And, as always, I look forward to your next AI model recommendation. KG, Johannesburg's finest gentleman, thank you for your sharp wit, incredible dance moves, and stocks investment tips. You have made life in Gothenburg infinitely more entertaining. King Pharaoh Wael Zozo, thank you for always bringing the heat and showing us how things are done in the business world, Habibi. Cate, thank you for being the supportive big sister I never knew I needed. I have always enjoyed your witty captions, admired your advocacy for women, and appreciated your passion for African growth and prosperity. Newton, thank you for setting the pace with your PhD and showing me that dreams do come true. Damian, thank you for bringing a little Ohio and New York flavour to our gatherings. Tove, thank you for bringing the Swedish touch and for your thoughtful, reflective questions, which have often made our conversations richer. Sammy and Felicia, thank you for welcoming me to South Africa for your wedding and reminding me of the beauty of friendship, love, and family. Amir, our resident aviation and techno enthusiast, thank you for your generosity and wonderful friendship. Here's to more brunches, Yakida nights, finally making our gym plans a reality, and, of course, sharing that fine bottle of Scotch whisky you gifted me. I'm saving it for a special occasion with you. Admire, your bubbly personality and infectious energy always bring life to our gatherings. Thank you for always making time to catch up whenever I'm in Stockholm. I've enjoyed our conversations about startups and business ideas, and I look forward to seeing some of those ideas come to life someday. Salome, thank you for helping me settle into Gothenburg during my early days. Having someone who understood the joys, frustrations, and

uncertainties of PhD life made those first steps much easier.

To my *aşkım*, Dilan, thank you for being a beautiful part of this journey. Your smile, your simple yet thoughtful "How's your day?", and your unwavering support have brought comfort and happiness into my life in more ways than I can express. Through the demanding days, endless writing sessions, and moments when everything seemed to be happening at once, you reminded me to appreciate the little things. I have always admired the dedication and care you bring to everything you do, no matter how big or small, and that commitment to giving your best continually rubs off on me to do the same. *Teşekkür ederim, aşkım.* I look forward to making many more beautiful memories together, sharing plenty more delicious *kaz eti* from Kars with the family and Beni.

Alexiaaa Boss! Little did I know when I arrived at Sapienza that I would leave not only with a degree but also with a best friend for life. You have been there through my highest highs and lowest lows, listening patiently and without judgement, well, perhaps with just a little judgement when the situation called for it. You have celebrated my victories as though they were your own and stood beside me when things were not going according to plan. I was deeply honoured when you asked me to be your best man. Our friendship has always been a two-way street, built on trust, honesty, laughter, and the willingness to show up for each other when it matters most. Thank you for being someone I can count on without hesitation, for telling me what I need to hear rather than simply what I want to hear, and for reminding me that distance means very little when a friendship is genuine. I also greatly appreciated you travelling from the Netherlands with Naa to celebrate my birthday. Gestures like that mean the world to me. Here's to many more milestones, adventures, and stories we will probably laugh about for years to come.

Miiister Skeegro! I have rarely encountered anyone as genuinely selfless as Albert. My memories of you go back to my very first day at Chalmers, in room 5343. True to form, you enthusiastically explained the structure of Chalmers, the Swedish study periods, and the courses I should consider taking. Even then, I could tell that you were someone who genuinely wanted others to succeed. Over the years, your knowledge across an astonishing range of subjects has enriched my life tremendously. You have taught me that being a classic gentleman involves much more than academic achievements: skincare, fragrances, personal style, and, most importantly, dressing the way one wishes

to be addressed. You have been a dependable friend, a generous mentor, and a brother from a Croatian mother. *Hvala ti, brate!* Look at us now. We managed to make our dream of defending our theses within weeks of each other and visiting Canada a reality. I look forward to many more adventures together, including levelling the score between Ghana and Croatia after that World Cup defeat and, perhaps an even greater challenge, levelling up your tolerance for properly spiced Ghanaian jollof!

Cheers to *La Squadra della Bella Vita*, my wonderful Italian crew who made Chalmers memorable: Francesco Gigante, Lorenzo, my chief MPC consultant, Sabino, Michela, Mattia, Martina, and Francesco Popolizio, who never misses an opportunity to remind the world how beautiful his hometown, Matera, is. I treasure our time at the CSAE conference in Shanghai, including that memorable DiDi ride to the Bund and time at INS. I wish you continued success in your PhD, *fra*! May the good life continue wherever our paths lead us.

Chalmers would also have been incomplete without the endless banter from Maharaja Anand, the Turkish delights from Sultan Ahmet, and the ever-cheerful now Professor Faris. The fifth floor became much more than a collection of offices. It was a place of friendship, support, spontaneous discussions, jokes, shared frustrations, and countless memories. To Gabriel, Kilian, Qingbo, Bethlehem, Yao, Nishant, Ying, Adam, Alvin, Amal, Benedick, Rikard, Esi, Xiaolei, Ilya, Yizhou, Huang, Carl-Johan, Sondre, Attila, Daniel, Lasse, Josef, Filip, Isac, Quentin, and Meng, thank you for being part of the community that made those years feel like home.

My gratitude extends well beyond Gothenburg. To my wonderful friends from Ghana—Dzovor, Cornelius, Greg, Buari and Eulalia, Baffour, Dabo, Ababio, Krevises, Anane, Atsino, Azu, Wilson, Agya, Ashbel, Derrick, Magdiel, Kankam, Nana Yaw, Flow, Bounty—thank you for keeping me connected to home, no matter how far away I have travelled. A special shout-out to Dzovor for having my back since our Olam days. One memory I'll always cherish is our road trip last year, when you rented a car and drove us all the way from New York to Miami. That trip will always be one for the books! To my friends from Italy—Any33mi Gerald, Nana Aryah, Collins, Claribel, and Dillion—thank you for the friendship and memories that continue long after our time together in Rome. And to my good friends in Canada—Beri-Winkles, Gigi, Dennis, Owura, Scrilla, Elikem, and Nanor—thank you for adding another

beautiful chapter to my journey. Beri-Winkles, appreciate you for the wonderful hospitality, the endless banter, and the memorable nights on Ossington and King Street with Dennis. I've enjoyed our business discussions, and I hope some of those ideas come to fruition soon. Here's to more drives in your Mercedes and fewer parking tickets! Although I'm convinced that if you'd let me take the wheel more often, we might have avoided a few of those tickets. Across continents and time zones, you have all reminded me that home is not simply a place but also the people who make you feel that you belong.

To my family, I dedicate this thesis. Long before I understood the meaning of research, academic qualifications, or doctoral degrees, you taught me the values that would ultimately carry me here: faith, discipline, hard work, humility, and perseverance. You have been my backbone through every stage of this journey, celebrating my successes, encouraging me through disappointments, and believing in me even when the road ahead was uncertain. Whatever I have achieved is as much a reflection of your love and sacrifices as it is of my own efforts.

To my dear mother, Florence, words will never adequately express how much your prayers, encouragement, and unconditional love have meant to me. Our Sunday evening calls have become one of the constants I treasure most, a reminder that regardless of where I am in the world, there is always someone at home cheering me on. Your faith in me has often been greater than my faith in myself, and your strength and sacrifices have shaped the person I am today. I hope this achievement brings you even a fraction of the joy and pride that you have given me throughout my life. I love you, Mum.

To my fathers of blessed memory, Emmanuel Attah Peprah Snr and Jnr, the seeds you sowed through your love, sacrifices, and guidance continue to bear fruit. Though you are no longer here to witness this milestone, your legacy lives on in the person I have become. This achievement is as much yours as it is mine. May your souls continue to rest in perfect peace.

To Sisi Dora, my siblings and their families, Sampep and Gilberto, Godwill and Elizabeth, Sista and AB, Adofoa, Gyemfua, Serwaa, and Frema, thank you for the love, support, encouragement, and sense of belonging that no distance could ever diminish. I am immensely grateful for each of you. To my uncle, Ayyy Kumah, thank you for being a father figure and a pillar of strength for our family. To Blessings, Elder Charles, Elder Steve, Ma Doris, Rev Akomea, thank you for your prayers and spiritual nourishment.

Gilberto, my brother, where do I even begin? I could write an entire thesis documenting everything you have done for me, and it still would not be enough. You have been a confidant, a best friend, a brother, a father figure, a code debugger, and a part-time PhD supervisor, all rolled into one. You have supported my ambitions, helped me navigate difficult decisions, listened when I needed someone to talk to, and never hesitated to lend a hand. Your presence has been one of the greatest blessings of my life. Thank you for being the kind of brother-in-law anyone would be fortunate to have. And here's to you embarking on your own PhD journey soon! I look forward to cheering you on, returning the favour as your part-time supervisor, and celebrating your own doctoral milestone someday.

And to all my wonderful nieces and nephews, who seem to have mastered the art of draining my bank account with remarkable efficiency, I love you all dearly! Your laughter, energy, and excitement are reminders of the joy that makes all the hard work worthwhile. Keep dreaming big. I look forward to watching each of you achieve wonderful things, although perhaps we can negotiate the financial arrangements along the way!

Above all, I give thanks to Almighty God, whose grace, mercy, and faithfulness have sustained me throughout this journey.

As a young boy growing up in Dansoman, Accra, I had a dream. I did not know precisely where that dream would take me or how many sacrifices, setbacks, opportunities, and extraordinary people I would encounter along the way. Years later, sitting on that flight to Rome, I was still uncertain of what lay ahead. But I chose to leave the safety of the harbour and set sail anyway. Today, looking back on the journey from Dansoman to Rome, Gothenburg, and beyond, I can say with a grateful heart that every uncertainty, every sacrifice, every challenge, and every leap of faith was worth it. The journey has given me far more than a doctoral degree. It has given me purpose, friendships that transcend borders, mentors who became family, and a deeper appreciation of what it means to dream and persevere.

Godwin Kwadwo Peprah
Gothenburg, November 2026

Acronyms

2D:	Two-dimensional
aTSC:	All-tab and surface cooling
BMS:	Battery management system
bTC:	Bottom-tab cooling
bTSC:	Bottom-tab and surface cooling
btTC:	Bottom- and top-tab cooling
BTMS:	Battery thermal management system
CSG:	Chebyshev spectral–Galerkin
DC:	Direct current
DOCAP:	Degradation-aware optimal coolant allocation problem
ECM:	Equivalent circuit model
EFC:	Equivalent full cycle
EoL:	End of life
ES:	Equal split
EV:	Electric vehicle
FE:	Finite element
FEM:	Finite element method
ITSC:	Integrated tab–surface cooling
LFP:	Lithium iron phosphate
Li:	Lithium
LiB:	Lithium-ion battery

MPC:	Model predictive control
mWLTP:	Modified WLTP
NMC:	Nickel manganese cobalt
NMPC:	Nonlinear model predictive control
OCAP:	Optimal coolant allocation problem
OCV:	Open-circuit voltage
ODE:	Ordinary differential equation
PDE:	Partial differential equation
PI:	Proportional–integral
PWM:	Pulse-width modulation
QP:	Quadratic programme
RC:	Resistor–capacitor
RMS:	Root mean square
RTI:	Real-time iteration
SC:	Surface cooling
SEI:	Solid electrolyte interphase
SoC:	State of charge
SoH:	State of health
TEC:	Thermal equivalent circuit
WLTP:	Worldwide harmonised light vehicles test procedure

Part I

Overview

CHAPTER 1

Introduction

1.1 Motivation

Sustainable electrification and battery temperature effects

Recent geopolitical tensions and disruptions to global energy supply have given renewed urgency to the transition towards electrification. Beyond its established role in reducing greenhouse-gas emissions, electrification can reduce dependence on imported fossil fuels and exposure to fuel price volatility, thereby contributing to improved energy security and resilience [1], [2].

The transport sector is particularly relevant in this context due to its strong dependence on fossil fuels. Electrification through electric vehicles (EVs) offers a means of reducing this dependence, with lithium-ion batteries (LiBs) remaining the dominant energy storage technology for electric mobility [3], owing to their high energy and power density, high efficiency, and comparatively long cycle life [4]. Despite these advantages, the performance, safety, and lifetime of LiBs are strongly influenced by temperature [5]. Elevated operating temperatures accelerate irreversible degradation and increase the risk of thermal runaway, whereas low temperatures reduce power capability,

charge acceptance, and can promote undesirable ageing mechanisms, particularly lithium plating during charging. Beyond the absolute temperature, spatial temperature variations or thermal gradients within a cell can induce non-uniform electrochemical reaction rates, localised ageing, and mechanical stress [6]. Effective battery thermal management must therefore address not only the overall cell temperature but also its spatial distribution.

Complementary tab and surface cooling

Battery thermal management systems (BTMSs) [7] employ different heat-transfer mechanisms and cooling interfaces to regulate battery thermal behaviour. Within liquid-cooled systems, heat can be removed through the lateral surface of a battery cell, through its electrical connection tabs or terminals, or through a combination of these interfaces [8]. As illustrated in Fig. 1.1, surface and tab cooling can result in distinctly different spatial temperature distributions owing to the anisotropic thermal conductivity of LiBs. Several studies [9], [10], [11], [12], [13] have investigated the thermal performance of these two cooling approaches. Lateral surface cooling generally benefits from a relatively large heat-transfer area and is effective at limiting the overall temperature rise. Tab cooling, in contrast, exploits the high thermal conductivity of the current collectors to promote a more uniform internal temperature distribution. The relative effectiveness of these cooling methods depends on the cell geometry, thermal properties, electrical loading, and imposed cooling conditions. Consequently, neither approach necessarily minimises both the overall temperature rise and internal thermal gradients simultaneously.

Although these complementary characteristics have been investigated extensively, their systematic exploitation through dynamic coordination of the cooling interfaces remains limited. Integrating tab and lateral surface cooling within a common thermal management framework introduces a coolant allocation problem. With a limited total coolant flow capacity, the appropriate distribution among the cooling interfaces depends on the electrical load, thermal state, coolant conditions, and thermal management objectives. Since these conditions evolve during battery operation, a fixed coolant distribution may not provide the desired thermal performance across different operating conditions. This motivates the development of a dynamic optimal coolant allocation strategy.

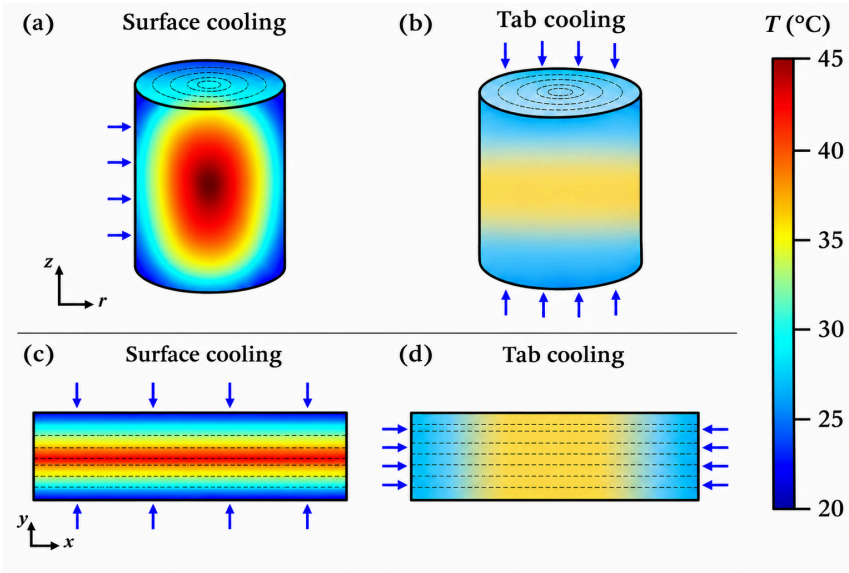

Figure 1.1: Illustrative temperature distributions under surface cooling: (a), (c) and tab cooling: (b), (d) for cylindrical and pouch cells. For the cylindrical cells, r and z denote the radial and axial directions, respectively, whereas for pouch cells, x and y denote the in-plane and through-plane directions, respectively.

Control-oriented thermal modelling

Addressing the coolant allocation problem requires models that provide sufficient information about the internal temperature distribution while remaining computationally tractable for real-time control. Lumped thermal equivalent circuit (TEC) models [14], [15], [16] are attractive for battery management because of their low computational complexity. However, they provide only a limited number of representative temperatures and do not explicitly reconstruct the spatial temperature field. Moreover, the assumptions underlying highly lumped thermal representations become increasingly restrictive when appreciable internal temperature gradients develop. High-fidelity first-principles thermal models based on partial differential equations (PDEs) [17] and numerical discretisation methods such as the finite element method (FEM) [18] can resolve the spatial thermal dynamics, but their computational

burden can be prohibitive for online optimisation. Reduced-order spatial thermal models [19], [20] have been developed to bridge this divide. However, existing formulations do not allow independent and targeted cooling control of battery thermal boundaries, limiting their direct suitability for control problems in which cooling must be routed between multiple battery surfaces. There is therefore a need for control-oriented thermal models that retain the relevant spatial information and facilitate independent actuation of the tab and lateral surface cooling interfaces without sacrificing computational efficiency.

Predictive coolant allocation

The control problem becomes more challenging when the dynamics of the cooling system itself are considered. Coolant temperature evolves as it exchanges heat with the battery, while physical actuators regulate the routing of coolant among the cooling channels. A control-oriented representation of the coupled battery and cooling system is therefore needed to establish the relationship between the manipulated coolant flow rates and the resulting spatial temperature distribution. Model predictive control (MPC) [16], [21], [22] is particularly suited to this task because it can anticipate future thermal loads, coordinate multiple manipulated inputs, and explicitly handle operating constraints.

However, allocating a finite coolant resource among multiple thermally coupled cooling interfaces differs from conventional predictive cooling problems involving a single cooling input. In particular, the coupling between coolant flow rates and thermal states introduces nonlinearities into the prediction model. Solving the resulting nonlinear optimisation problem repeatedly during battery operation can impose a substantial computational burden. This motivates the development of a computationally efficient predictive control formulation that explicitly accounts for the coupled battery, coolant, and actuator models while facilitating real-time coolant allocation.

Degradation-aware thermal control

Temperature regulation and thermal gradient reduction are commonly used as objectives for battery thermal control because both are closely associated with battery safety and health [23], [24], [25]. Nevertheless, battery degradation [26], [27], [28] is governed by a complex cocktail of interacting mechanisms

whose dependence on temperature, state of charge (SoC), current, and operating history can differ substantially. In addition to influencing local degradation rates directly, intra-cell temperature variations can induce spatial variations in electrical resistance, current distribution, SoC, and heat generation. These electrothermal interactions can cause different regions of the cell to age at different rates. The resulting non-uniform degradation can, in turn, influence the subsequent electrical and thermal behaviour of the cell. Conventional lumped electrical and degradation models cannot explicitly capture these spatial interactions. Moreover, owing to the nonlinear dependence of degradation rates on the local operating conditions, evaluating degradation using spatially averaged quantities is not generally equivalent to spatially averaging the local degradation rates. A distributed electrothermal–degradation modelling framework is therefore needed to investigate the influence of intra-cell non-uniformity on battery ageing.

From a control perspective, minimising the mean temperature and internal thermal gradients does not necessarily minimise the resulting capacity loss, particularly when individual ageing mechanisms exhibit competing dependencies on the operating conditions. This motivates extending predictive battery cooling beyond conventional thermal-proxy objectives towards control strategies that explicitly account for the predicted degradation. However, incorporating distributed electrothermal and degradation dynamics into the coolant allocation problem introduces additional computational challenges for real-time control.

Against this backdrop, this thesis develops a unified modelling and predictive control framework for hybrid battery cooling. The work progresses from computationally efficient spatial thermal modelling to the dynamic integration of tab and lateral surface cooling, real-time predictive coolant allocation, and ultimately degradation-aware predictive cooling. The overall aim is to exploit the complementary characteristics of different cooling interfaces while balancing thermal performance, coolant utilisation, computational efficiency, and battery lifetime.

1.2 Research questions

Based on the research gaps identified in Section 1.1, this thesis addresses the following research questions:

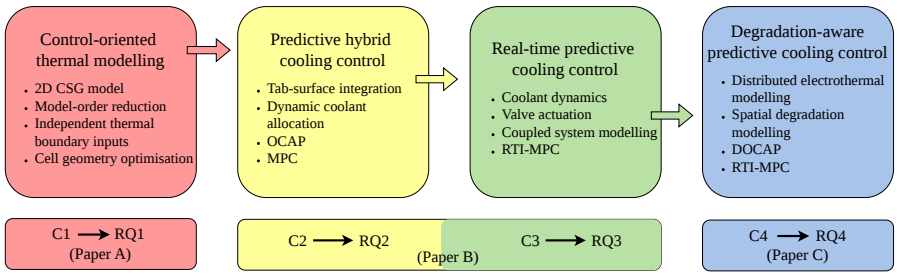

Figure 1.2: Overview of the research progression.

RQ1 How can a computationally efficient, spatially resolved thermal modelling framework be developed for battery cells that enables independent cooling control at the tab and lateral surface interfaces, remains applicable across various battery form factors and chemistries, and is suitable for real-time control?

RQ2 How can the complementary thermal characteristics of tab and lateral surface cooling be exploited through dynamic coolant allocation to balance temperature regulation and thermal uniformity under a limited coolant resource?

RQ3 How can the nonlinear optimal coolant allocation problem, accounting for the coupled battery, coolant, and actuator dynamics, be solved with sufficient computational efficiency for real-time predictive control?

RQ4 How can spatial variations in the internal states of battery cells be incorporated into degradation modelling, and how can the resulting model be exploited in predictive cooling control to mitigate battery degradation?

1.3 Thesis contributions

The research questions in the preceding section are addressed through four interconnected contributions, summarised in Fig. 1.2.

C1 Control-oriented spatial thermal modelling. A computationally efficient two-dimensional (2D) thermal modelling framework is developed

by combining the Chebyshev spectral-Galerkin (CSG) method with model component decomposition. The framework enables reconstruction of the spatial temperature distribution while allowing independent cooling control at the tab and lateral surface interfaces. Models of different orders can be generated to balance prediction accuracy and computational complexity, and the framework can be parametrised for different battery cell geometries. The models are evaluated against high-fidelity FE and conventional lumped TEC models, and their applicability to several model-based thermal management problems is investigated. This contribution addresses RQ1 and is presented in Paper A.

C2 Integration of tab and lateral surface cooling. The complementary thermal characteristics of tab and surface cooling are exploited through an integrated multi-channel cooling architecture. The dynamic distribution of coolant among the different cooling interfaces is formulated as an optimal coolant allocation problem (OCAP) to regulate cell temperature and reduce internal thermal gradients subject to coolant availability and operational constraints. The problem is solved using MPC. This contribution addresses RQ2 and is presented in Paper B.

C3 Real-time predictive coolant allocation. The battery modelling framework is extended to incorporate coolant and valve actuation dynamics, enabling a unified representation of the coupled battery-cooling system. A real-time iteration MPC (RTI-MPC) scheme is developed to efficiently solve the resulting nonlinear optimal coolant allocation problem. The scheme substantially reduces the online computational burden while retaining comparable closed-loop performance to nonlinear MPC (NMPC). This contribution addresses RQ3 and is presented in Paper B.

C4 Degradation-aware predictive cooling. A spatially distributed electrothermal and degradation modelling framework is developed to capture the influence of intra-cell variations in temperature, current distribution, and SoC on battery degradation. The framework is incorporated into a degradation-aware optimal coolant allocation problem (DOCAP), enabling direct consideration of battery degradation in predictive cooling control. The resulting approach is implemented using RTI-MPC and evaluated against conventional thermal-proxy control to investigate the benefits of direct degradation minimisation. This contribution addresses

RQ4 and is presented in Paper C.

1.4 Thesis outline

This thesis is organised into two parts. Part I provides an overview of the research, including the theoretical background, modelling framework, control methodologies, and main findings. It comprises six chapters, organised as follows:

- **Chapter 1** introduces the research motivation, research questions, and main contributions of the thesis.
- **Chapter 2** provides a background on lithium-ion batteries, the influence of temperature on battery performance and degradation, and battery thermal management systems, with particular emphasis on integrated tab-surface cooling.
- **Chapter 3** presents the electrical, thermal, and degradation modelling frameworks, together with the coolant and valve actuation models.
- **Chapter 4** presents the optimal coolant allocation and predictive control strategies, including their real-time implementation.
- **Chapter 5** summarises the appended papers and outlines their individual contributions to the thesis.
- **Chapter 6** presents the main conclusions in relation to the research questions and discusses directions for future research.

Part II comprises the three appended papers (Papers A–C), which provide a detailed account of the work presented in this thesis.

CHAPTER 2

Battery thermal management

2.1 Lithium-ion batteries and temperature effects

The basic operating principle of the lithium-ion battery (LiB) is illustrated in Fig. 2.1. LiBs store and release energy through the reversible transport of lithium (Li)-ions between the negative and positive electrodes. During discharge, Li-ions migrate from the negative electrode through the electrolyte and separator to the positive electrode, while electrons flow through the external circuit to supply the load. During charging, these processes are reversed [29]. A typical commercial vehicle cell consists of a graphite-based negative electrode, a positive electrode such as nickel-manganese-cobalt (NMC) or lithium-iron-phosphate (LFP), an electrolyte, and an electronically insulating but ionically conductive separator [30]. The electrode materials are coated onto metallic current collectors, typically copper for the negative electrode and aluminium for the positive electrode. The current collectors provide electronic conduction paths between the electrode materials and the external circuit through the electrical tabs or terminals.

Battery behaviour is strongly temperature dependent. At low temperatures, reduced ionic transport and increased internal resistance decrease available

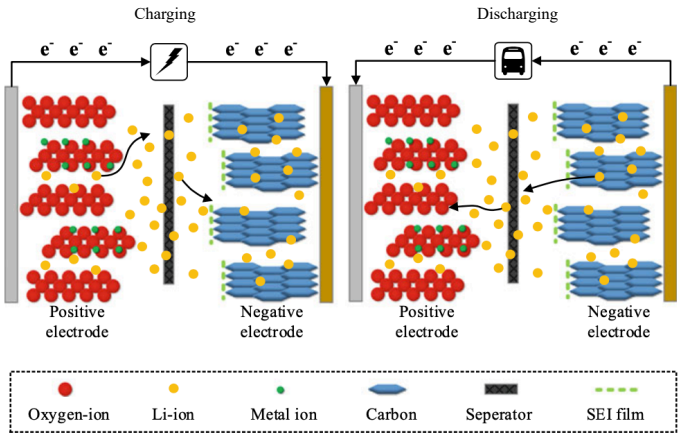

Figure 2.1: Basic structure and operating principle of a LiB [31].

power and capacity, and charging under cold conditions can promote lithium plating at the negative electrode [32], [33]. At elevated temperatures, electrochemical and parasitic reaction rates increase, accelerating degradation processes such as solid electrolyte interphase (SEI) growth [5], [34]. At extremely high temperatures, self-accelerating exothermic reactions can cause heat generation to exceed heat dissipation, potentially leading to thermal runaway [35].

In addition to the absolute temperature level, spatial temperature non-uniformity plays an important role in battery performance and degradation [6]. Owing to the temperature dependence of electrochemical kinetics and internal resistance, intra-cell temperature gradients can induce non-uniform current distributions, local variations in SoC, and heterogeneous degradation. These spatial variations can, in turn, influence local heat generation and further alter the temperature distribution within the cell. Such coupled phenomena highlight the importance of accounting for spatial variations in battery states when modelling battery degradation and designing thermal management strategies. Battery thermal management must address both the overall temperature level and intra-cell temperature gradients. These requirements motivate the battery cooling strategies discussed in the following section.

2.2 Battery thermal management systems

BTMSs are designed to maintain battery cells within an appropriate operating temperature range while limiting temperature non-uniformity within and among cells. Depending on the operating and ambient conditions, this may involve heating the battery at low temperatures or removing excess heat during operation [32]. Depending on the thermal load and system requirements, different cooling technologies can be employed, including air cooling, liquid cooling, phase-change materials, and combinations of these approaches [7]. In line with the scope of this thesis, the following discussion focuses on liquid-based battery cooling.

Battery cooling

In liquid-cooled battery systems, heat generated within the cell is conducted towards one or more external boundaries and subsequently transferred to a circulating coolant. The resulting thermal response is governed by the internal heat generation, thermal conductivity of the cell, effective heat-transfer coefficient, coolant temperature, coolant flow rate, and location of the cooled boundary. Liquid cooling is particularly suitable for applications with high thermal loads owing to its greater heat-transfer capability compared with conventional air cooling [7], [8].

Beyond the choice of cooling medium, the location of the cooling interface is key in determining the spatial temperature distribution within the cell. Owing to the anisotropic thermal conductivity of LiBs, different cooling interfaces provide different heat-conduction paths between the heat-generating regions and the coolant [8], [9], [12]. Consequently, cooling configurations with comparable heat removal rates can result in substantially different internal temperature distributions. This distinction is particularly relevant when both temperature regulation and intra-cell thermal uniformity are considered as thermal management objectives.

Integrated tab-surface cooling

As discussed in Section 1.1, surface and tab cooling exhibit distinct thermal characteristics owing to the anisotropic thermal conductivity of LiBs. This anisotropy arises primarily from the layered internal cell structure, consisting

of metallic current collectors, porous electrodes, and separators with different thermophysical properties. The relatively high thermal conductivities of the copper and aluminium current collectors facilitate heat conduction along the electrode layers, whereas heat conduction perpendicular to these layers is restricted by the relatively low thermal conductivities of the electrode and separator materials [9], [10].

In cylindrical cells, the winding of the electrode layers results in a higher axial thermal conductivity than radial thermal conductivity. Similarly, in pouch cells, heat conduction along the electrode layers is more effective than through the thickness of the cell. Surface cooling benefits from a relatively large heat-transfer area, making it effective at limiting the overall temperature rise. However, heat generated in the cell interior must travel predominantly along the direction of lower thermal conductivity. Consequently, relatively low surface temperatures can coexist with substantially higher core temperatures, resulting in pronounced internal temperature gradients, as illustrated in Figs. 1.1(a) and (c). Tab cooling, on the other hand, utilises the metallic current collectors as heat-conduction paths towards the electrical tabs or terminals. The favourable thermal conductivity along these paths promotes a more uniform internal temperature distribution, as illustrated in Figs. 1.1(b) and (d). However, the overall heat removal capability can be limited by the thermal resistance between the current collectors and the tabs, the dimensions and positioning of the tabs, and the relatively small heat-transfer area available at these interfaces [11], [36], [37], [38], [39]. Consequently, despite achieving better thermal uniformity, tab cooling may result in a higher average cell temperature rise than surface cooling, particularly under high thermal loads.

The effectiveness of each cooling configuration is therefore closely linked to the geometrical design of the cell. For instance, the cell dimensions determine both the heat-conduction distances and the available heat-transfer areas, while the arrangement and dimensions of the tabs influence the thermal resistance along the tab cooling path. Consequently, a cooling configuration that is favourable for one cell geometry may not necessarily be optimal for another.

The complementary strengths and limitations of these cooling configurations motivate their integration within a common thermal management architecture. By combining the heat removal capability of surface cooling with the ability of tab cooling to reduce internal temperature gradients, both tem-

perature regulation and thermal uniformity can be addressed simultaneously. However, the relative cooling demand at each interface may vary with the operating conditions and evolving thermal state of the cell. Rather than selecting a single cooling interface or imposing a fixed coolant distribution between the interfaces, the proposed integrated tab-surface cooling (ITSC) system allows the cooling effort at each interface to be adjusted independently. This provides additional control degrees of freedom for regulating both the overall temperature rise and the spatial temperature distribution within the cell.

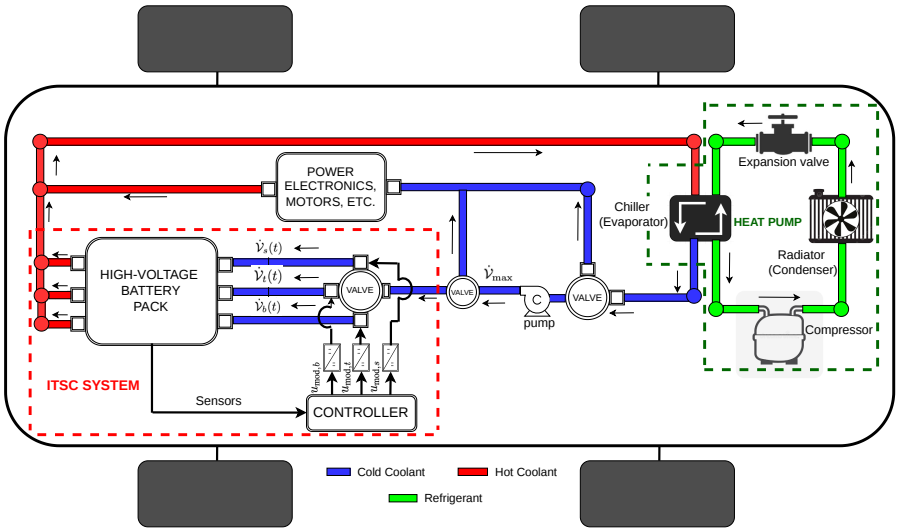

Figure 2.2: Schematic of the proposed ITSC system integrated into an EV thermal management system. Coolant from a common supply is dynamically distributed among the lateral surface, top, and bottom cooling channels. Reproduced from Paper B.

A schematic of the proposed ITSC system within an EV thermal management system is illustrated in Fig. 2.2. The system comprises a coolant circulation loop, a pump, dedicated cooling channels interfacing with the battery, and flow control valves that regulate the coolant distribution among the channels. Heat absorbed by the coolant is subsequently dissipated through the vehicle's heat exchange system. For the cylindrical cell configuration considered in this thesis, the coolant supplied by a common pump is distributed among three

cooling channels interfacing with the lateral surface, top, and bottom of the cell. The flow rate through each channel is regulated through valve actuation. Depending on the coolant distribution, the system can operate under surface-only cooling, tab-only cooling, or simultaneous tab and surface cooling. The individual channel flow rates are coupled through the total available coolant flow rate, requiring a coordinated allocation of the available cooling resource among the interfaces.

The dynamic allocation of coolant introduces additional dynamics beyond those of the battery itself. As heat is transferred from the cell to the coolant, the coolant temperature evolves along each cooling channel, while the channel flow rates influence the convective heat-transfer conditions at the corresponding cell boundaries. Furthermore, the valve actuation dynamics govern the redistribution of coolant among the channels. Effective control of the ITSC system therefore requires a modelling framework that captures the spatial temperature distribution within the battery and its interaction with the coolant and valve actuation dynamics. The development of such a framework is the subject of the following chapter.

CHAPTER 3

Battery and cooling system modelling

The analysis and control of the ITSC system require a model that captures the interactions between the battery and the cooling system. The battery model describes the electrical, thermal, and degradation dynamics, while the cooling system model describes the coolant and valve actuation dynamics governing heat transfer at the independently actuated cooling interfaces.

This chapter presents the development of these models, with particular emphasis on capturing spatial variations in the internal battery states while maintaining computational efficiency for control-oriented applications. The individual subsystem models are subsequently combined into a coupled battery–cooling system model for the predictive control formulations presented in the following chapter.

3.1 Battery model

Electrical modelling

Battery electrical models can be broadly classified into physics-based electrochemical [40], equivalent circuit [41], and data-driven [42] models. Elec-

trochemical models describe the internal transport and reaction processes of the cell through conservation laws and electrochemical kinetics. The Doyle–Fuller–Newman model [40] is a representative example, resolving quantities such as solid- and electrolyte-phase concentrations and potentials through a set of coupled partial differential equations (PDEs). Such models provide detailed information about the internal electrochemical states, but their relatively high computational complexity can limit their direct use in real-time control applications. Data-driven models approximate the battery input–output behaviour directly from experimental data, but their accuracy and generalisability depend on the quality and operating range of the training data. Equivalent circuit models (ECMs) provide an intermediate representation, in which the dominant electrical dynamics are described using resistive, capacitive, and voltage source elements. Their relatively low computational complexity, combined with sufficient prediction accuracy over relevant operating conditions, makes them particularly suitable for control-oriented applications.

Conventional ECMs, however, represent the electrical behaviour of the battery using lumped parameters and states, without explicitly accounting for spatial variations within the cell. While such a representation is suitable for describing the overall electrical response, it cannot capture the non-uniform current distributions and local SoC variations associated with intra-cell temperature gradients. Accordingly, both lumped and spatially distributed ECM formulations are considered in this thesis.

Lumped electrical model

The electrical behaviour of the battery is represented by the ECM illustrated in Fig. 3.1. The model consists of an open-circuit voltage source V_{ocv}, which is a nonlinear function of the SoC, a series resistance R_0, and p parallel resistance–capacitance (RC) branches connected in series. Each branch is characterised by a resistance R_i and capacitance C_i, representing the transient polarisation dynamics of the cell [43]. The model is expressed in continuous-time state-space form as

$$\dot{X}_e(t) = \mathcal{A}_e X_e(t) + \mathcal{B}_e U_e(t), \tag{3.1a}$$

$$\mathcal{Y}_e(t) = \mathcal{C}_e X_e(t) + \mathcal{D}_e U_e(t) + V_{\mathrm{ocv}}\left(\mathrm{SoC}(t)\right), \tag{3.1b}$$

where the state vector is defined as

$$X_e(t) = \begin{bmatrix} \mathrm{SoC}(t), & V_1(t), & \cdots & V_p(t) \end{bmatrix}^T, \tag{3.2}$$

with V_i denoting the voltage across the ith RC branch. The model input is the applied battery current, $U_e = I_b$, with positive current corresponding to discharge and negative current to charge. The output $\mathcal{Y}_e(t)$ represents the terminal voltage. The corresponding system matrices are

$$A_e = \begin{bmatrix} 0 & 0 & \cdots & 0 \\ 0 & -\frac{1}{R_1 C_1} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & -\frac{1}{R_p C_p} \end{bmatrix}, \quad B_e = \begin{bmatrix} -\frac{1}{3600 C_{\text{nom}}} \\ \frac{1}{C_1} \\ \vdots \\ \frac{1}{C_p} \end{bmatrix}, \tag{3.3a}$$

$$C_e = \begin{bmatrix} 0 & -1 & \cdots & -1 \end{bmatrix}, \qquad D_e = -R_0, \tag{3.3b}$$

where C_{nom} is the nominal cell capacity expressed in ampere-hours (Ah) and the factor 3600 converts the capacity from Ah to ampere-seconds (coulombs). This lumped formulation describes the overall electrical response of the cell using a single SoC state and spatially uniform electrical parameters. It is employed in the thermal modelling and coolant allocation frameworks developed in Papers A and B.

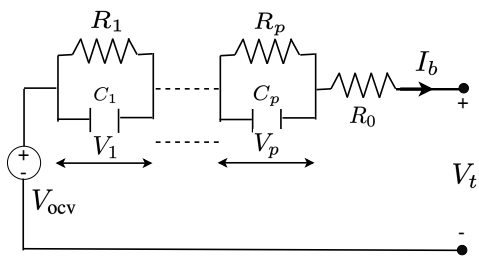

Figure 3.1: Equivalent circuit model of the battery cell with p RC branches.

Spatially distributed electrical model

The lumped ECM neglects the influence of intra-cell temperature variations on the local electrical behaviour. To address this limitation, the ECM is extended

to a spatially distributed formulation over the cylindrical cell domain Ω_b, as developed in Paper C. The model accounts for spatial variations in current, SoC, and polarisation voltages, while retaining the computational efficiency of the equivalent circuit representation.

Assuming ideal electrical current collectors, all spatial locations share a common terminal voltage V_t. However, variations in local temperature, SoC, polarisation voltages, and degradation states can result in non-uniform current distributions within the cell. Let $i(r, z, t)$ denote the local volumetric current density, defined over radial r and axial z coordinates at time t, with positive values corresponding to discharge, and negative to charge. The corresponding cell-equivalent local current is defined as

$$I_\ell(r, z, t) = i(r, z, t)V_b, \tag{3.4}$$

where V_b is the cell volume. Using a two-RC-branch ECM, the local electrical dynamics are expressed as

$$\dot{Z}(r, z, t) = -\frac{I_\ell(r, z, t)}{3600 C_{\text{avail}}(r, z, t)}, \tag{3.5a}$$

$$\dot{V}_j(r, z, t) = -\frac{V_j(r, z, t)}{R_j(r, z, t)C_j(r, z, t)} + \frac{I_\ell(r, z, t)}{C_j(r, z, t)}, \quad j = 1, 2, \tag{3.5b}$$

$$V_t(t) = V_{\text{ocv}}(Z(r, z, t)) - \sum_{j=1}^{2} V_j(r, z, t) - R_0(r, z, t)I_\ell(r, z, t), \tag{3.5c}$$

where $Z(r, z, t)$ denotes the local SoC, $C_{\text{avail}}(r, z, t)$ is the available cell-equivalent local capacity, and R_0, R_j, and C_j are the local ECM parameters. These parameters depend on the local operating conditions, including temperature and SoC. Rearranging the terminal voltage equation gives the local current density as

$$i(r, z, t) = \frac{V_{\text{ocv}}(Z(r, z, t)) - \sum_{j=1}^{2} V_j(r, z, t) - V_t(t)}{R_0(r, z, t)V_b}. \tag{3.6}$$

The local current distribution is constrained by the applied battery current through the global current conservation condition

$$I_b(t) = \int_{\Omega_b} i(r, z, t)\, d\Omega_b = \frac{1}{V_b} \int_{\Omega_b} I_\ell(r, z, t)\, d\Omega_b. \tag{3.7}$$

The volume-averaged SoC is defined as

$$Z_v(t) = \frac{1}{V_b} \int_{\Omega_b} Z(r, z, t)\, d\Omega_b. \tag{3.8}$$

For numerical implementation, the spatial integrals are evaluated using tensor-product Gauss–Legendre quadrature in the radial and axial directions. This approach provides accurate numerical integration using relatively few quadrature nodes for sufficiently smooth spatial distributions, while accounting for the cylindrical volume weighting. The quadrature nodes and weights can be precomputed, thereby reducing the computational burden during online simulation and control. With N_e quadrature nodes and normalised volume weights w_i, the current conservation condition becomes

$$I_b(t) \approx \sum_{i=1}^{N_e} w_i I_{\ell,i}(t), \qquad \sum_{i=1}^{N_e} w_i = 1. \tag{3.9}$$

The resulting electrical model comprises $3N_e$ dynamic states representing the local SoC and two polarisation voltages at each quadrature node. The local electrical states and parameters determine the current distribution, which subsequently influences the spatial heat generation and degradation rates. Under spatially uniform operating conditions and identical local electrical states and parameters, the model reduces to the conventional lumped ECM.

Thermal modelling

Battery thermal models describe the generation, storage, and transfer of heat within a cell and between the cell and its surroundings. These models range from lumped representations to spatially distributed models that resolve the temperature field throughout the battery. For battery management and control applications, the main modelling trade-off lies between computational cost and the level of thermal detail required.

Lumped thermal equivalent circuit (TEC) models represent heat storage and transfer using thermal capacitances and resistances, respectively. Their relatively low computational cost makes them attractive for control-oriented applications [15], [41], [44], [45]. However, lumped modelling is based on assumptions regarding the spatial temperature distribution within the battery. In particular, each thermal node is assumed to have a spatially uniform temperature. The validity of such an approximation is commonly assessed using the Biot number [17],

$$\mathrm{Bi} = \frac{hL_c}{k}, \tag{3.10}$$

where h is the convective heat-transfer coefficient, L_c is a characteristic length, and k is the thermal conductivity. When $\mathrm{Bi} \ll 1$, the internal thermal resistance is small relative to the surface convective resistance, and a lumped temperature representation is generally appropriate. For larger Biot numbers, internal temperature gradients become increasingly important, particularly under non-uniform cooling conditions.

Starting from the conservation of energy, physics-based thermal models can describe the transient heat conduction through partial differential equations (PDEs) subject to appropriate initial and boundary conditions [17], [46], [47]. Such models can explicitly account for the battery geometry, anisotropic thermal properties, internal heat generation, and the thermal conditions imposed at the cell boundaries. Although conventional numerical methods, such as finite element (FE) and finite difference methods, can provide high spatial resolution, their computational requirements may limit their direct application in real-time control.

Spectral methods provide a means of obtaining reduced-order physics-based thermal models that balance thermal fidelity and computational complexity [48], [49]. Rather than discretising the spatial domain into a large number of local elements, the temperature field is approximated using a finite set of global basis functions. Low-order spectral-Galerkin models [19], [20] have previously been developed for cylindrical and pouch battery cells, retaining spatial temperature information with fewer states. These formulations, however, do not provide independent thermal inputs at the different cell boundaries, which hinders direct application to the integrated tab–surface cooling problem considered in this thesis.

To address this limitation, Paper A develops a control-oriented thermal modelling framework for battery cells that combines the spectral-Galerkin method with Chebyshev basis functions [48] and model component decomposition. For a cylindrical cell, the temperature distribution $T(r, z, t)$ at time t in the radial r and axial z coordinates is governed by [46]

$$\rho_b c_{p,b} \frac{\partial T(r,z,t)}{\partial t} - k_{r,b} \frac{\partial^2 T(r,z,t)}{\partial r^2} - k_{r,b} \frac{\partial T(r,z,t)}{r \partial r} - k_{z,b} \frac{\partial^2 T(r,z,t)}{\partial z^2} = q(r,z,t), \tag{3.11}$$

subject to convection boundary conditions at the lateral surface, top, and

bottom of the cell, together with an adiabatic inner radial boundary,

$$-k_{r,b}\frac{\partial T}{\partial r} = h_s(T - T_{\infty,s}), \qquad r = R_{\text{out}}, \tag{3.12a}$$

$$k_{r,b}\frac{\partial T}{\partial r} = 0, \qquad r = R_{\text{in}}, \tag{3.12b}$$

$$-k_{z,b}\frac{\partial T}{\partial z} = h_t(T - T_{\infty,t}), \qquad z = L, \tag{3.12c}$$

$$k_{z,b}\frac{\partial T}{\partial z} = h_b(T - T_{\infty,b}), \qquad z = 0. \tag{3.12d}$$

Here, ρ_b and $c_{p,b}$ are the volume-averaged density and specific heat capacity, respectively, while $k_{r,b}$ and $k_{z,b}$ denote the effective radial and axial thermal conductivities. The variables h and T_∞ denote the convection coefficient and fluid free-stream temperature, respectively. R_{in} and R_{out} are the inner and outer radii of the cylindrical cell, while L denotes its length. The subscripts s, t, and b refer to the lateral surface, top, and bottom boundaries, respectively. $q(r, z, t)$ is the heat generation rate.

To obtain a reduced-order representation while retaining independent manipulation of the cooling boundaries, the temperature field is decomposed into homogeneous and particular components,

$$T(\tilde{r}, \tilde{z}, t) = T_h(\tilde{r}, \tilde{z}, t) + T_p(\tilde{r}, \tilde{z}, t). \tag{3.13}$$

The spatial coordinates are mapped from the physical domains $r \in [R_{\text{in}}, R_{\text{out}}]$ and $z \in [0, L]$ to the dimensionless domains $\tilde{r}, \tilde{z} \in [-1, 1]$ according to

$$\tilde{r} = \frac{2(r - R_{\text{in}})}{R_{\text{out}} - R_{\text{in}}} - 1, \qquad \tilde{z} = \frac{2z}{L} - 1. \tag{3.14}$$

The homogeneous component is approximated using a finite expansion of Chebyshev basis functions,

$$\hat{T}_h(\tilde{r}, \tilde{z}, t) = \sum_{m=0}^{M-1}\sum_{n=0}^{N-1} c_{mn}(t)\phi_m^{\tilde{r}}(\tilde{r})\phi_n^{\tilde{z}}(\tilde{z}), \tag{3.15}$$

where $\hat{T}_h$ denotes the approximation of T_h, $c_{mn}(t)$ are the time-varying expansion coefficients, and $\phi_m^{\tilde{r}}$ and $\phi_n^{\tilde{z}}$ are the Chebyshev basis functions in the radial and axial directions, respectively. The integers M and N determine the number of basis functions in the two spatial directions, with the model order

$O = MN$. The particular solution is similarly approximated, but decomposed into boundary-specific components. In the cylindrical cell formulation,

$$\hat{T}_p(\tilde{r}, \tilde{z}, t) = \hat{T}_p^s(\tilde{r}, \tilde{z})u_s(t) + \hat{T}_p^t(\tilde{r}, \tilde{z})u_t(t) + \hat{T}_p^b(\tilde{r}, \tilde{z})u_b(t), \tag{3.16}$$

where $\hat{T}_p$ denotes the approximation of the particular solution, and $\hat{T}_p^s$, $\hat{T}_p^t$, and $\hat{T}_p^b$ are the boundary-specific particular solution expansions associated with the lateral surface, top, and bottom boundaries, respectively. The corresponding thermal boundary inputs are

$$u_s(t) = h_s(t)T_{s,\infty}(t), \quad u_t(t) = h_t(t)T_{t,\infty}(t), \quad u_b(t) = -h_b(t)T_{b,\infty}(t). \tag{3.17}$$

This decomposition allows the contributions of the individual cooling interfaces to be represented independently, providing the basis for dynamic coolant allocation.

Applying the CSG method to the non-dimensionalised forms of (3.11)–(3.12) yields a finite-dimensional system of coupled ordinary differential equations (ODEs) of the form

$$\mathcal{G}_T\dot{X}_T(t) = \mathcal{A}_T X_T(t) + \mathcal{B}_T U_T(t) + \mathcal{F}_T W_T(t) \tag{3.18a}$$

$$\mathcal{Y}_T(t) = \mathcal{C}_T X_T(t) + \mathcal{D}_T U_T(t) \tag{3.18b}$$

where the state vector X_T is defined as

$$X_T = [c_{11},\cdots, c_{1N}, c_{21},\cdots, c_{2N},\cdots, c_{M1},\cdots, c_{MN}]^T. \tag{3.19}$$

The input vector U_T is composed of the cooling inputs applied at the sides of the cell, leading to

$$U_T(t) = [u_s(t),\ u_t(t),\ u_b(t)]^T. \tag{3.20}$$

The cooling inputs in (3.17) and (3.20) capture the thermal interaction between the battery cell and the surrounding coolant. The disturbance input is given by $W_T(t) = q(r, z, t)$, representing the volumetric heat generation rate. In Papers A and B, the volumetric heat generation is assumed spatially uniform, such that $q(r, z, t) = q(t)$. In Paper C, this assumption is relaxed by incorporating the spatially distributed heat generation obtained from the electrical model. The corresponding source projection is expressed as

$$\mathcal{F}_T W_T(t) = \int_{R_{\text{in}}}^{R_{\text{out}}} \int_0^L \phi_m^r(\tilde{r})\phi_n^z(\tilde{z})q(r, z, t)\, r\, dz\, dr, \tag{3.21}$$

This extension allows local electrical non-uniformity to influence the spatial temperature distribution while retaining the reduced-order thermal model structure. The output vector $\mathcal{Y}_T$ consists of representative temperatures within and on the cell, including the quantities required for battery–coolant coupling. In particular, the volume-averaged cell temperature is defined as

$$T_v(t) = \frac{1}{V_b} \int_{\Omega_b} T(r, z, t)\, d\Omega_b, \tag{3.22}$$

where V_b and Ω_b are the cell volume and battery physical domain, respectively. The area-averaged battery–coolant interface temperatures at the lateral surface, top, and bottom sides are defined as

$$T_{bc,s}(t) = \frac{1}{A_{ht,s}} \int_{A_{ht,s}} T(R_{\text{out}}, z, t)\, dA_{ht,s}, \tag{3.23a}$$

$$T_{bc,t}(t) = \frac{1}{A_{ht,t}} \int_{A_{ht,t}} T(r, L, t)\, dA_{ht,t}, \tag{3.23b}$$

$$T_{bc,b}(t) = \frac{1}{A_{ht,b}} \int_{A_{ht,b}} T(r, 0, t)\, dA_{ht,b}, \tag{3.23c}$$

respectively, where A_{ht} denotes the effective heat transfer area. In addition to the mean and interface temperatures, the root-mean-square (RMS) spatial temperature gradient is used to quantify thermal non-uniformity within the cell and is defined as

$$\delta T(t) = \left(\frac{1}{V_b} \int_{\Omega_b} \|\nabla T(r, z, t)\|_2^2 \, d\Omega_b \right)^{1/2}, \tag{3.24}$$

where

$$\nabla T(r, z, t) = \begin{bmatrix} \dfrac{\partial T(r, z, t)}{\partial r} \\ \dfrac{\partial T(r, z, t)}{\partial z} \end{bmatrix}. \tag{3.25}$$

The detailed derivations of the system matrices $\mathcal{G}_t$, $\mathcal{A}_t$, $\mathcal{B}_t$, $\mathcal{C}_t$, $\mathcal{D}_t$, and $\mathcal{F}_t$ are provided in Papers A and B.

Degradation modelling

Battery degradation is a cumulative process that manifests primarily through capacity loss and resistance. It is influenced by operating conditions such as

temperature, elapsed time, SoC, current, and charge throughput [26]. Battery degradation models can broadly be classified as physics-based, empirical, and semi-empirical [28], [50], [51].

Physics-based degradation models describe the underlying electrochemical and physicochemical ageing mechanisms, including solid-electrolyte interphase (SEI) growth, lithium plating, loss of active material, and electrolyte degradation. Although these models can provide detailed insight into processes responsible for battery ageing, they are complex and have a relatively high computational cost. Empirical models instead establish relationships between measured degradation and operating conditions using experimental ageing data. These models are generally computationally efficient, but their predictive accuracy depends on the operating conditions represented in the experimental data. Semi-empirical models combine experimentally identified parameters with physically motivated relationships, such as Arrhenius temperature dependence and SoC- and current-dependent ageing stress factors. They provide a compromise between computational efficiency and physical interpretability.

The degradation modelling adopted in this thesis is based on the semi-empirical model developed by Schimpe et al. for Lithium-Iron-Phosphate (LFP)/graphite cells [27]. In Paper C, the original lumped formulation is expressed as a system of dynamic states and subsequently extended to account for spatially non-uniform degradation arising from local variations in temperature, SoC, and current.

Lumped degradation model

The adopted degradation model represents the total normalised capacity loss as the sum of calendar ageing and three cycling-ageing mechanisms,

$$Q_{\text{loss}} = Q_C + Q_H + Q_L + Q_{LS}, \tag{3.26}$$

where Q_C denotes calendar ageing, Q_H represents high-temperature cycling degradation, Q_L accounts for degradation associated with low-temperature charging, and Q_{LS} represents the additional degradation associated with low-temperature charging at high SoC. Under constant operating conditions, the accumulated capacity loss is expressed as

$$Q_{\text{loss}} = k_C(T_v, Z)\sqrt{\tau} + k_H(T_v)\sqrt{\varphi_{\text{tot}}} + k_L(T_v, I_{\text{ch}})\sqrt{\varphi_{\text{ch}}} + k_{LS}(T_v, I_{\text{ch}}, Z)\varphi_{\text{ch}}, \tag{3.27}$$

where $\tau = t/3600$ denotes the elapsed time in hours, φ_{tot} and φ_{ch} are the accumulated total and charging throughputs in ampere-hours (Ah), respectively, $I_{\text{ch}}(t) = \max(-I_b(t), 0)$ denotes the charging current, and k_C, k_H, k_L, and k_{LS} are the corresponding degradation stress factors. The volume-averaged temperature T_v and SoC Z_v are used as the representative cell temperature and SoC, respectively. The stress factors capture the dependence of the individual ageing mechanisms on temperature, SoC, and charging current and are expressed as

$$k_C(T_v, Z_v) = k_{C,\text{ref}} \exp\left[-\frac{E_{a,C}}{R_g}\left(\frac{1}{T_v} - \frac{1}{T_{\text{ref}}}\right)\right] g_C(Z_v), \tag{3.28a}$$

$$k_H(T_v) = k_{H,\text{ref}} \exp\left[-\frac{E_{a,H}}{R_g}\left(\frac{1}{T_v} - \frac{1}{T_{\text{ref}}}\right)\right], \tag{3.28b}$$

$$k_L(T_v, I_{\text{ch}}) = k_{L,\text{ref}} \exp\left[\frac{E_{a,L}}{R_g}\left(\frac{1}{T_v} - \frac{1}{T_{\text{ref}}}\right)\right] \cdot \exp\left[\beta_L \frac{I_{\text{ch}} - I_{\text{ch,ref}}}{C_0}\right], \tag{3.28c}$$

$$k_{LS}(T_v, I_{\text{ch}}, Z_v) = k_{LS,\text{ref}} \exp\left[\frac{E_{a,LS}}{R_g}\left(\frac{1}{T_v} - \frac{1}{T_{\text{ref}}}\right)\right] \cdot \exp\left[\beta_{LS} \frac{I_{\text{ch}} - I_{\text{ch,ref}}}{C_0}\right] g_{LS}(Z_v). \tag{3.28d}$$

The SoC-dependent stress functions are defined as

$$g_C(Z_v) = \exp\left[\frac{\alpha F \left(U_{a,\text{ref}} - U_a(Z_v)\right)}{R_g T_{\text{ref}}} + k_0\right], \tag{3.29a}$$

$$g_{LS}(Z_v) = \frac{\operatorname{sgn}(Z_v - Z_{\text{ref}}) + 1}{2}. \tag{3.29b}$$

Here, E_a denotes the activation energy, R_g is the universal gas constant, F is the Faraday constant, α is the transfer coefficient, U_a is the anode open-circuit potential, and k_0 is a constant offset. The subscript ref denotes reference values, C_0 is the nominal cell capacity at the beginning of life, and β is the charge-current sensitivity coefficient.

To accommodate time-varying operating conditions and facilitate integration with the electrothermal model, Paper C expresses the degradation model in differential form. The accumulated throughputs are described by

$$\dot{\varphi}_{\text{tot}}(t) = \frac{|I_b(t)|}{3600}, \qquad \dot{\varphi}_{\text{ch}}(t) = \frac{I_{\text{ch}}(t)}{3600}, \tag{3.30}$$

and the corresponding capacity loss rates are

$$\dot{Q}_C = \frac{k_C(T_v, Z_v)}{7200\sqrt{\tau + \epsilon_\tau}}, \tag{3.31a}$$

$$\dot{Q}_H = \frac{k_H(T_v)|I_b|}{7200\sqrt{\varphi_{\text{tot}} + \epsilon_{\text{tot}}}}, \tag{3.31b}$$

$$\dot{Q}_L = \frac{k_L(T_v, I_{\text{ch}})I_{\text{ch}}}{7200\sqrt{\varphi_{\text{ch}} + \epsilon_{\text{ch}}}}, \tag{3.31c}$$

$$\dot{Q}_{LS} = \frac{k_{LS}(T_v, I_{\text{ch}}, Z_v)I_{\text{ch}}}{3600}, \tag{3.31d}$$

where ϵ_τ, ϵ_{tot}, and ϵ_{ch} are small positive regularisation constants introduced to avoid singularities at zero elapsed time and throughput. The degradation state vector is defined as

$$X_{\text{deg}}(t) = \begin{bmatrix} \varphi_{\text{tot}}, & \varphi_{\text{ch}}, & Q_C, & Q_H, & Q_L, & Q_{LS} \end{bmatrix}^T, \tag{3.32}$$

and the degradation model output is the total normalised capacity loss obtained as

$$\mathcal{Y}_{\text{deg}}(t) = Q_{\text{loss}}(t) = \mathcal{C}_{\text{deg}} X_{\text{deg}}(t), \tag{3.33}$$

where $C_{\text{deg}} = \begin{bmatrix} 0 & 0 & 1 & 1 & 1 & 1 \end{bmatrix}$. Accordingly, the capacity-based state of health (SoH) is defined as

$$\text{SoH}(t) = 1 - \mathcal{Y}_{\text{deg}}(t). \tag{3.34}$$

Spatially distributed degradation model

The lumped degradation model evaluates the ageing stress factors using representative cell-level quantities, thereby neglecting spatial variations in temperature, SoC, and current. However, the non-uniform electrothermal conditions within a battery cell can cause different regions to experience different degradation rates. To capture these effects, Paper C extends the degradation model to the spatial domain Ω_b, where the individual ageing mechanisms are evaluated using the local temperature $T(r, z, t)$, SoC $Z(r, z, t)$, and current $I_\ell(r, z, t)$.

To account for differences in nominal capacity between the cell under consideration and the reference cell used for degradation parameterisation, the

local cell-equivalent current is scaled as

$$\tilde{I}_\ell(r,z,t) = \frac{C_{\mathrm{ref}}}{C_0} I_\ell(r,z,t), \tag{3.35}$$

where C_{ref} denotes the nominal capacity of the reference cell and C_0 is the nominal capacity of the considered cell at the beginning of life. The corresponding local charging current is defined as $I_{\mathrm{ch}}(r,z,t) = \max\left(-\tilde{I}_\ell(r,z,t), 0\right)$. The accumulated total and charging throughputs are evaluated locally, with the corresponding dynamics given by

$$\dot{\varphi}_{\mathrm{tot}}(r,z,t) = \frac{|\tilde{I}_\ell(r,z,t)|}{3600}, \qquad \dot{\varphi}_{\mathrm{ch}}(r,z,t) = \frac{I_{\mathrm{ch}}(r,z,t)}{3600}. \tag{3.36a}$$

The local capacity loss rates are obtained by evaluating the stress factors defined in (3.28) using the corresponding spatially varying operating conditions. The total local capacity loss rate is expressed as

$$\dot{Q}_{\mathrm{loss}}(r,z,t) = \dot{Q}_C(r,z,t) + \dot{Q}_H(r,z,t) + \dot{Q}_L(r,z,t) + \dot{Q}_{LS}(r,z,t), \tag{3.37}$$

where the individual contributions retain the functional forms of Eq. 3.31, with the lumped quantities replaced by their corresponding local values. The local degradation state vector is defined as

$$X_{\mathrm{deg}}(r,z,t) = \left[\varphi_{\mathrm{tot}}, \quad \varphi_{\mathrm{ch}}, \quad Q_C, \quad Q_H, \quad Q_L, \quad Q_{LS}\right]^T, \tag{3.38}$$

where all states depend on (r,z,t). Thus, the same degradation mechanisms are retained as in the lumped model, but their evolution is determined by the local electrothermal conditions.

Assuming a uniform initial capacity density, the cell-level normalised capacity loss is obtained by spatially averaging the local capacity losses,

$$Q_{\mathrm{loss}}^{\mathrm{dist}}(t) = \frac{1}{V_b} \int_{\Omega_b} Q_{\mathrm{loss}}(r,z,t)\, d\Omega_b. \tag{3.39}$$

For numerical implementation, the local degradation states are propagated at the Gauss–Legendre quadrature nodes introduced in Section 3.1. The cell-level capacity loss is approximated as

$$Q_{\mathrm{loss}}^{\mathrm{dist}}(t) \approx \sum_{i=1}^{N_e} w_i Q_{\mathrm{loss},i}(t), \tag{3.40}$$

where w_i are the normalised volume weights and $Q_{\mathrm{loss},i}$ is the accumulated capacity loss at the ith quadrature node. The resulting distributed degradation model comprises $6N_e$ states. Its output represents the cell-level normalised capacity loss,

$$\mathcal{Y}_{\mathrm{deg}}^{\mathrm{dist}}(t) = Q_{\mathrm{loss}}^{\mathrm{dist}}(t), \tag{3.41}$$

and the corresponding capacity-based SoH is

$$\mathrm{SoH}^{\mathrm{dist}}(t) = 1 - \mathcal{Y}_{\mathrm{deg}}^{\mathrm{dist}}(t). \tag{3.42}$$

The accumulated local capacity loss influences the electrical dynamics through the available cell-equivalent local capacity,

$$C_{\mathrm{avail}}(r, z, t) = C_0 \left[1 - Q_{\mathrm{loss}}(r, z, t)\right]. \tag{3.43}$$

In addition to capacity fade, battery ageing can increase the internal resistance. Such effects can be represented by

$$R_j(r, z, t) = R_j^{\mathrm{BOL}}(Z, T) \left[1 + \mu_{R,j}(r, z, t)\right], \quad j \in \{0, 1, 2\}, \tag{3.44}$$

where R_j^{BOL} denotes the beginning-of-life resistance evaluated at the local SoC and temperature, and $\mu_{R,j}$ represents the accumulated local relative resistance increase due to ageing. The resulting variations in available capacity and internal resistance can influence the local current distribution and heat generation, which subsequently affect the temperature distribution and degradation rates. The electrical, thermal, and degradation dynamics are coupled through the evolving spatial distributions of the battery states. Under spatially uniform operating conditions and identical initial degradation states, the distributed model reduces to the lumped degradation formulation.

3.2 Cooling system model

The cooling system model describes the thermal dynamics of the coolant and the actuation mechanism governing its distribution among the independently controlled cooling interfaces of the ITSC system. The coolant model captures the exchange of heat between the battery and circulating coolant, while the valve actuation model relates the controller commands to the effective coolant flow rates supplied to the lateral surface, top, and bottom cooling channels.

Coolant modelling

Paper B develops a dynamic model for the mean coolant temperature in each cooling channel of the ITSC system, capturing the effects of coolant flow rate on convective heat transfer and advective energy transport. Heat transfer between the battery and the coolant is modelled as internal forced convection, with a water–glycol mixture circulated by a pump through cooling plates or jackets in thermal contact with the cell.

The coolant temperature dynamics are obtained by applying the first law of thermodynamics to a control volume comprising the coolant flowing through a channel in thermal contact with the battery. The coolant is assumed incompressible with constant density ρ_c and specific heat capacity $c_{p,c}$. Only sensible thermal energy is considered, while kinetic and potential energy changes and pressure work effects are neglected. The coolant temperature within each channel is represented by its bulk mean temperature, assuming dominant streamwise heat transfer and sufficient mixing across the channel cross-section. Under these simplifications, the energy balance yields the following first-order dynamical model:

$$\frac{d\bar{T}_{cl}(t)}{dt} = \frac{h(t)A_{ht}(T_{bc}(t) - \bar{T}_{cl}(t))}{\rho_c V_c c_{p,c}} + \frac{\dot{\mathcal{V}}(t)(T_{cl,\mathrm{in}}(t) - \bar{T}_{cl}(t))}{V_c}, \tag{3.45}$$

where $\bar{T}_{cl}$ is the mean coolant temperature, T_{bc} is the area-averaged battery–coolant interface temperature defined in (3.23), $T_{cl,\mathrm{in}}$ is the coolant inlet temperature, $\dot{\mathcal{V}}$ is the volumetric coolant flow rate, V_c is the coolant volume in the channel, A_{ht} is the effective heat transfer area, and h is the convective heat transfer coefficient.

The first term represents the convective heat exchange between the battery and coolant, and the second accounts for the transport of thermal energy associated with coolant entering and leaving the channel. The model is applied independently to the lateral surface, top, and bottom cooling channels, with the corresponding mean coolant temperatures forming the state vector

$$X_{cl}(t) = \begin{bmatrix} \bar{T}_{cl,s}(t) & \bar{T}_{cl,t}(t) & \bar{T}_{cl,b}(t) \end{bmatrix}^T. \tag{3.46}$$

The volumetric flow rates constitute the effective manipulated inputs to the coolant model. The dynamics are nonlinear owing to the product between the coolant temperature state $\bar{T}_{cl}$ and the flow rate input $\dot{\mathcal{V}}$.

Valve actuation modelling

The volumetric flow rates in (3.45) are regulated through valves that distribute the available coolant among the lateral surface, top, and bottom cooling channels. In Paper B, the valve actuation is modelled using pulse-width modulation (PWM) [52], [53], which provides a relationship between the controller commands and the effective coolant flow rates supplied to the individual channels.

For each cooling channel, the controller generates a modulating signal $u_{\mathrm{mod}}(t) \in [0, 1]$, which is compared with a high-frequency carrier signal to produce a binary switching signal $s_w(t) \in \{0, 1\}$. To avoid explicitly modelling the high-frequency switching behaviour, switching-period averaging [16], [53] is applied over the switching period t_{sp}, thereby replacing the switching signals with their time-averaged equivalents:

$$u_{\mathrm{mod}}(t) := \frac{1}{t_{sp}} \int_{t-t_{sp}}^{t} s_w(\tau)\, d\tau = \frac{s_{w,\mathrm{ON}}(t)}{t_{sp}}, \tag{3.47}$$

where $s_{w,\mathrm{ON}}$ is the total ON duration of s_w within one PWM period. The resulting continuous signals u_{mod} represent the valve duty cycles, corresponding to the fraction of time each valve remains open per switching period. Let $\dot{\mathcal{V}}_{\max}$ denote the maximum available volumetric flow capacity supplied by the pump. The effective volumetric flow rates supplied to the three cooling channels are

$$\dot{\mathcal{V}}_s(t) = \dot{\mathcal{V}}_{\max} u_{\mathrm{mod},s}(t),\ \dot{\mathcal{V}}_t(t) = \dot{\mathcal{V}}_{\max} u_{\mathrm{mod},t}(t),\ \dot{\mathcal{V}}_b(t) = \dot{\mathcal{V}}_{\max} u_{\mathrm{mod},b}(t). \tag{3.48}$$

The duty cycles satisfy

$$0 \le u_{\mathrm{mod},j}(t) \le 1, \qquad j \in \{s, t, b\}, \tag{3.49}$$

and the finite coolant flow capacity imposes

$$u_{\mathrm{mod},s}(t) + u_{\mathrm{mod},t}(t) + u_{\mathrm{mod},b}(t) \le 1. \tag{3.50}$$

The inequality in (3.50) allows the controller to use less than the total available coolant flow when the thermal demand is low. The valve actuation model thus provides a continuous mapping between the controller commands u_{mod} and the effective volumetric flow rates used in the coolant dynamics, without explicitly simulating individual valve switching events.

3.3 Coupled battery–cooling system

The electrical, thermal, degradation, coolant, and valve actuation models are coupled through the battery heat generation, the dependence of degradation on local electrothermal conditions, the heat exchange between the battery and coolant, and the allocation of the available coolant flow among the cooling interfaces. These interactions form the coupled battery–cooling system used for the subsequent predictive control formulations.

The electrical and thermal models are coupled through the heat generated during battery operation. In the lumped formulation employed in Papers A and B, the volumetric heat generation rate is assumed spatially uniform and is expressed as [54]

$$q(t) = \frac{q_{\mathrm{rev}}(t) + q_{\mathrm{irr}}(t)}{V_b}, \tag{3.51a}$$

$$q_{\mathrm{rev}}(t) = -I_b(t) T_v(t) \frac{\partial V_{\mathrm{ocv}}}{\partial T_v}, \tag{3.51b}$$

$$q_{\mathrm{irr}}(t) = I_b^2(t) R_0 + \sum_{j=1}^{p} \frac{V_j^2(t)}{R_j}, \tag{3.51c}$$

where V_b denotes the cell volume, q_{rev} and q_{irr} represent the reversible and irreversible heat generation rates, respectively. The reversible component accounts for the entropic heat associated with the electrochemical reactions, where $\partial V_{\mathrm{ocv}}/\partial T$ denotes the entropic temperature coefficient of the open-circuit voltage. Depending on the current direction and the sign of this coefficient, the reversible contribution may either generate or absorb heat. The irreversible component accounts for ohmic and polarisation losses, which contribute to heat generation during both charging and discharging.

In Paper C, the heat generation model is extended to account for spatial variations in the electrical states and parameters. The volumetric heat generation rate comprises the corresponding local reversible and irreversible

components, expressed as

$$q(r,z,t) = q_{\text{rev}}(r,z,t) + q_{\text{irr}}(r,z,t), \tag{3.52a}$$

$$q_{\text{rev}}(r,z,t) = -\frac{T}{V_b}\left(\frac{V_{\text{ocv}} - V_t - \sum_{j=1}^{2} V_j}{R_0}\right)\frac{\partial V_{\text{ocv}}}{\partial T}, \tag{3.52b}$$

$$q_{\text{irr}}(r,z,t) = \frac{1}{V_b}\left[\frac{\left(V_{\text{ocv}} - V_t - \sum_{j=1}^{2} V_j\right)^2}{R_0} + \sum_{j=1}^{2}\frac{V_j^2}{R_j}\right]. \tag{3.52c}$$

All spatially varying quantities on the right-hand side are evaluated at the corresponding local position (r, z) and time t, while V_t is the common terminal voltage. The local irreversible heat generation is expressed in terms of the terminal and polarisation voltages using the distributed ECM. The resulting spatially distributed heat generation is incorporated into the thermal model through the Galerkin source projection defined in (3.21). Conversely, the temperature distribution influences the local ECM parameters and current distribution, establishing a bidirectional electrothermal coupling.

The degradation model is coupled to the electrical and thermal models through the operating conditions that determine the ageing stress factors and capacity loss rates. In the lumped formulation, these quantities are evaluated using representative cell-level variables, including T_v and Z_v. In the distributed formulation developed in Paper C, the local temperature, SoC, current, and accumulated throughputs determine the degradation rates at individual spatial locations. The accumulated local capacity loss modifies the available cell-equivalent capacity and thereby influences the subsequent electrical dynamics. Ageing-induced changes in internal resistance can also be incorporated into the electrical model. Consequently, the coupled formulation captures the interaction between electrothermal non-uniformity and battery degradation.

Heat exchange between the battery and coolant introduces a bidirectional coupling between the thermal and coolant models. The mean coolant temperatures in the lateral surface, top, and bottom channels determine the corresponding thermal boundary inputs,

$$u_s(t) = h_s(t)\bar{T}_{cl,s}(t), \quad u_t(t) = h_t(t)\bar{T}_{cl,t}(t), \quad u_b(t) = -h_b(t)\bar{T}_{cl,b}(t), \tag{3.53}$$

These inputs govern the thermal boundary conditions and influence the spatial temperature distribution within the cell. Conversely, the thermal model

provides the battery–coolant interface temperatures $T_{bc,s}$, $T_{bc,t}$, and $T_{bc,b}$ defined in (3.23), which determine the convective heat exchange in the respective coolant channels.

The valve actuation model links the controller commands to the coolant dynamics. The valve duty cycles determine the effective channel flow rates according to (3.48)–(3.50). The controller input can be expressed as

$$u_{\mathrm{mod}}(t) = \left[u_{\mathrm{mod},s}(t),\, u_{\mathrm{mod},t}(t),\, u_{\mathrm{mod},b}(t)\right]^T, \tag{3.54}$$

and the corresponding effective battery–coolant system input as

$$U(t) = \left[\dot{\mathcal{V}}_s(t), \quad \dot{\mathcal{V}}_t(t), \quad \dot{\mathcal{V}}_b(t)\right]^T. \tag{3.55}$$

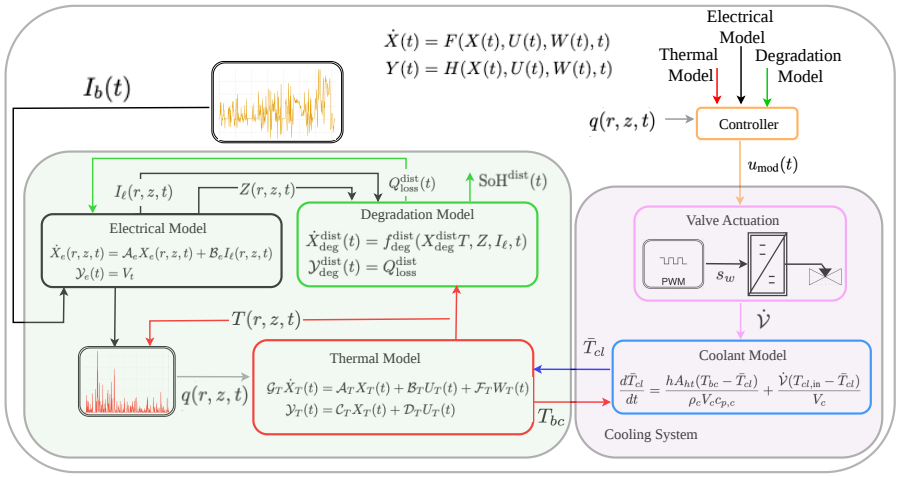

Figure 3.2: Interconnections among the electrical, thermal, degradation, coolant, and valve actuation models in the coupled battery–cooling system.

The resulting interconnections among the electrical, thermal, degradation, coolant, and valve actuation models are illustrated in Fig. 3.2. The complete nonlinear battery–cooling system is represented as

$$\dot{X}(t) = F(X(t), U(t), W(t), t), \tag{3.56a}$$
$$Y(t) = H(X(t), U(t), W(t), t), \tag{3.56b}$$

where $F(\cdot)$ and $H(\cdot)$ denote the nonlinear state and output mappings, respectively. For the spatially distributed formulation, the complete state vector is assembled as

$$X(t) = \begin{bmatrix} X_e^T(t) & X_T^T(t) & (X_{\text{deg}}^{\text{dist}})^T(t) & X_{cl}^T(t) \end{bmatrix}^T, \tag{3.57}$$

where X_e, X_T, $X_{\text{deg}}^{\text{dist}}$, and X_{cl} denote the electrical, thermal, distributed degradation, and coolant state vectors, respectively. The exogenous input vector is

$$W(t) = \begin{bmatrix} I_b(t) & T_{cl,\text{in},s}(t) & T_{cl,\text{in},t}(t) & T_{cl,\text{in},b}(t) \end{bmatrix}^T. \tag{3.58}$$

The system outputs comprise the terminal voltage, representative battery temperatures, degradation output, and coolant temperatures,

$$Y(t) = \begin{bmatrix} \mathcal{Y}_e^T(t) & \mathcal{Y}_T^T(t) & Y_{\text{deg}}^{\text{dist}}(t) & X_{cl}^T(t) \end{bmatrix}^T. \tag{3.59}$$

The coupled model (3.56) forms the basis for the predictive cooling control formulations developed in the following chapter.

CHAPTER 4

Predictive battery cooling control

The ITSC system enables dynamic coolant allocation among multiple cooling interfaces to regulate battery temperature and thermal uniformity. Model predictive control (MPC) provides a framework for optimising this allocation while accounting for the coupled battery–cooling dynamics and operating constraints. Paper B develops a temperature- and gradient-based coolant allocation strategy, while Paper C extends the formulation to explicitly minimise battery degradation.

This chapter presents the corresponding optimal coolant allocation problems, their MPC formulations, and the real-time iteration (RTI) approach adopted to enhance computational efficiency.

4.1 Optimal coolant allocation

Section 2.2 discussed the complementary thermal characteristics of tab and surface cooling, which motivate the dynamic allocation of coolant among the different cooling interfaces. The coolant distribution is adjusted according to the evolving battery operating conditions and the desired control objectives. At each discrete sampling instant k, the coolant allocation problem is

described by the valve duty-cycle vector

$$u_{\text{mod}}(k) = \begin{bmatrix} u_{\text{mod},s}(k) & u_{\text{mod},t}(k) & u_{\text{mod},b}(k) \end{bmatrix}^T. \tag{4.1}$$

As described in Chapter 3, the duty cycles determine the effective coolant flow rates supplied to the lateral surface, top, and bottom cooling channels through

$$\dot{V}_j(k) = \dot{V}_{\max} u_{\text{mod},j}(k), \qquad j \in \{s, t, b\}, \tag{4.2}$$

subject to

$$0 \leq u_{\text{mod},j}(k) \leq 1, \qquad \sum_{j \in \{s,t,b\}} u_{\text{mod},j}(k) \leq 1. \tag{4.3}$$

Coolant allocation objectives

Two optimal coolant allocation problems are formulated based on different representations of battery health. The first regulates battery temperature and thermal uniformity as indirect indicators of battery health, while the second explicitly incorporates battery degradation into the control objective.

Temperature- and gradient-based coolant allocation. The optimal coolant allocation problem (OCAP), developed in Paper B, seeks to regulate battery temperature towards a desired reference while minimising internal thermal gradients. For comparison with the degradation-aware formulation, Paper C augments the original objective to include a coolant utilisation penalty, giving

$$\begin{aligned} \min_{u_{\text{mod},k}} J_{\text{OCAP}}(k) := \sum_{i=0}^{N_p-1} \Bigg[& \lambda_T \left\| T^r(k+i) - \hat{\mathcal{Y}}_T(k+i) \right\|_{Q_e}^2 + \lambda_\delta \left\| \delta\hat{T}(k+i) \right\|_{Q_\delta}^2 \\ & + \lambda_u \left(\mathbf{1}^T u_{\text{mod}}(k+i) \right)^2 + \lambda_{\Delta u} \left\| \Delta u_{\text{mod}}(k+i) \right\|_R^2 \Bigg], \end{aligned} \tag{4.4}$$

where N_p is the prediction horizon, $u_{\text{mod},k} := \{u_{\text{mod}}(k+i)\}_{i=0}^{N_p-1}$ denotes the sequence of predicted valve duty-cycle commands, T^r is the temperature reference, $\hat{\mathcal{Y}}_T$ contains the predicted thermal outputs, and $\delta\hat{T}$ denotes the predicted RMS spatial temperature gradient. The scalar weights λ_T, λ_δ, λ_u, and $\lambda_{\Delta u}$ determine the relative importance of the objectives, while Q_e, Q_δ, and R are weighting matrices. The first two terms penalise deviations from the desired

temperature reference and internal thermal non-uniformity, respectively. The third penalises coolant utilisation, expressed as the fraction of the maximum available coolant flow,

$$\mathbf{1}^T u_{\mathrm{mod}}(k) = \frac{\sum_{j\in\{s,t,b\}} \dot{\mathcal{V}}_j(k)}{\dot{\mathcal{V}}_{\max}}, \tag{4.5}$$

where $\mathbf{1} = [1, 1, 1]^T$. The final term penalises variations in the valve duty-cycle commands, defined as

$$\Delta u_{\mathrm{mod}}(k+i) = u_{\mathrm{mod}}(k+i) - u_{\mathrm{mod}}(k+i-1). \tag{4.6}$$

Although temperature regulation and gradient minimisation can mitigate operating conditions associated with accelerated battery ageing, these thermal quantities do not directly quantify the degradation accumulated under a given operating trajectory.

Degradation-aware coolant allocation. Battery degradation depends on several interacting operating variables, including temperature, SoC, current, and accumulated charge throughput. Their nonlinear and history-dependent effects mean that minimising temperature or thermal gradients does not necessarily minimise battery degradation. Furthermore, spatial variations in these variables can lead to non-uniform ageing within the cell.

To address these limitations, Paper C extends the optimal coolant allocation framework by explicitly incorporating the spatially distributed degradation model developed in Chapter 3 into the control objective. The degradation-aware optimal coolant allocation problem (DOCAP) determines the coolant distribution that minimises the predicted increase in battery capacity loss while accounting for coolant utilisation and variations in the valve duty cycles. The objective function is formulated as

$$\begin{aligned}\min_{u_{\mathrm{mod},k}} J_{\mathrm{DOCAP}}(k) :=& \lambda_1 \left[\hat{\mathcal{Y}}_{\mathrm{deg}}(k+N_p) - \hat{\mathcal{Y}}_{\mathrm{deg}}(k)\right] \\ &+ \sum_{i=0}^{N_p-1} \left[\lambda_2 \left(\mathbf{1}^T u_{\mathrm{mod}}(k+i)\right)^2 + \lambda_3 \left\|\Delta u_{\mathrm{mod}}(k+i)\right\|^2_{R_\Delta}\right].\end{aligned} \tag{4.7}$$

where $\hat{\mathcal{Y}}^{\mathrm{dist}}_{\mathrm{deg}}$ denotes the predicted accumulated capacity loss. The weights λ_1, λ_2, and λ_3 determine the relative importance of degradation minimisation, coolant utilisation, and control variation, respectively. The first term

penalises the additional capacity loss predicted over the horizon. Through the dependence of degradation on the spatial distributions of temperature, SoC, and current, as well as the accumulated local charge throughputs, the coolant allocation is determined by the predicted ageing response rather than thermal performance alone. The second and third terms penalise coolant utilisation and variations in the valve duty-cycle commands, respectively, as in the OCAP formulation.

Both OCAP and DOCAP retain the same coolant allocation degrees of freedom and actuator constraints, and are implemented within the MPC framework presented in the following section.

4.2 Model predictive control

MPC is particularly suitable for the ITSC system because it accounts for the coupled battery–cooling dynamics, anticipates future operating conditions, and explicitly handles battery safety and actuator constraints. At each sampling instant k, the controller determines an optimal sequence of valve duty-cycle commands over the prediction horizon by minimising either the OCAP objective in (4.4) or the DOCAP objective in (4.7). The optimisation is subject to the discrete-time prediction model

$$\hat{X}(k+i+1) = F_d\Big(\hat{X}(k+i), U(k+i), W(k+i)\Big), \tag{4.8a}$$

$$\hat{Y}(k+i) = H_d\Big(\hat{X}(k+i), U(k+i), W(k+i)\Big), \tag{4.8b}$$

where F_d and H_d denote the discrete-time state-transition and output mappings of the coupled battery–cooling system (3.56). For the OCAP, the prediction model comprises the electrical, thermal, and coolant subsystems, whereas DOCAP additionally propagates the spatially distributed degradation states. The optimisation is further subject to the battery operating and actuator constraints

$$0 \le \hat{Z}(k+i) \le 1, \tag{4.9a}$$

$$\hat{T}_{\max}(k+i+1) \le T_{\max}^{\lim}, \tag{4.9b}$$

$$0 \le u_{\mathrm{mod},j}(k+i) \le 1, \quad j \in \{s,t,b\}, \tag{4.9c}$$

$$\sum_{j\in\{s,t,b\}} u_{\mathrm{mod},j}(k+i) \le 1, \tag{4.9d}$$

for $i = 0, \ldots, N_p - 1$. The first constraint restricts the predicted SoC to its admissible range, while the second limits the maximum predicted cell temperature to ensure thermally safe operation. The remaining constraints enforce the valve duty-cycle bounds and the maximum available coolant flow.

The optimisation is implemented in a receding-horizon manner. At each sampling instant, the optimal sequence $\{u^*_{\mathrm{mod}}(k+i)\}_{i=0}^{N_p-1}$ is computed, but only the first duty cycle command $u^*_{\mathrm{mod}}(k)$ is applied to the valve actuation system. At the next sampling instant, the prediction horizon is shifted forward and the optimisation is repeated using the updated state information.

The coupled battery–cooling dynamics are nonlinear, particularly due to the dependence of coolant heat transfer on the allocated flow rates and coolant temperatures. For DOCAP, additional nonlinearities arise from the dependence of the degradation dynamics on local temperature, SoC, current, and accumulated charge throughput. Solving the resulting nonlinear MPC (NMPC) problem to convergence at every sampling instant can be computationally tasking. To address this challenge, Papers B and C employ a real-time iteration MPC (RTI-MPC) approach, presented in the following section. The NMPC serves as a numerical benchmark for evaluating its control performance and computational efficiency.

Real-time iteration model predictive control

The RTI-MPC scheme reduces the computational burden of NMPC by locally approximating the nonlinear optimisation problem and solving a single quadratic programme (QP) at each sampling instant. The implementation consists of a warm-start step, a preparation phase, and a feedback phase. For brevity, the individual cooling-channel subscripts are omitted in the following formulation, with u_{mod} representing the vector of valve duty cycles for all three cooling channels.

Warm-start step. At sampling instant k, the optimal valve duty-cycle sequence obtained at the preceding sampling instant is shifted forward by one step to initialise the nominal control trajectory,

$$\tilde{u}_{\mathrm{mod}}(k+i) = u^{*,\mathrm{prev}}_{\mathrm{mod}}(i+1), \quad i = 0, \ldots, N_p - 2, \tag{4.10a}$$

$$\tilde{u}_{\mathrm{mod}}(k+N_p-1) = u^{*,\mathrm{prev}}_{\mathrm{mod}}(N_p-1). \tag{4.10b}$$

The corresponding nominal coolant flow trajectory $\tilde{U}(k+i)$ is obtained using the valve actuation model. The nonlinear prediction model is then propagated over the horizon to obtain the nominal state and output trajectories,

$$\tilde{X}(k+i+1) = F_d\left(\tilde{X}(k+i), \tilde{U}(k+i), W(k+i)\right), \tag{4.11a}$$

$$\tilde{Y}(k+i) = H_d\left(\tilde{X}(k+i), \tilde{U}(k+i), W(k+i)\right). \tag{4.11b}$$

For OCAP, the nominal trajectories comprise the electrical, thermal, and coolant states, whereas DOCAP additionally includes the spatially distributed degradation states.

Preparation phase. The nonlinear prediction dynamics are locally linearised about the nominal state and control trajectories. Defining the perturbations

$$\delta\hat{X}(k+i) = \hat{X}(k+i) - \tilde{X}(k+i), \tag{4.12a}$$

$$\delta U(k+i) = U(k+i) - \tilde{U}(k+i), \tag{4.12b}$$

$$\delta u_{\mathrm{mod}}(k+i) = u_{\mathrm{mod}}(k+i) - \tilde{u}_{\mathrm{mod}}(k+i), \tag{4.12c}$$

a first-order Taylor expansion of the discrete-time prediction model yields

$$\delta\hat{X}(k+i+1) = A_i\delta\hat{X}(k+i) + B_i\delta U(k+i), \tag{4.13}$$

with the Jacobian matrices

$$A_i = \left.\frac{\partial F_d}{\partial X}\right|_{\tilde{\xi}_i}, \qquad B_i = \left.\frac{\partial F_d}{\partial U}\right|_{\tilde{\xi}_i}, \tag{4.14}$$

and

$$\tilde{\xi}_i = \left(\tilde{X}(k+i), \tilde{U}(k+i), W(k+i)\right). \tag{4.15}$$

The corresponding output mapping is linearised as

$$\delta\hat{Y}(k+i) = C_i\delta\hat{X}(k+i) + D_i\delta U(k+i), \tag{4.16}$$

where C_i and D_i are the Jacobians of H_d with respect to the states and inputs, respectively, evaluated along the nominal trajectory. The coolant flow perturbations are related to the valve duty-cycle perturbations through

$$\delta U(k+i) = \dot{\mathcal{V}}_{\max}\delta u_{\mathrm{mod}}(k+i). \tag{4.17}$$

The local approximations transform the nonlinear coolant allocation problem into a QP, with the duty-cycle correction sequence

$$\delta u_{\mathrm{mod},k} := \{\delta u_{\mathrm{mod}}(k+i)\}_{i=0}^{N_p-1} \tag{4.18}$$

as the decision variable. The QP is expressed as

$$\begin{aligned} \min_{\delta u_{\mathrm{mod},k}} \quad & J_{\mathrm{QP}}(k) \\ \text{s.t.} \quad & \text{linearised prediction dynamics,} \\ & \text{battery operating constraints,} \\ & 0 \le \tilde{u}_{\mathrm{mod}}(k+i) + \delta u_{\mathrm{mod}}(k+i) \le 1, \\ & \mathbf{1}^T \left(\tilde{u}_{\mathrm{mod}}(k+i) + \delta u_{\mathrm{mod}}(k+i)\right) \le 1, \end{aligned} \tag{4.19}$$

for $i = 0, \ldots, N_p - 1$, where J_{QP} denotes the local quadratic approximation of the corresponding OCAP or DOCAP objective.

Feedback phase. Once updated state information becomes available, the initial state perturbation is determined as $\delta\hat{X}(k) = X(k) - \tilde{X}(k)$. The prepared QP is solved once to obtain the optimal duty-cycle correction sequence. The first corrected command,

$$u^*_{\mathrm{mod}}(k) = \tilde{u}_{\mathrm{mod}}(k) + \delta u^*_{\mathrm{mod}}(k), \tag{4.20}$$

is applied to the valve actuation system. The remaining corrected commands are stored and shifted to initialise the optimisation at the next sampling instant.

By performing one local linearisation and solving a single QP per sampling instant, RTI-MPC reduces the computational effort required for online coolant allocation while retaining the receding-horizon and constraint-handling capabilities of NMPC.

CHAPTER 5

Summary of included papers

This chapter provides a summary of the included papers.

5.1 Paper A

Godwin K. Peprah, Yicun Huang, Torsten Wik, Faisal Altaf, Changfu Zou
Thermal Modeling of Battery Cells for Optimal Tab and Surface Cooling Control
Published in IEEE Transactions on Control Systems Technology, vol. 34, no. 4, pp. 1831–1844, July 2026.
doi: 10.1109/TCST.2026.3675726.

© 2026 The Authors. This work is licensed under a Creative Commons Attribution 4.0 License .

Paper A develops a control-oriented thermal modelling framework for battery cells of different geometries by combining the Chebyshev spectral-Galerkin method with model component decomposition. The resulting library of reduced-order thermal models can be configured with different numbers of states,

thereby allowing a trade-off between computational complexity and spatial temperature prediction accuracy. The models are validated against a high-fidelity finite element model and benchmarked against a conventional thermal equivalent circuit (TEC) model under realistic driving and cooling conditions. A four-state model is shown to accurately reproduce the two-dimensional thermal dynamics, while even a one-state model outperforms the commonly used two-state TEC model in both prediction accuracy and computational efficiency. The framework also enables independent manipulation of tab and surface cooling boundary conditions. Its versatility is demonstrated through the evaluation of different cooling configurations, closed-loop temperature control, and assessment of cell aspect ratio on thermal behaviour. The results show that surface cooling is generally more effective at reducing the mean cell temperature, whereas tab cooling promotes a more uniform internal temperature distribution, particularly when cooling is applied symmetrically through both tabs. Furthermore, analysis reveals that the thermal performance of a given cell geometry depends strongly on the cooling configuration. For cylindrical cells, shorter and wider geometries favour tab-involved cooling due to shorter axial heat-conduction paths, whereas taller and slimmer geometries favour surface cooling through increased lateral heat-transfer area. Consequently, no single cell aspect ratio is universally optimal, highlighting the need to co-design battery cell geometry and the intended cooling strategy rather than optimising them independently.

Godwin K. Peprah led the development and implementation of the thermal model, model validation, simulation studies, analysis of results, and preparation of the manuscript. Torsten Wik contributed to the conceptualisation of the work, modelling methodology, interpretation of the results, supervision, and revision of the manuscript. Yicun Huang contributed to the thermal modelling methodology and technical discussions. Faisal Altaf contributed application-oriented technical discussions and revision of the manuscript.
Changfu Zou contributed to the conceptualisation of the work, modelling methodology, interpretation of the results, supervision, and revision of the manuscript.

5.2 Paper B

Godwin K. Peprah, Torsten Wik, Masood Tamadondar, Changfu Zou
Model Predictive Coolant Allocation for Integrated Tab–Surface Cooling of Battery Cells
Manuscript, 2026. Submitted to IEEE Transactions on Control Systems Technology for review. Available as *arXiv preprint arXiv: 2607.10872.*
The version reproduced in this thesis is the manuscript dated 12 July 2026.
A later revised or published version may differ.

© 2026 The Authors .

Paper B develops an integrated tab-surface cooling (ITSC) system in which a finite coolant resource is dynamically allocated among the lateral surface, top, and bottom cooling channels of a battery cell. A first-principles dynamic coolant model is developed and coupled with the battery electrical and spatial thermal models, while a valve-actuation model captures the dynamic routing of coolant among the channels. The coolant allocation problem is formulated as an optimal control problem that regulates battery temperature towards a desired reference while suppressing internal thermal gradients. To enable real-time implementation, the resulting nonlinear optimisation problem is solved using a real-time iteration model predictive control (RTI-MPC) scheme and benchmarked against nonlinear (N)MPC. Under realistic driving conditions, RTI-MPC closely reproduces the NMPC thermal response, with absolute temperature errors below 0.0035 °C and a mean online computation time of 19.3 ms, demonstrating its potential for real-time implementation. Comparisons with conventional cooling configurations show that ITSC achieves the lowest thermal gradients and temperature tracking errors among the feedback-controlled configurations. In particular, it reduces the maximum cell temperature to 39.47 °C, compared with 47.66 °C under surface cooling, while achieving a substantially more uniform temperature distribution. Although a fixed equal-split cooling strategy achieves lower temperatures through continuous full-flow cooling, it produces larger thermal gradients and lacks adaptive coolant allocation. These findings demonstrate that dynamically combining tab and surface cooling provides a more effective balance between temperature regulation and thermal uniformity than conventional fixed cooling configurations.

Godwin K. Peprah led the development of the integrated tab-surface cooling framework, coolant and actuator modelling, formulation and implementation of the RTI-MPC controller, simulations, analysis of results, and preparation of the manuscript. Torsten Wik contributed to the modelling and control methodology, interpretation of the results, supervision, and revision of the manuscript. Masood Tamadondar contributed technical discussions concerning battery cooling systems and their practical implementation and reviewed the manuscript. Changfu Zou contributed to the conceptualisation of the work, modelling and control methodology, supervision, interpretation of the results, and revision of the manuscript.

5.3 Paper C

Godwin K. Peprah, Phil Kollmeyer, Torsten Wik, Changfu Zou
Degradation-Aware Model Predictive Coolant Allocation for Integrated Tab-Surface Cooling of Battery Cells
Manuscript. To be submitted to IEEE Transactions on Control Systems Technology.
The version reproduced in this thesis is the manuscript dated 20 November 2026.
A later submitted or published version may differ.

© 2026 The Authors .

Paper C extends the ITSC framework to explicitly account for battery degradation in the coolant allocation objective. A coupled spatially distributed electrothermal–degradation model is developed to capture the non-uniform distributions of temperature, current density, state of charge (SoC), and capacity loss within battery cells. The conventional equivalent circuit model is extended to a distributed formulation, while the degradation mechanisms are reformulated as dynamic states and evaluated locally. The resulting model captures the coupled effects of thermal non-uniformity, current redistribution, and spatially varying degradation. The model is evaluated for two cylindrical LFP cells with capacities of 2.3 Ah and 45 Ah. The results reveal that spatial temperature differences induce non-uniform current distributions, local SoC variations, and uneven degradation, including internal current redistribution during rest periods despite zero net battery current. These effects

become more pronounced in larger cells, where local degradation differences are substantially greater even when spatially averaged capacity losses are comparable. For instance, under surface cooling at 3C, the 45 Ah cell exhibits an internal temperature difference approaching 9 °C, compared with approximately 2 °C for the smaller cell. These findings highlight the importance of accounting for local electrothermal behaviour when predicting battery degradation, particularly in larger-format cells. A degradation-aware optimal coolant allocation problem (DOCAP) is subsequently formulated to minimise the predicted increase in capacity loss while accounting for coolant utilisation and control variation. The resulting nonlinear optimisation problem is solved using RTI-MPC, reducing the mean online execution time from 4.6 s for NMPC to 0.1 s, corresponding to a 46-fold speed-up, while maintaining comparable closed-loop performance. Evaluations under repeated driving and charging sequences demonstrate that ITSC achieves the lowest accumulated capacity loss among the feedback-controlled cooling configurations. Compared with continuous equal-split cooling, ITSC achieves comparable degradation performance while reducing coolant utilisation by 85% and 62.5% for the 2.3 Ah and 45 Ah cells, respectively. Furthermore, comparison with the temperature- and gradient-based coolant allocation problem (OCAP) reveals that minimising thermal quantities does not necessarily minimise battery degradation. Under comparable coolant utilisation, DOCAP achieves approximately 0.40% lower accumulated capacity loss than OCAP, despite the latter maintaining a lower mean temperature and smaller spatial thermal gradients. These findings demonstrate the importance of accounting for the competing degradation mechanisms when optimally allocating coolant among complementary cooling interfaces.

Godwin K. Peprah led the development and implementation of the spatially distributed electrothermal–degradation modelling framework, including the extension of the electrical equivalent circuit and degradation models to account for local temperature, current, SoC, and capacity loss distributions. He also led the formulation and implementation of the degradation-aware RTI-MPC controller, simulation studies, analysis and interpretation of results, and preparation of the manuscript. Torsten Wik contributed to the degradation and control methodology, interpretation of the results, supervision, and revision of the manuscript. Phil Kollmeyer contributed battery degradation and experimental perspectives, technical discussions, interpretation of the results,

and revision of the manuscript. Changfu Zou contributed to the conceptualisation of the work, modelling and control methodology, supervision, interpretation of the results, and revision of the manuscript.

CHAPTER 6

Concluding remarks and future work

6.1 Conclusions

This thesis has developed a unified modelling and predictive control framework for battery thermal management, integrating spatial thermal modelling, dynamic coolant allocation, real-time optimisation, and battery degradation. Building on the research objectives and contributions outlined in Sections 1.2–1.3, the work addresses not only how heat is removed from battery cells, but also how the available cooling resource should be optimally distributed in response to evolving electrothermal conditions and their influence on battery health. The main conclusions, corresponding to the four research questions are summarised below.

RQ1: Control-oriented spatial thermal modelling

First, a control-oriented spatial thermal modelling framework was developed by combining the Chebyshev spectral-Galerkin method with model component decomposition. The resulting library of reduced-order models enables independent control of cooling applied to different cell interfaces while accurately

predicting the spatial temperature distribution within the cell. The framework was demonstrated for both cylindrical and pouch cells under various combinations of tab and surface cooling and realistic driving conditions. The results showed that accurate spatial temperature predictions can be achieved with a small number of states, while maintaining computational requirements comparable to, and in some cases lower than, conventional thermal equivalent circuit (TEC) models. The developed framework provides a viable alternative to TEC models for model-based battery management applications requiring spatial thermal information.

The utility of the developed models was further demonstrated through cooling configuration and cell geometry studies. The results revealed that thermal performance depends strongly on the interaction between cell geometry and the applied cooling configuration. For cylindrical cells, lower aspect ratios, corresponding to shorter and wider geometries, favour tab-involved cooling by reducing the axial heat-conduction distance, whereas surface-only cooling benefits from higher aspect ratios due to the increased lateral heat-transfer area. Consequently, no single cell aspect ratio is universally optimal, highlighting the need to co-design battery cell geometry and the intended cooling strategy rather than optimising them independently.

Beyond battery applications, the combination of spectral-Galerkin discretisation and model component decomposition offers potential applicability to other distributed-parameter systems governed by partial differential equations with non-homogeneous boundary conditions.

RQ2: Integration of tab and surface cooling

Building on the developed thermal models, the complementary characteristics of tab and lateral surface cooling were combined through the integrated tab–surface cooling (ITSC) architecture. Surface cooling is generally more effective at suppressing the overall temperature rise, whereas tab cooling provides greater authority in shaping the internal temperature uniformity. Rather than selecting either strategy a priori, the proposed approach formulates their integration as a constrained coolant allocation problem, allowing the available coolant resource to be dynamically distributed among the cooling interfaces according to the evolving battery operating conditions.

Evaluation against conventional cooling configurations demonstrated that ITSC achieves the lowest thermal gradients and temperature tracking er-

rors among the feedback-controlled configurations. The controller predominantly favours the tab cooling channels to suppress internal temperature non-uniformity, while retaining the surface channel as an additional degree of freedom for temperature regulation. These findings demonstrate that dynamically combining complementary cooling interfaces provides a more effective balance between temperature regulation and thermal uniformity than conventional fixed cooling configurations.

RQ3: Real-time predictive coolant allocation

The predictive control framework was developed by coupling the battery electrical and spatial thermal models with the dynamics of the coolant and valve actuation system. The resulting nonlinear battery–cooling model captures the interaction between the allocated coolant flow rates, the thermal boundary conditions of the battery, and the evolving temperature distribution that governs heat transfer to the coolant. This enabled the formulation of a constrained model predictive control (MPC) problem for dynamically allocating the available coolant among the cooling interfaces.

To address the computational demands of nonlinear (N)MPC, a real-time iteration (RTI)-MPC approach was developed. By locally linearising the nonlinear prediction dynamics and solving a single quadratic programme at each sampling instant, RTI-MPC closely reproduced the closed-loop behaviour of NMPC while substantially reducing the online computational burden. In Paper B, the mean execution time was reduced to approximately 19.3 ms for the temperature- and gradient-based controller, while Paper C achieved a mean execution time of 0.1 s for the more complex degradation-aware controller, corresponding to a 46-fold speed-up relative to NMPC. Both implementations satisfied the prescribed 1 s sampling interval under the evaluated conditions, demonstrating that the additional control authority of ITSC can be exploited without sacrificing the computational requirements imposed by real-time battery thermal management.

RQ4: Degradation-aware predictive cooling

The final part of the thesis extended the predictive cooling framework beyond conventional temperature regulation and thermal gradient minimisation by explicitly incorporating battery degradation into the control objective. A semi-

empirical degradation model was reformulated as dynamic states and coupled with spatially distributed electrical and thermal models to capture local variations in temperature, current density, SoC, and capacity loss. The resulting framework revealed that thermal non-uniformity induces current redistribution and local SoC differences, including internal currents during rest periods despite zero net battery current. These effects become more pronounced in larger cells, leading to substantially greater spatial variations in degradation even when the spatially averaged capacity losses remain comparable. These findings highlight the importance of accounting for local electrothermal behaviour when modelling and controlling battery degradation.

Building on this framework, the degradation-aware optimal coolant allocation problem (DOCAP) was formulated to minimise the predicted increase in capacity loss while accounting for coolant utilisation and actuator constraints. Evaluation against conventional cooling configurations demonstrated that ITSC achieves the lowest accumulated capacity loss among the feedback-controlled configurations, while attaining comparable degradation performance to continuous equal-split cooling with substantially lower coolant utilisation. These results demonstrate the benefits of dynamically allocating the available coolant according to its predicted influence on battery degradation rather than relying on fixed cooling configurations.

Importantly, the degradation-aware formulation revealed a fundamental distinction between thermal performance and battery health optimisation. Although lower temperatures and smaller thermal gradients are generally desirable, their minimisation does not necessarily minimise accumulated capacity loss. The individual degradation mechanisms depend jointly on temperature, SoC, current, and accumulated charge throughput, and exhibit different and sometimes competing sensitivities to these operating variables. Consequently, a cooling action that improves a prescribed thermal metric may provide little degradation benefit or adversely affect another degradation mechanism. Under comparable coolant utilisation, DOCAP achieved a relative reduction of approximately 0.40% in accumulated capacity loss compared with the thermal-proxy controller (OCAP), despite the latter maintaining a lower mean temperature and smaller spatial thermal gradients. These findings demonstrate that direct degradation minimisation can lead to fundamentally different and more effective coolant allocation decisions for preserving battery health than conventional thermal-proxy control.

6.2 Synthesis and broader implications

Taken together, the findings of this thesis demonstrate that effective battery thermal management requires more than simply increasing heat rejection or minimising temperature. The three papers establish a unified framework connecting spatial thermal modelling, the integration of complementary cooling interfaces, real-time predictive control, and degradation-aware optimisation. The central finding is that cell geometry, cooling architecture, and control objectives are closely interdependent and should therefore be considered jointly rather than as separate design problems.

Relative to conventional approaches, the framework combines spatially resolved battery behaviour with dynamic coolant allocation and explicit consideration of local degradation effects. This enables cooling decisions to account for internal thermal non-uniformity and its consequences, rather than relying solely on average temperature or prescribed thermal metrics. Importantly, the results demonstrate that improvements in thermal performance do not necessarily translate into reduced degradation, highlighting the importance of directly incorporating battery health into the control objective.

From a practical perspective, integrating multiple cooling interfaces introduces additional design considerations, including cooling-system complexity and packaging requirements. Nevertheless, conventional single-interface cooling may be insufficient to simultaneously achieve effective temperature regulation and thermal uniformity, particularly under demanding operating conditions. The complementary cooling characteristics demonstrated in this thesis suggest that a suitably designed ITSC system could offer a more favourable balance between thermal performance, cooling resource utilisation, and battery lifetime. Furthermore, the computational efficiency of RTI-MPC provides a promising basis for real-time implementation, although system-level benefits and practical feasibility remain to be experimentally investigated.

6.3 Future work

Although the proposed framework demonstrates promising results, several opportunities remain for extending its scope and practical applicability. The present work has primarily focused on cell-level modelling and control, with the complete integrated cooling architecture evaluated numerically under pre-

scribed operating conditions. Further research could address the following directions:

- **Extension to module- and pack-level systems.** The proposed framework could be extended to battery modules and packs by incorporating cell-to-cell thermal interactions, non-uniform operating conditions, and the coupling introduced by shared coolant networks. Such an extension would enable assessment of the scalability of ITSC and the trade-offs between cooling system complexity, packaging requirements, and battery performance.
- **Joint optimisation of coolant allocation and pump operation.** The present formulation distributes the available coolant among the cooling interfaces without explicitly modelling pump operation and its associated energy consumption. Future work could incorporate total coolant flow and pump power into the optimisation problem, allowing both the magnitude and spatial distribution of cooling effort to be optimised while balancing thermal performance, degradation, and auxiliary energy consumption.
- **Vehicle-level integrated thermal management.** A further extension is to integrate the proposed battery cooling framework with the vehicle heating, ventilation, and air-conditioning (HVAC) system. Such an approach would enable coordinated management of battery and cabin thermal demands, accounting for their shared thermal resources and competing energy requirements. A system-level predictive control strategy could optimise battery temperature regulation, cabin thermal comfort, and auxiliary energy consumption while considering the implications for battery degradation and vehicle driving range.
- **Degradation-aware fast charging and thermal management.** The proposed framework could be extended to jointly optimise charging current and coolant allocation during fast charging. By incorporating spatially distributed electrothermal and degradation dynamics, such a strategy could balance charging time, temperature regulation, thermal uniformity, and battery lifetime while enforcing local operating constraints. This would enable the cooling system to adapt to the evolving charging conditions and potentially reduce degradation associated with aggressive charging.

- **Experimental validation and real-time implementation.** Experimental investigation of the integrated cooling architecture is necessary to assess the predicted thermal and degradation benefits under practical operating conditions. Implementing RTI-MPC on representative embedded hardware would further enable evaluation of computational feasibility and robustness to model uncertainty, measurement noise, parameter variations, and practical actuator limitations.

Code availability

The code developed for the modelling, simulation, and predictive control frameworks presented in this thesis is available from the author upon reasonable request.

References

[1] International Energy Agency, *The energy crisis creates even stronger impetus for EU electrification*, IEA Commentary, 2026.

[2] International Energy Agency, "Global EV outlook 2026," International Energy Agency, Paris, Tech. Rep., 2026.

[3] T. E. Jefferson, S. Rana, and N. Banerjee, "Lithium-ion batteries: A review of materials, design advances, sustainability challenges, and future prospects," *Energy Storage*, vol. 8, no. 1, e70320, 2026.

[4] G. Sankaran and S. Venkatesan, "An overview of lithium-ion batteries for electric mobility and energy storage applications," in *IOP Conference Series: Earth and Environmental Science*, IOP Publishing, vol. 1042, 2022, p. 012 012.

[5] S. Ma et al., "Temperature effect and thermal impact in lithium-ion batteries: A review," *Prog. Nat. Sci.: Mater. Int.*, vol. 28, no. 6, pp. 653–666, 2018.

[6] C. Bibin, M. Vijayaram, V. Suriya, R. S. Ganesh, and S. Soundarraj, "A review on thermal issues in Li-ion battery and recent advancements in battery thermal management system," *Mater. Today Proc.*, vol. 33, pp. 116–128, 2020.

[7] P. Shrivastava and A. Tiwari, "Advances in battery thermal management systems and energy storage technologies," *Int. Commun. Heat Mass Transf.*, vol. 175, p. 111 050, 2026.

[8] Y. Zhao, L. B. Diaz, Y. Patel, T. Zhang, and G. J. Offer, "How to cool lithium ion batteries: Optimising cell design using a thermally coupled model," *J. Electrochem. Soc.*, vol. 166, no. 13, A2849–A2859, 2019.

[9] Y. Zhao, Y. Patel, T. Zhang, and G. J. Offer, "Modeling the effects of thermal gradients induced by tab and surface cooling on lithium ion cell performance," *J. Electrochem. Soc.*, vol. 165, no. 13, A3169–A3178, 2018.

[10] O. Dondelewski et al., "The role of cell geometry when selecting tab or surface cooling to minimise cell degradation," *eTransportation*, vol. 5, p. 100 073, 2020.

[11] S. Li et al., "Optimal cell tab design and cooling strategy for cylindrical lithium-ion batteries," *J. Power Sources*, vol. 492, p. 229 594, 2021.

[12] C. Bolsinger and K. P. Birke, "Effect of different cooling configurations on thermal gradients inside cylindrical battery cells," *J. Energy Storage*, vol. 21, pp. 222–230, 2019.

[13] T. Ahmad, A. Mishra, S. Ghosh, and C. S. Casari, "Identifying efficient cooling approach of cylindrical lithium-ion batteries," *Energy Technol.*, vol. 10, no. 2, p. 2 100 888, 2022.

[14] X. Lin et al., "Online parameterization of lumped thermal dynamics in cylindrical lithium ion batteries for core temperature estimation and health monitoring," *IEEE Trans. Control Syst. Technol.*, vol. 21, no. 5, pp. 1745–1755, 2013.

[15] X. Lin et al., "A lumped-parameter electro-thermal model for cylindrical batteries," *J. Power Sources*, vol. 257, pp. 1–11, 2014.

[16] G. K. Peprah, F. Liberati, F. Altaf, G. Osei-Dadzie, A. Di Giorgio, and A. Pietrabissa, "Optimal load sharing in reconfigurable battery systems using an improved model predictive control method," in *Proc. 29th Mediterranean Conf. Control Autom. (MED)*, IEEE, 2021, pp. 979–984.

[17] F. P. Incropera and D. P. DeWitt, *Fundamentals of Heat and Mass Transfer*. New York, NY: Wiley, 1985.

[18] U. S. Kim, J. Yi, C. B. Shin, T. Han, and S. Park, "Modelling the thermal behaviour of a lithium-ion battery during charge," *J. Power Sources*, vol. 196, no. 11, pp. 5115–5121, 2011.

[19] R. R. Richardson, S. Zhao, and D. A. Howey, "On-board monitoring of 2-D spatially-resolved temperatures in cylindrical lithium-ion batteries: Part i. low-order thermal modelling," *J. Power Sources*, vol. 326, pp. 377–388, 2016.

[20] X. Hu, W. Liu, X. Lin, Y. Xie, A. M. Foley, and L. Hu, "A control-oriented electrothermal model for pouch-type electric vehicle batteries," *IEEE Trans. Power Electron.*, vol. 36, no. 5, pp. 5530–5544, 2020.

[21] L. Magni, D. M. Raimondo, and F. Allgöwer, *Nonlinear model predictive control*. Springer, 2009, vol. 384.

[22] Y.-G. Xi, D.-W. Li, and S. Lin, "Model predictive control-status and challenges," *Acta Autom. Sin.*, vol. 39, no. 3, pp. 222–236, 2013.

[23] Y. Kim, "Critical review of optimal control methods for Li-Ion batteries in electric vehicles," *ChemElectroChem*, vol. 11, no. 3, e202300497, 2024.

[24] Z. Lu, H. Tu, H. Fang, Y. Wang, and S. Mou, "Integrated optimal fast charging and active thermal management of lithium-ion batteries in extreme ambient temperatures," *IEEE Trans. Control Syst. Technol.*, vol. 33, no. 2, pp. 714–728, 2025.

[25] F. Altaf, B. Egardt, and L. J. Mårdh, "Load management of modular battery using model predictive control: Thermal and state-of-charge balancing," *IEEE Trans. Control Syst. Technol.*, vol. 25, no. 1, pp. 47–62, 2016.

[26] J. S. Edge et al., "Lithium ion battery degradation: What you need to know," *Phys. Chem. Chem. Phys.*, vol. 23, no. 14, pp. 8200–8221, 2021.

[27] M. Schimpe, M. E. von Kuepach, M. Naumann, H. C. Hesse, K. Smith, and A. Jossen, "Comprehensive modeling of temperature-dependent degradation mechanisms in lithium iron phosphate batteries," *J. Electrochem. Soc.*, vol. 165, no. 2, A181–A193, 2018.

[28] W. Vermeer, G. R. C. Mouli, and P. Bauer, "A comprehensive review on the characteristics and modeling of lithium-ion battery aging," *IEEE Trans. Transp. Electrific.*, vol. 8, no. 2, pp. 2205–2232, 2022.

[29] J. B. Goodenough and Y. Kim, "Challenges for rechargeable Li batteries," *Chem. Mater.*, vol. 22, no. 3, pp. 587–603, 2010.

[30] N. Nitta, F. Wu, J. T. Lee, and G. Yushin, "Li-ion battery materials: Present and future," *Mater. Today*, vol. 18, no. 5, pp. 252–264, 2015.

[31] R. Xiong, *Battery management algorithm for electric vehicles*. Springer, 2020.

[32] J. Jaguemont, L. Boulon, and Y. Dubé, "A comprehensive review of lithium-ion batteries used in hybrid and electric vehicles at cold temperatures," *Appl. Energy*, vol. 164, pp. 99–114, 2016.

[33] X. Lin, K. Khosravinia, X. Hu, J. Li, and W. Lu, "Lithium plating mechanism, detection, and mitigation in lithium-ion batteries," *Prog. Energy Combust. Sci.*, vol. 87, p. 100 953, 2021.

[34] E. Peled and S. Menkin, "SEI: Past, present and future," *J. Electrochem. Soc.*, vol. 164, no. 7, A1703, 2017.

[35] X. Feng, M. Ouyang, X. Liu, L. Lu, Y. Xia, and X. He, "Thermal runaway mechanism of lithium ion battery for electric vehicles: A review," *Energy Storage Mater.*, vol. 10, pp. 246–267, 2018.

[36] T. G. Tranter, R. Timms, P. R. Shearing, and D. Brett, "Communication—prediction of thermal issues for larger format 4680 cylindrical cells and their mitigation with enhanced current collection," *J. Electrochem. Soc.*, vol. 167, no. 16, p. 160 544, 2020.

[37] A. Samba, N. Omar, H. Gualous, O. Capron, P. Van den Bossche, and J. Van Mierlo, "Impact of tab location on large format lithium-ion pouch cell based on fully coupled three-dimensional electrochemical-thermal modeling," *Electrochim. Acta*, vol. 147, pp. 319–329, 2014.

[38] J.-J. Lee, J.-S. Kim, D.-C. Lee, H. Chang, and C.-W. Kim, "Design optimization of tab attachment positions and cell aspect ratio to minimize temperature difference in 45-Ah LFP large-format lithium-ion pouch cells," *Appl. Therm. Eng.*, vol. 182, p. 116 143, 2021.

[39] H. Lundgren et al., "Thermal management of large-format prismatic lithium-ion battery in PHEV application," *J. Electrochem. Soc.*, vol. 163, no. 2, A309–A317, 2016.

[40] K. Liu et al., "Electrochemical modeling and parameterization towards control-oriented management of lithium-ion batteries," *Control Eng. Pract.*, vol. 124, p. 105 176, 2022.

[41] S. Nejad, D. Gladwin, and D. Stone, "A systematic review of lumped-parameter equivalent circuit models for real-time estimation of lithium-ion battery states," *J. Power Sources*, vol. 316, pp. 183–196, 2016.

[42] V. Lucaferri, M. Quercio, A. Laudani, and F. Riganti Fulginei, "A review on battery model-based and data-driven methods for battery management systems," *Energies*, vol. 16, no. 23, p. 7807, 2023.

[43] X. Hu, S. Li, and H. Peng, "A comparative study of equivalent circuit models for Li-ion batteries," *J. Power Sources*, vol. 198, pp. 359–367, 2012.

[44] J. Liu, S. Yadav, M. Salman, S. Chavan, and S. C. Kim, "Review of thermal coupled battery models and parameter identification for lithium-ion battery heat generation in EV battery thermal management system," *Int. J. Heat Mass Transf.*, vol. 218, p. 124 748, 2024.

[45] W. Allafi et al., "A lumped thermal model of lithium-ion battery cells considering radiative heat transfer," *Appl. Therm. Eng.*, vol. 143, pp. 472–481, 2018.

[46] D. W. Hahn and M. N. Özisik, *Heat conduction*. John Wiley & Sons, 2012.

[47] M. Astaneh et al., "Calibration optimization methodology for lithium-ion battery pack model for electric vehicles in mining applications," *Energies*, vol. 13, no. 14, p. 3532, 2020.

[48] J. Shen, T. Tang, and L.-L. Wang, *Spectral methods: algorithms, analysis and applications*. Springer Science & Business Media, 2011, vol. 41.

[49] L. N. Trefethen, *Spectral methods in MATLAB*. SIAM, 2000.

[50] X. Jin et al., "Applicability of available Li-ion battery degradation models for system and control algorithm design," *Control Eng. Pract.*, vol. 71, pp. 1–9, 2018.

[51] T. Mesbahi, N. Rizoug, P. Bartholomeus, R. Sadoun, F. Khenfri, and P. L. Moigne, "Dynamic model of Li-ion batteries incorporating electrothermal and ageing aspects for electric vehicle applications," *IEEE Trans. Ind. Electron.*, vol. 65, no. 2, pp. 1298–1305, Feb. 2018.

[52] M. Taghizadeh, A. Ghaffari, and F. Najafi, "Modeling and identification of a solenoid valve for PWM control applications," *C. R. Mecanique*, vol. 337, no. 3, pp. 131–140, 2009.

[53] M. K. Kazimierczuk, *Pulse-Width Modulated DC-DC Power Converters*. John Wiley & Sons, 2015.

[54] Y. Xie, S. Shi, J. Tang, H. Wu, and J. Yu, "Experimental and analytical study on heat generation characteristics of a lithium-ion power battery," *Int. J. Heat Mass Transf.*, vol. 122, pp. 884–894, 2018.

Part II

Papers

PAPER A

Thermal Modeling of Battery Cells for Optimal Tab and Surface Cooling Control

Godwin K. Peprah, Yicun Huang, Torsten Wik, Faisal Altaf, Changfu Zou

Published in IEEE Transactions on Control Systems Technology, vol. 34, no. 4, pp. 1831–1844, July 2026.
doi: 10.1109/TCST.2026.3675726.

© 2026 The Authors. This work is licensed under a Creative Commons Attribution 4.0 License

The layout has been revised.

Correction note

Paper A: G. K. Peprah, Y. Huang, T. Wik, F. Altaf, and C. Zou, "Thermal Modeling of Battery Cells for Optimal Tab and Surface Cooling Control," *IEEE Trans. Control Syst. Technol.*, vol. 34, no. 4, pp. 1831–1844, 2026.
DOI: `https://doi.org/10.1109/TCST.2026.3675726`

The published version of this paper contains sign errors in the conductive heat-flux terms of the thermal boundary conditions in Eqs. (2), (4), (7), (11), (13), and (42). These errors have been corrected in the version reproduced in this thesis, following Fourier's law of heat conduction and the outward-normal sign convention,

$$-k\nabla T \cdot \mathbf{n} = h(T - T_\infty),$$

where $\mathbf{n}$ denotes the outward unit normal vector.

For the cylindrical cell, the corrected boundary conditions in Eq. (2) are

$$\begin{aligned}
-k_r \frac{\partial T}{\partial r} &= h_s(T - T_{s,\infty}), \quad r = R_{\text{out}}, \\
k_r \frac{\partial T}{\partial r} &= h_c(T - T_{c,\infty}), \quad r = R_{\text{in}}, \\
-k_z \frac{\partial T}{\partial z} &= h_t(T - T_{t,\infty}), \quad z = L, \\
k_z \frac{\partial T}{\partial z} &= h_b(T - T_{b,\infty}), \quad z = 0.
\end{aligned}$$

Accordingly, the corrected transformed boundary conditions in Eq. (7) are

$$\begin{aligned}
h_s T + \alpha k_r \frac{\partial T}{\partial \tilde{r}} &= u_s, \quad \tilde{r} = 1, \\
-h_c T + \alpha k_r \frac{\partial T}{\partial \tilde{r}} &= u_c, \quad \tilde{r} = -1, \\
h_t T + \beta k_z \frac{\partial T}{\partial \tilde{z}} &= u_t, \quad \tilde{z} = 1, \\
-h_b T + \beta k_z \frac{\partial T}{\partial \tilde{z}} &= u_b, \quad \tilde{z} = -1.
\end{aligned}$$

The corresponding sign corrections have also been applied to the pouch-cell boundary conditions in Eqs. (4) and (42), and the homogeneous and particular-solution boundary conditions in Eqs. (11) and (13), respectively.

The numerical implementation employed the correct sign convention, and the reported simulation results and conclusions remain unchanged.

Abstract

Optimal cooling that minimises thermal gradients and the average temperature is essential for enhanced battery safety and health. This work presents a new modelling approach for battery cells of different shapes by integrating the Chebyshev spectral-Galerkin method and model component decomposition. As a result, a library of reduced-order, computationally efficient battery thermal models is obtained, characterised by different numbers of states. These models are validated against a high-fidelity finite element model and are compared with a thermal equivalent circuit (TEC) model under real-world vehicle driving and battery cooling scenarios. Illustrative results demonstrate that the proposed model with four states can faithfully capture the two-dimensional thermal dynamics, while the model with only one state significantly outperforms the widely used two-state TEC model in both accuracy and computational efficiency, reducing computation time by 28.7%. Furthermore, our developed models allow for independent control of tab and surface cooling channels, enabling effective thermal performance optimisation. Additionally, the proposed model's versatility and effectiveness are demonstrated through various applications, including the evaluation of different cooling scenarios, closed-loop temperature control, and the thermal assessment of cell aspect ratio.

1 Introduction

The pursuit of carbon-neutral energy sources has emerged as a compelling and sustainable alternative to traditional non-renewable energy such as fossil fuels. This shift has been driven by the growing global commitment to combat climate change through the reduction of carbon emissions, particularly in the transportation sector [1]. Central to achieving this cleaner and more efficient mobility solution is the adoption of electric vehicles (EVs), whose performance strongly depends on the battery pack performance. Temperature affects sev-

eral aspects of the battery, including safety, electrochemical processes, charge acceptance, round-trip efficiency, power and energy capability, and lifetime [2]. This underscores the need for a sophisticated thermal management system [3] capable of controlling the battery temperature to desired values regardless of operating conditions.

Subject to various thermal boundary conditions such as liquid or air convection [4], battery cells can be thermally controlled at different interfaces, including the electrical connection tabs (terminals), selected regions of the external cell casing, or both [5]. Prior literature studies have investigated the impact of tab and lateral surface cooling methods on battery thermal performance. Hunt *et al.* [6] and Zhao *et al.* [7] concluded that lateral surface cooling is more effective at maintaining lower average temperatures in lithium (Li)-ion pouch cells, especially under high current rates, due to the larger cooling area. In a case study, these authors also demonstrated that tab cooling extended the lifetime of a battery pack by a factor of three, primarily due to the more uniform temperature distribution it provides. In a follow-up study, Dondelewski *et al.* [8] compared both methods and found that surface cooling resulted in lower average temperatures and slower degradation but reduced usable capacity. Conversely, tab cooling produced higher average temperatures and faster degradation, yet allowed for greater usable capacity due to more effective thermal distribution. Similar results were achieved for cylindrical Li-ion cells in [9], [10], where tab cooling was shown to reduce internal temperature inhomogeneities by approximately 25%, though it resulted in a higher average cell temperature. Furthermore, Ahmad *et al.* [11] reported that surface cooling causes a faster drop in temperature, driven by the larger lateral surface area available for heat transfer. However, tab cooling offered improved thermal uniformity, with a recorded temperature difference of 1.9°C, compared to 3.8°C under surface cooling.

It is clear that both lateral surface and tab cooling offer distinct advantages; surface cooling can enable faster heat removal and lower average cell temperatures, while tab cooling enhances temperature uniformity and thermal balance. However, to the best of our knowledge, there is currently no battery control framework in the state-of-the-art literature that systematically investigates the optimal integration of these two cooling methods for advanced battery thermal management.

Thermal equivalent circuit (TEC) models with lumped parameters have

been extensively utilized for control-oriented modelling of batteries due to their ease of implementation and computational efficiency [12], [13]. However, TECs can only predict average and surface temperatures and their applicability is limited to cells with small Biot numbers. Physics-based models [14] result in partial differential equations (PDEs) governing the underlying heat diffusion. They can predict the spatially distributed temperature field throughout the cell but are typically implemented via computationally expensive numerical schemes, such as finite element method (FEM) and finite difference method (FDM), rendering them impractical for real-time control purposes. Spectral methods [15], [16] are alternative numerical methods for finding solutions to PDEs. They belong to the class of weighted residual methods, and unlike FEM and FDM, they make use of global (rather than local) approximating functions in the discretisation of the spatial domain, rendering them computationally efficient. With this benefit in mind, the spectral method based on the Galerkin approach was adopted to develop low-order 2-dimensional (2D) thermal models for cylindrical cells in [17] and for pouch cells in [18]. However, these models do not allow independent and targeted cooling control of battery tabs and surfaces, preventing their use for optimally combining tab and surface cooling for battery thermal management.

To bridge the identified research gap, we propose new 2D battery thermal models that stem from the decomposition of the particular solution of the non-homogeneous PDE-based 2D heat equation. Unlike [17], [18], which used a lumped particular solution to capture all sides of the cell in one go, our model decomposes the particular solution into four components, each representing a distinct side of the battery cell. This method and its resulting model formulation enable us to obtain independent control signals for each side of the battery, including the tabs and surfaces.

For a given battery system and usage profile, the choice of the most suitable cooling scenario can significantly impact its key characteristics. Hence, we conduct a model-based analysis to study the performance trade-offs amongst different cooling scenarios. Additionally, we demonstrate the application of the model to study thermal gradients under various C-rates, through a simple feedback temperature control strategy based on multiple proportional-integral (PI) controllers. Finally, using cylindrical cells as an example, we conduct a model-based study on the influence of varying dimensions on battery thermal performance.

The novel models can accurately predict the spatially resolved full temperature distribution throughout battery cells of any geometry (cylindrical, pouch, and prismatic) with higher computational efficiency compared to TEC models. Hence, it can be readily used for real-time optimal control in battery management systems, particularly for the management of the cooling effort applied on battery tabs and surfaces. It can be an efficient simulation tool for understanding battery thermal dynamics to guide the design of next-generation cooling systems and battery cells.

The remainder of this article is organized as follows. Section 2 details the problem underlying the battery thermal modelling. Section 3 presents the proposed modelling framework for joint tab and surface cooling control. Section 4 presents the model validation and evaluation results. Section 5 showcases the efficacy of the model through various model-based applications and analyses. Finally, Section 6 concludes the work.

2 Overview of the PDE-based thermal model

Battery cells of cylindrical and pouch types are widely used for electrical energy storage, as schematized in Figs. 1(a)-(c). Considering the symmetric structure along the radial (r) direction of a cylindrical cell, and by assuming the same thermal conductivity in the y- and z-directions of a pouch cell, the temperature distribution of both cell types can be modelled in 2D. For completeness, this section briefly reviews the first-principle PDE-based models for battery thermal dynamics, including mathematical formulations as well as their underlying assumptions and considerations.

2.1 Cylindrical cell modelling

For the cylindrical cell, the governing equation for temperature $T(r, z, t)$ at time t and position (r, z) can be represented by the following 2D boundary value problem [19]:

$$\begin{aligned} \rho c_p \frac{\partial T(r,z,t)}{\partial t} - k_r \frac{\partial^2 T(r,z,t)}{\partial r^2} \\ -k_r \frac{\partial T(r,z,t)}{r \partial r} - k_z \frac{\partial^2 T(r,z,t)}{\partial z^2} = q(t), \end{aligned} \tag{A.1}$$

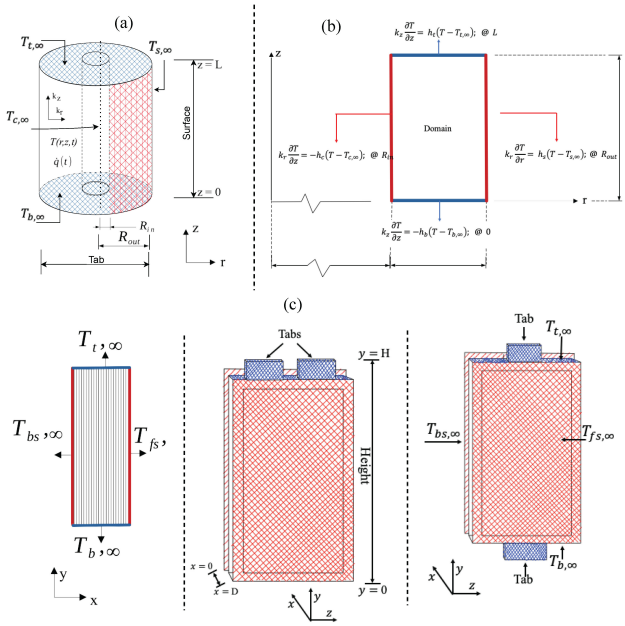

Figure 1: Schematics of different shapes of battery cells for 2D thermal modelling, with the blue and red-shaded areas representing tab and surface cooling channels, respectively. (a) A cylindrical cell cooled on the top and bottom tabs and part of the curved surface. (b) Illustration of the boundary conditions of (A.2) for cylindrical battery cell modelling. (c) A pouch cell whose tabs can either be on the top, bottom, or both sides.

subject to the non-homogeneous convection Robin boundary conditions given by

$$-k_r \frac{\partial T}{\partial r} = h_s \left(T - T_{s,\infty}\right) \qquad \text{at } r = R_{\text{out}}, \tag{A.2a}$$

$$k_r \frac{\partial T}{\partial r} = h_c \left(T - T_{c,\infty}\right) \qquad \text{at } r = R_{\text{in}}, \tag{A.2b}$$

$$-k_z \frac{\partial T}{\partial z} = h_t \left(T - T_{t,\infty}\right) \qquad \text{at } z = L, \tag{A.2c}$$

$$k_z \frac{\partial T}{\partial z} = h_b \left(T - T_{b,\infty}\right) \qquad \text{at } z = 0. \tag{A.2d}$$

where ρ, c_p, k, h, and T_∞ represent the volume-average density, volume-average specific heat capacity, thermal conductivity, convection coefficient, and fluid free-stream temperature, respectively. R_{in} and R_{out} are the cylindrical cell's inner and outer radius, and L is the length. The subscripts s, c, t, and b denote the cell's surface, core, top, and bottom, respectively. Equation (A.2) is illustrated in Fig. 1(b), where the convective heat transfer is presumed to occur on the external surfaces, and the heat transfer coefficient and fluid free-stream temperatures may differ for each surface. The cooling occurring in the core area is considered as zero, leading to the corresponding convection coefficient $h_c = 0$ in (A.2b).

The cylindrical cell's multi-layer composition is considered a homogeneous solid with anisotropic thermal conductivities along the r- and z-directions. Due to the spirally wound structure of the electrodes of the cylindrical cell, the thermal conductivity k_r would differ slightly from the inner to the outer radius but is assumed constant in this study for simplicity. The heat generation rate $q(t)$ in (A.1) is assumed to be uniformly distributed in space.

For cylindrical cells, the tabs placed on the top and bottom sides can have different sizes. Existing designs make the tab's area much smaller than the cell's top/bottom area. It has been demonstrated that these designs of cylindrical cells are not efficient for battery cooling merely from the tabs [9], [20], [21]. In this regard, this work considers the entire area of the top/bottom as the tab, forming the so-called all-tab or tab-less cell design, as illustrated in Figs. 1a and b. Correspondingly, the thermal convection coefficients h_t and h_b in (A.2) are affected by the entire top and bottom areas, respectively, in addition to the applied cooling media.

2.2 Pouch cell modelling

The thermal model for pouch cells illustrated in Fig. 1(c) can be derived in a similar fashion as for the cylindrical cell. The thermal conductivity in the y-direction is assumed to be identical to that of the z-direction because of the stacked structure of the pouch cell. In addition, symmetric cooling is assumed along the thin side edges of the pouch cell (the z-direction boundaries). By treating the pouch cell's multi-layer structure as a homogeneous solid with anisotropic thermal conductivities k_x in the x-direction and k_y in the y-direction, the governing heat equation in Cartesian coordinates is given

by the 2D boundary value problem [19]

$$\rho c_p \frac{\partial T(x,y,t)}{\partial t} - k_x \frac{\partial^2 T(x,y,t)}{\partial x^2} - k_y \frac{\partial^2 T(x,y,t)}{\partial y^2} = q(t), \tag{A.3}$$

subject to the non-homogeneous convection Robin boundary conditions expressed as

$$-k_x \frac{\partial T}{\partial x} = h_{fs}\left(T - T_{fs,\infty}\right) \quad \text{at } x = D, \tag{A.4a}$$

$$k_x \frac{\partial T}{\partial x} = h_{bs}\left(T - T_{bs,\infty}\right) \quad \text{at } x = 0, \tag{A.4b}$$

$$-k_y \frac{\partial T}{\partial y} = h_t\left(T - T_{t,\infty}\right) \quad \text{at } y = H, \tag{A.4c}$$

$$k_y \frac{\partial T}{\partial y} = h_b\left(T - T_{b,\infty}\right) \quad \text{at } y = 0. \tag{A.4d}$$

where D, H, the subscripts fs, and bs represent the pouch cell's width, height, front side, and back side, respectively.

3 Control-oriented battery thermal modelling

In this section, we demonstrate the control-oriented modelling process for the thermal behaviour of cylindrical cells. The model for pouch cells, which can be derived in a similar fashion, is introduced briefly afterwards.

3.1 Model non-dimensionalization

To facilitate model approximation, reduction, and the subsequent analyses, we first remove all physical dimensions involved in the model's governing equation (A.1). Specifically, new coordinate variables are introduced along the radius and height directions

$$\tilde{r} = \alpha r - \frac{R_{\text{out}} + R_{\text{in}}}{R_{\text{out}} - R_{\text{in}}}, \tag{A.5a}$$

$$\tilde{z} = \beta z - 1, \tag{A.5b}$$

where $\alpha = 2/(R_{\text{out}} - R_{\text{in}})$ and $\beta = 2/L$. As a result, both $\tilde{r}$ and $\tilde{z}$ fall in the range of $[-1, 1]$. Replacing all r and z by $\tilde{r}$ and $\tilde{z}$, respectively, the initial 2D

thermal model for cylindrical cells is scaled and becomes

$$\begin{aligned}\rho c_p \frac{\partial T(\tilde{r},\tilde{z},t)}{\partial t} - \alpha^2 k_r \frac{\partial^2 T(\tilde{r},\tilde{z},t)}{\partial \tilde{r}^2} \\ -\gamma \frac{\partial T(\tilde{r},\tilde{z},t)}{\partial \tilde{r}} - \beta^2 k_z \frac{\partial^2 T(\tilde{r},\tilde{z},t)}{\partial \tilde{z}^2} = q(t),\end{aligned} \tag{A.6}$$

subject to the boundary conditions

$$h_s T + \alpha k_r \frac{\partial T}{\partial \tilde{r}} = u_s \qquad \text{at } \tilde{r} = 1, \tag{A.7a}$$

$$-h_c T + \alpha k_r \frac{\partial T}{\partial \tilde{r}} = u_c \qquad \text{at } \tilde{r} = -1, \tag{A.7b}$$

$$h_t T + \beta k_z \frac{\partial T}{\partial \tilde{z}} = u_t \qquad \text{at } \tilde{z} = 1, \tag{A.7c}$$

$$-h_b T + \beta k_z \frac{\partial T}{\partial \tilde{z}} = u_b \qquad \text{at } \tilde{z} = -1. \tag{A.7d}$$

with $\gamma = \alpha^2 k_r/(1 + \tilde{r} + \alpha R_{\text{in}})$. The cooling power applied per unit area is denoted by u, and their values at the different sides of the cell are given, respectively, by

$$u_s(t) = h_s(t) T_{s,\infty}(t), \qquad u_c(t) = -h_c(t) T_{c,\infty}(t) = 0, \tag{A.8a}$$

$$u_t(t) = h_t(t) T_{t,\infty}(t), \qquad u_b(t) = -h_b(t) T_{b,\infty}(t). \tag{A.8b}$$

The above cooling power can be considered as the dynamic thermal system's input. The coolant's temperature and flow rate can influence each of these inputs. Furthermore, (A.8) can be integrated with a flow model to better capture the convective cooling effects and dynamic interactions between the coolant and the battery system under various cooling scenarios.

3.2 The overall solution

According to the boundary lifting principle [16], [22], [23], the overall solution to the boundary value problem formulated in (A.6)–(A.8) is composed of the homogeneous solution T_h and the particular solution T_p. Namely, there exists

$$T(\tilde{r},\tilde{z},t) = T_h(\tilde{r},\tilde{z},t) + T_p(\tilde{r},\tilde{z},t). \tag{A.9}$$

Here, T_h satisfies the modified problem for the cylindrical cell with the governing equation

$$\rho c_p \frac{\partial T_h}{\partial t} - \alpha^2 k_r \frac{\partial^2 T_h}{\partial \tilde{r}^2} - \gamma \frac{\partial T_h}{\partial \tilde{r}} - \beta^2 k_z \frac{\partial^2 T_h}{\partial \tilde{z}^2} = q^*, \tag{A.10}$$

and the homogeneous boundary conditions

$$h_s T_h + \alpha k_r \frac{\partial T_h}{\partial \tilde{r}} = 0 \qquad \text{at } \tilde{r} = 1, \tag{A.11a}$$

$$-h_c T_h + \alpha k_r \frac{\partial T_h}{\partial \tilde{r}} = 0 \qquad \text{at } \tilde{r} = -1, \tag{A.11b}$$

$$h_t T_h + \beta k_z \frac{\partial T_h}{\partial \tilde{z}} = 0 \qquad \text{at } \tilde{z} = 1, \tag{A.11c}$$

$$-h_b T_h + \beta k_z \frac{\partial T_h}{\partial \tilde{z}} = 0 \qquad \text{at } \tilde{z} = -1. \tag{A.11d}$$

where q^* in (A.10) is defined as

$$q^* \triangleq q - \left(-\alpha^2 k_r \frac{\partial^2 T_p}{\partial \tilde{r}^2} - \gamma \frac{\partial T_p}{\partial \tilde{r}} - \beta^2 k_z \frac{\partial^2 T_p}{\partial \tilde{z}^2} \right). \tag{A.12}$$

The particular solution for the locally distributed temperature T_p in (A.9) satisfies the original boundary conditions defined in (A.7), leading to

$$h_s T_p + \alpha k_r \frac{\partial T_p}{\partial \tilde{r}} = u_s \qquad \text{at } \tilde{r} = 1, \tag{A.13a}$$

$$-h_c T_p + \alpha k_r \frac{\partial T_p}{\partial \tilde{r}} = u_c \qquad \text{at } \tilde{r} = -1, \tag{A.13b}$$

$$h_t T_p + \beta k_z \frac{\partial T_p}{\partial \tilde{z}} = u_t \qquad \text{at } \tilde{z} = 1, \tag{A.13c}$$

$$-h_b T_p + \beta k_z \frac{\partial T_p}{\partial \tilde{z}} = u_b \qquad \text{at } \tilde{z} = -1. \tag{A.13d}$$

After the particular solution T_p is solved from (A.13), it can be substituted into (A.12) to facilitate the determination of the homogeneous solution T_h governed by (A.10)–(A.12).

3.3 Model reduction for the homogeneous solution

The homogeneous problem given by (A.10)–(A.12) can be solved using the Chebyshev spectral-Galerkin (CSG) method. Although Jacobi polynomials, such as Legendre or Chebyshev polynomials, could be effective candidates, Chebyshev functions are employed in this work because they can be readily defined to satisfy a wide range of boundary conditions [15], including the Robin boundary conditions encountered in this homogeneous problem.

Within the CSG method, the homogeneous solution T_h to the 2D thermal model of cylindrical battery cells can be approximated by a finite sum of basis functions in the form of

$$\hat{T}_h(\tilde{r},\tilde{z},t)=\sum_{m=0}^{M}\sum_{n=0}^{N}c_{mn}(t)\phi_m^{\tilde{r}}(\tilde{r})\phi_n^{\tilde{z}}(\tilde{z}), \tag{A.14}$$

where $c_{mn}(t)$ are unknown expansion (solution) coefficients. $\phi_m^{\tilde{r}}(\tilde{r})$ is the m-th Chebyshev basis function along the dimension $\tilde{r}$, satisfying the boundary conditions (A.11a)–(A.11b). M is the number of Chebyshev basis functions, and $m\in\{0,\cdots,M\}$. Along the dimension $\tilde{z}$, $\phi_n^{\tilde{z}}(\tilde{z})$ is the n-th Chebyshev basis function satisfying the boundary conditions (A.11c)–(A.11d), where $n\in\{0,\cdots,N\}$. M and N are tuning parameters depending on the required accuracy.

Let $P_m(\tilde{r})=\cos\left(m\theta^{\tilde{r}}\right)$ and $P_n(\tilde{z})=\cos\left(n\theta^{\tilde{z}}\right)$, with $\theta^{\tilde{r}}=\arccos\left(\tilde{r}\right)$ and $\theta^{\tilde{z}}=\arccos\left(\tilde{z}\right)$. Here, $P_m(\tilde{r})$ and $P_n(\tilde{z})$ represent the Chebyshev polynomials of the first kind, of degrees m and n, respectively. In spectral methods, to enable efficient solution computations and satisfaction of boundary conditions, neighbouring orthogonal polynomials should be used to form the basis functions. Therefore, we seek the basis functions as a compact combination of Chebyshev polynomials in the form [15]

$$\phi_m^{\tilde{r}}(\tilde{r})=P_m(\tilde{r})+a_m^{\tilde{r}}(\tilde{r})P_{m+1}(\tilde{r})+b_m^{\tilde{r}}(\tilde{r})P_{m+2}(\tilde{r}), \tag{A.15a}$$

$$\phi_n^{\tilde{z}}(\tilde{z})=P_n(\tilde{z})+a_n^{\tilde{z}}(\tilde{z})P_{n+1}(\tilde{z})+b_n^{\tilde{z}}(\tilde{z})P_{n+2}(\tilde{z}), \tag{A.15b}$$

where the coefficients $a_m^{\tilde{r}}$, $b_m^{\tilde{r}}$, $a_n^{\tilde{z}}$ and $b_n^{\tilde{z}}$ are defined according to Lemma 4.3 of [15].

Substituting (A.14) for T_h into (A.10) yields the residual denoted by R

$$R=\rho c_p\frac{\partial\hat{T}_h}{\partial t}-\alpha^2k_r\frac{\partial^2\hat{T}_h}{\partial\tilde{r}^2}-\gamma\frac{\partial\hat{T}_h}{\partial\tilde{r}}-\beta^2k_z\frac{\partial^2\hat{T}_h}{\partial\tilde{z}^2}-q^*\neq 0. \tag{A.16}$$

The principle of the CSG method is to force an integral of the resulting residual R to zero as

$$\begin{aligned}\langle rR,\eta\rangle=\Bigg\langle r\Big[&\rho c_p\frac{\partial\hat{T}_h}{\partial t}-\alpha^2k_r\frac{\partial^2\hat{T}_h}{\partial\tilde{r}^2}\\&-\gamma\frac{\partial\hat{T}_h}{\partial\tilde{r}}-\beta^2k_z\frac{\partial^2\hat{T}_h}{\partial\tilde{z}^2}-q^*\Big],\eta\Bigg\rangle=0,\end{aligned} \tag{A.17}$$

where we use the notation $\langle f, \eta \rangle$ to represent the inner product of f and a test function η in the domain of $\tilde{r}$ and $\tilde{z}$, namely $\langle f, \eta \rangle = \int_{-1}^{1} \int_{-1}^{1} f(\tilde{r}, \tilde{z}) \eta(\tilde{r}, \tilde{z}) d\tilde{r} d\tilde{z}$. Note that r is included in (A.17) to account for cylindrical coordinates. For the CSG method, the test function η must belong to the same set of Chebyshev basis functions, thereby giving $\eta = \phi_m^{\tilde{r}} \phi_n^{\tilde{z}}$.

Substituting $\hat{T}_h$ in (A.17) with its actual form in (A.14), we have

$$\begin{aligned}
&\Bigg\langle r \Big[\rho c_p \frac{\partial \big(\sum_{m=0}^{M} \sum_{n=0}^{N} c_{mn}(t) \phi_m^{\tilde{r}} \phi_n^{\tilde{z}} \big)}{\partial t} \\
&- \alpha^2 k_r \frac{\partial^2 \big(\sum_{m=0}^{M} \sum_{n=0}^{N} c_{mn}(t) \phi_m^{\tilde{r}} \phi_n^{\tilde{z}} \big)}{\partial \tilde{r}^2} \\
&- \gamma \frac{\partial \big(\sum_{m=0}^{M} \sum_{n=0}^{N} c_{mn}(t) \phi_m^{\tilde{r}} \phi_n^{\tilde{z}} \big)}{\partial \tilde{r}} \\
&- \beta^2 k_z \frac{\partial^2 \big(\sum_{m=0}^{M} \sum_{n=0}^{N} c_{mn}(t) \phi_m^{\tilde{r}} \phi_n^{\tilde{z}} \big)}{\partial \tilde{z}^2} - q^* \Big], \eta \Bigg\rangle = 0.
\end{aligned} \tag{A.18}$$

Finally, rewriting (A.18) in matrix-vector notations, this gives

$$\begin{aligned}
&\Bigg\langle r \Bigg[\rho c_p \Big[\phi_1^{\tilde{r}} \phi_1^{\tilde{z}}, \cdots, \phi_M^{\tilde{r}} \phi_N^{\tilde{z}} \Big] \begin{bmatrix} \dot{c}_{11}(t) \\ \vdots \\ \dot{c}_{MN}(t) \end{bmatrix} \\
&\quad - \alpha^2 k_r \Bigg[\frac{\partial^2 \phi_1^{\tilde{r}} \phi_1^{\tilde{z}}}{\partial \tilde{r}^2}, \cdots, \frac{\partial^2 \phi_M^{\tilde{r}} \phi_N^{\tilde{z}}}{\partial \tilde{r}^2} \Bigg] \begin{bmatrix} c_{11}(t) \\ \vdots \\ c_{MN}(t) \end{bmatrix} \\
&\quad - \gamma \Bigg[\frac{\partial \phi_1^{\tilde{r}} \phi_1^{\tilde{z}}}{\partial \tilde{r}}, \cdots, \frac{\partial \phi_M^{\tilde{r}} \phi_N^{\tilde{z}}}{\partial \tilde{r}} \Bigg] \begin{bmatrix} c_{11}(t) \\ \vdots \\ c_{MN}(t) \end{bmatrix} \\
&\quad - \beta^2 k_z \Bigg[\frac{\partial^2 \phi_1^{\tilde{r}} \phi_1^{\tilde{z}}}{\partial \tilde{z}^2}, \cdots, \frac{\partial^2 \phi_M^{\tilde{r}} \phi_N^{\tilde{z}}}{\partial \tilde{z}^2} \Bigg] \begin{bmatrix} c_{11}(t) \\ \vdots \\ c_{MN}(t) \end{bmatrix} - q^* \Bigg], \eta \Bigg\rangle = 0.
\end{aligned} \tag{A.19}$$

3.4 Model reduction for the particular solution

Remark 1. *Unlike the homogeneous solution, which often exhibits smoother spatial variations governed by the battery system's intrinsic properties, the particular solution T_p may require a more flexible representation to accommodate abrupt changes or spatial gradients induced by external factors such as boundary conditions. Additionally, T_p must satisfy not only the PDE* (A.6) *but also the non-homogeneous boundary conditions* (A.7) *and any forcing terms, such as the heat generation q, presented in the problem. These additional constraints necessitate a richer and more flexible approach than a straightforward finite-sum expansion as employed in the homogeneous case.*

Based on *Remark 1*, we develop a highly adaptable approach for solving the particular solution T_p. This approach will be able to effectively handle abrupt changes and spatial gradients while satisfying all constraints imposed by (A.6)–(A.7).

Approximation of T_p from a polynomial vector space

As the first step of model reduction for the particular solution T_p, we project and approximate it from a polynomial vector space denoted by $\mathcal{V}_{rz}$. Following a similar approach as in [24], [25], $\mathcal{V}_{rz}$ is defined as

$$\begin{aligned}\mathcal{V}_{rz} = &\left[\{\phi_m^{\tilde{r}}(\tilde{r}), 1 \le m \le M\} \otimes \{\tilde{z}, \tilde{z}^2\}\right] \\ &\oplus \left[\{\phi_n^{\tilde{z}}(\tilde{z}), 1 \le n \le N\} \otimes \{\tilde{r}, \tilde{r}^2\}\right],\end{aligned} \tag{A.20}$$

where $\otimes$ and $\oplus$ are the outer product and direct sum, respectively. $\tilde{z}$, $\tilde{z}^2$, $\tilde{r}$, and $\tilde{r}^2$ are the basis vectors of $\mathcal{V}_{rz}$. (A.20) explicitly couples the $\tilde{r}$- and $\tilde{z}$-dimensions by combining basis functions from one dimension with polynomials in the other. This coupling is essential because the heat generation and boundary conditions can create interactions between $\tilde{r}$ and $\tilde{z}$ that a simpler product space might not capture effectively.

To elucidate the idea of finite-sum approximation in the vector space $\mathcal{V}_{rz}$, we start with a simple case. Consider a vector space $\mathcal{V}_y$ consisting of the basis vectors y and y^2. This space can be written as

$$\mathcal{V}_y = \text{span}\{y, y^2\}. \tag{A.21}$$

Any function, $f(y) \in \mathcal{V}_y$ can be represented as a linear combination of these

basis vectors. Specifically, we can write any such $f(y)$ using a finite-sum approximation

$$f(y) \approx \sum_{i=0}^{2} p_i y^i, \tag{A.22}$$

where p_i are the unknown coefficients, which can be determined using various methods, such as least squares or spectral methods.

Following the idea of finite-sum approximation introduced in (A.21)-(A.22), approximating T_p in the vector space $\mathcal{V}_{rz}$, i.e., in two space dimensions, leads to the expansion

$$\begin{aligned}\hat{T}_p(\tilde{r},\tilde{z},t) = &\sum_{n=0}^{N} \Big(D_{1,n}(\tilde{z},t)\tilde{r} + D_{2,n}(\tilde{z},t)\tilde{r}^2\Big)\phi_n^{\tilde{z}}(\tilde{z}) \\ &+ \sum_{m=0}^{M} \Big(D_{3,m}(\tilde{r},t)\tilde{z} + D_{4,m}(\tilde{r},t)\tilde{z}^2\Big)\phi_m^{\tilde{r}}(\tilde{r}),\end{aligned} \tag{A.23}$$

where $\hat{T}_p$ is the approximate solution of T_p, and $D_{1,n}, D_{2,n}, D_{3,m}$, and $D_{4,m}$ are unknown expansion coefficients. The first and second terms on the right-hand side of (A.23) represent the lumped particular solutions in the vertical and horizontal directions of the battery cell, respectively.

Remark 2. *The deviation between $\hat{T}_p$ and T_p can arise from two primary sources: the chosen vector space along $\tilde{z}$ and $\tilde{r}$, and the finite-sum approximation, which depends on the tuning parameters N and M. Generally, using higher-order polynomials and including more terms in the finite sum results in smaller approximation errors but increases the number of expansion coefficients that need to be determined. The trade-offs involved, as well as the applicability of this approximation method, will be quantitatively investigated in Section 4.*

Determination of the expansion coefficients

Taking advantage of the orthogonality and approximation properties of the Chebyshev polynomials in (A.15), we employ the CSG method to solve for the boundary conditions formulated in (A.13). The process is similar to that for the homogenous problem, with the details given in the following. Substituting

(A.23) into the boundary condition at the surface side (A.13a), yields the residual R_s given by

$$R_s = h_s \hat{T}_p - \alpha k_r \frac{\partial \hat{T}_p}{\partial \tilde{r}} - u_s \neq 0, \text{ at } \tilde{r} = 1. \tag{A.24}$$

As the CSG method demands, we force the integral of the resulting residual R_s to zero as

$$\langle R_s, \phi_i^{\tilde{z}} \rangle_{\tilde{z}} = \left\langle \left[h_s \hat{T}_p - \alpha k_r \frac{\partial \hat{T}_p}{\partial \tilde{r}} - u_s \right], \phi_i^{\tilde{z}} \right\rangle_{\tilde{z}} = 0, \text{ at } \tilde{r} = 1, \tag{A.25}$$

where the notation $\langle f, \phi_i^{\tilde{z}} \rangle_{\tilde{z}} = \int_{-1}^{1} f(\tilde{z}) \phi_i^{\tilde{z}}(\tilde{z}) d\tilde{z}$ represents the one-dimensional inner product in the $\tilde{z}$-direction, and $\phi_i^{\tilde{z}}$ with $i \in \{0, \cdots, N\}$, are the test functions. Substituting $\hat{T}_p$ in (A.25) with its actual form in (A.23), and with $\tilde{r} = 1$, simplifies to

$$\begin{aligned} &\sum_{n=0}^{N} \Big[D_{1,n}(h_s - \alpha k_r) + D_{2,n}(h_s - 2\alpha k_r) \Big] \Big\langle \phi_n^{\tilde{z}}, \phi_i^{\tilde{z}} \Big\rangle_{\tilde{z}} \\ &\qquad = \Big\langle u_s, \phi_i^{\tilde{z}} \Big\rangle_{\tilde{z}}. \end{aligned} \tag{A.26}$$

Then, we adopt a similar process for the other vertical side (i.e., the core) and the horizontal sides (i.e., the top and bottom sides) of the battery cell with the boundary conditions formulated in (A.13b)–(A.13d). Consequently, we get

$$\begin{aligned} &\sum_{n=0}^{N} \Big[D_{1,n}(h_c - \alpha k_r) + D_{2,n}(-h_c + 2\alpha k_r) \Big] \Big\langle \phi_n^{\tilde{z}}, \phi_i^{\tilde{z}} \Big\rangle_{\tilde{z}} \\ &\qquad = \Big\langle u_c, \phi_i^{\tilde{z}} \Big\rangle_{\tilde{z}}, \end{aligned} \tag{A.27a}$$

$$\begin{aligned} &\sum_{m=0}^{M} \Big[D_{3,m}(h_t - \beta k_z) + D_{4,m}(h_t - 2\beta k_z) \Big] \Big\langle \phi_m^{\tilde{r}}, \phi_j^{\tilde{r}} \Big\rangle_{\tilde{r}} \\ &\qquad = \Big\langle u_t, \phi_j^{\tilde{r}} \Big\rangle_{\tilde{r}}, \end{aligned} \tag{A.27b}$$

$$\begin{aligned} &\sum_{m=0}^{M} \Big[D_{3,m}(h_b - \beta k_z) + D_{4,m}(-h_b + 2\beta k_z) \Big] \Big\langle \phi_m^{\tilde{r}}, \phi_j^{\tilde{r}} \Big\rangle_{\tilde{r}} \\ &\qquad = \Big\langle u_b, \phi_j^{\tilde{r}} \Big\rangle_{\tilde{r}}, \end{aligned} \tag{A.27c}$$

where the notation $\langle f, \phi_j^{\tilde{r}} \rangle_{\tilde{r}} = \int_{-1}^{1} r f(\tilde{r}) \phi_j^{\tilde{r}}(\tilde{r}) d\tilde{r}$, represents the one-dimensional inner product in the $\tilde{r}$-direction, and r is included in the inner product to account for cylindrical coordinates. $\phi_j^{\tilde{r}}$ with $j \in \{0, \cdots, M\}$ are the test functions.

For all values of i and j in (A.26)–(A.27), there are $2N + 2M$ unknown coefficients in $2N+2M$ equations. By solving these equations, we can uniquely determine $D_{1,n}$, $D_{2,n}$, $D_{3,m}$, and $D_{4,m}$ for all n and m. To simplify the notations, along the cell's surface, core, top, and bottom sides, we define

$$
\begin{aligned}
s_1 &= (h_s - \alpha k_r), \quad s_2 = (h_s - 2\alpha k_r), \\
t_1 &= (h_t - \beta k_z), \quad t_2 = (h_t - 2\beta k_z), \\
c_1 &= (h_c - \alpha k_r), \quad c_2 = (-h_c + 2\alpha k_r), \\
b_1 &= (h_b - \beta k_z), \quad b_2 = (-h_b + 2\beta k_z).
\end{aligned}
$$

Then, (A.26)-(A.27) can be re-written as

$$\Phi_v \left(D_1 s_1 + D_2 s_2 \right) = u_s S_v, \tag{A.28a}$$
$$\Phi_v \left(D_1 c_1 + D_2 c_2 \right) = u_c S_v, \tag{A.28b}$$
$$\Phi_h \left(D_3 t_1 + D_4 t_2 \right) = u_t S_h, \tag{A.28c}$$
$$\Phi_h \left(D_3 b_1 + D_4 b_2 \right) = u_b S_h, \tag{A.28d}$$

where the matrices for the vertical sides $\Phi_v(\tilde{z}) \in \mathcal{R}^{N \times N}$ and the horizontal sides $\Phi_h(\tilde{r}) \in \mathcal{R}^{N \times N}$ are given by $\Phi_v = \langle \phi_n^{\tilde{z}}, \phi_i^{\tilde{z}} \rangle_{\tilde{z}}$ and $\Phi_h = \langle \phi_m^{\tilde{r}}, \phi_j^{\tilde{r}} \rangle_{\tilde{r}}$, respectively. The source terms for the vertical sides $S_v(\tilde{z}) \in \mathcal{R}^N$ and the horizontal sides $S_h(\tilde{r}) \in \mathcal{R}^M$ are given by $S_v = \langle 1, \phi_i^{\tilde{z}} \rangle_{\tilde{z}}$ and $S_h = \langle 1, \phi_j^{\tilde{r}} \rangle_{\tilde{r}}$, respectively. D_1, D_2, D_3, and D_4 are vectors of the unknown coefficients, explicitly defined as

$$
\begin{aligned}
D_1 &= \begin{bmatrix} D_{1,1} & D_{1,2} & \cdots & D_{1,N} \end{bmatrix}^T, \\
D_2 &= \begin{bmatrix} D_{2,1} & D_{2,2} & \cdots & D_{2,N} \end{bmatrix}^T, \\
D_3 &= \begin{bmatrix} D_{3,1} & D_{3,2} & \cdots & D_{3,M} \end{bmatrix}^T, \\
D_4 &= \begin{bmatrix} D_{4,1} & D_{4,2} & \cdots & D_{4,M} \end{bmatrix}^T.
\end{aligned}
$$

By solving (A.28), the coefficient vectors can be determined

$$D_1 = \frac{1}{s_1c_2 - s_2c_1}(c_2u_s - s_2u_c)\Phi_v^{-1}S_v, \tag{A.29a}$$

$$D_2 = \frac{1}{s_1c_2 - s_2c_1}(-c_2u_s + s_1u_c)\Phi_v^{-1}S_v, \tag{A.29b}$$

$$D_3 = \frac{1}{t_1b_2 - t_2b_1}(b_2u_t - t_2u_b)\Phi_h^{-1}S_h, \tag{A.29c}$$

$$D_4 = \frac{1}{t_1b_2 - t_2b_1}(-b_2u_t + t_1u_b)\Phi_h^{-1}S_h. \tag{A.29d}$$

As a result, the particular solution $\hat{T}_p$ in (A.23) has been derived.

Model decomposition for independent and targeted cooling control

Remark 3. *Previous efforts in the battery community to derive* (A.9) *e.g., in [17], [18], utilised* (A.23) *as their particular solution. However,* (A.23) *in its present form does not allow independent and targeted cooling control of battery tabs and surfaces, preventing its use for optimally combining tab and surface cooling for enhanced battery thermal performance.*

Different from existing works in the literature as noted in *Remark 3*, we propose a new method to derive the particular solution T_p in battery thermal modelling. This method and its resulting model formulation offer the flexibility of independent controls for each side of the cell, including the tabs and surfaces. It is achieved by pertinent model decomposition, in which the key idea is to decompose (A.23) into four components, each representing a side of the cylindrical battery cell.

According to (A.29), the coefficient vectors, D_1 and D_2, are functions of u_s and u_c, while D_3 and D_4 are functions of u_t and u_b. With this in mind, we can re-arrange (A.29) so that each coefficient vector is expressed into two terms as

$$\begin{aligned} D_1 &= d_1^s u_s + d_1^c u_c,\ D_2 = d_2^s u_s + d_2^c u_c, \\ D_3 &= d_1^t u_t + d_1^b u_b,\ D_4 = d_2^t u_t + d_2^b u_b, \end{aligned} \tag{A.30}$$

where

$$\begin{aligned} d_1^s &= \frac{c_2}{s_1c_2 - s_2c_1}\Phi_v^{-1}S_v, \quad d_2^s = \frac{-c_2}{s_1c_2 - s_2c_1}\Phi_v^{-1}S_v, \\ d_1^c &= \frac{-s_2}{s_1c_2 - s_2c_1}\Phi_v^{-1}S_v, \quad d_2^c = \frac{s_1}{s_1c_2 - s_2c_1}\Phi_v^{-1}S_v, \\ d_1^t &= \frac{b_2}{t_1b_2 - t_2b_1}\Phi_h^{-1}S_h, \quad d_2^t = \frac{-b_2}{t_1b_2 - t_2b_1}\Phi_h^{-1}S_h, \\ d_1^b &= \frac{-t_2}{t_1b_2 - t_2b_1}\Phi_h^{-1}S_h, \quad d_2^b = \frac{t_1}{t_1b_2 - t_2b_1}\Phi_h^{-1}S_h. \end{aligned} \tag{A.31}$$

Substituting (A.30)–(A.31) into (A.23), we have the new expansion for the particular solution as

$$\begin{aligned} \hat{T}_p(\tilde{r},\tilde{z},t) = \sum_{n=0}^{N}\Bigg[\Big(d_1^s(\tilde{z})u_s(t) + d_1^c(\tilde{z})u_c(t)\Big)\tilde{r} + \Big(d_2^s(\tilde{z})u_s(t) \\ + d_2^c(\tilde{z})u_c(t)\Big)\tilde{r}^2\Bigg]\phi_n^{\tilde{z}}(\tilde{z}) + \sum_{m=0}^{M}\Bigg[\Big(d_1^t(\tilde{r})u_t(t) + d_1^b(\tilde{r})u_b(t)\Big)\tilde{z} \\ + \Big(d_2^t(\tilde{r})u_t(t) + d_2^b(\tilde{r})u_b(t)\Big)\tilde{z}^2\Bigg]\phi_m^{\tilde{r}}(\tilde{r}). \end{aligned} \tag{A.32}$$

The above result can be re-organized into four components as

$$\begin{aligned} \hat{T}_p(\tilde{r},\tilde{z},t) = \hat{T}_p^s(\tilde{r},\tilde{z})u_s(t) + \hat{T}_p^c(\tilde{r},\tilde{z})u_c(t) \\ + \hat{T}_p^t(\tilde{r},\tilde{z})u_t(t) + \hat{T}_p^b(\tilde{r},\tilde{z})u_b(t), \end{aligned} \tag{A.33}$$

where $\hat{T}_p^s$, $\hat{T}_p^c$, $\hat{T}_p^t$, and $\hat{T}_p^b$ represent the particular solution expansions projected onto the surface, core, top, and bottom, respectively. Specifically, these components are given by

$$\hat{T}_p^s = \sum_{n=0}^{N}\Big(d_1^s\tilde{r} + d_2^s\tilde{r}^2\Big)\phi_n^{\tilde{z}}, \;\; \hat{T}_p^c = \sum_{n=0}^{N}\Big(d_1^c\tilde{r} + d_2^c\tilde{r}^2\Big)\phi_n^{\tilde{z}}, \tag{A.34a}$$

$$\hat{T}_p^t = \sum_{m=0}^{M}\Big(d_1^t\tilde{z} + d_2^t\tilde{z}^2\Big)\phi_m^{\tilde{r}}, \;\; \hat{T}_p^b = \sum_{m=0}^{M}\Big(d_1^b\tilde{z} + d_2^b\tilde{z}^2\Big)\phi_m^{\tilde{r}}. \tag{A.34b}$$

We recall from (A.8) that $u_c = 0$ for the cylindrical cell, thus $\hat{T}_p^c$ does not contribute to the particular solution. Due to this, we neglect $\hat{T}_p^c$ in the sequel.

Remark 4. *The result obtained in* (A.33)–(A.34) *implies that for the particular problem in* (A.13)*, which leads to the particular solution T_p featuring linearity in two spatial dimensions, the solution trajectory in response to the cooling power applied at each individual boundary satisfies the superposition principle. Additionally, the particular solution $T_p(\tilde{r}, \tilde{z}, t)$ at each position $(\tilde{r}, \tilde{z})$ can be expressed as a linear combination of the cooling powers applied at different boundaries.*

3.5 State-space model representation

Substituting the particular solution in (A.33)–(A.34) into (A.12) and using the linearity of the partial differential operator, the term q^*, which is required for deriving the homogeneous solution T_h, can be obtained accordingly

$$\begin{aligned} q^* = q &+ \left[\alpha^2 k_r \frac{\partial^2 \hat{T}_p^s}{\partial \tilde{r}^2} + \gamma \frac{\partial \hat{T}_p^s}{\partial \tilde{r}} + \beta^2 k_z \frac{\partial^2 \hat{T}_p^s}{\partial \tilde{z}^2}\right] u_s \\ &+ \left[\alpha^2 k_r \frac{\partial^2 \hat{T}_p^t}{\partial \tilde{r}^2} + \gamma \frac{\partial \hat{T}_p^t}{\partial \tilde{r}} + \beta^2 k_z \frac{\partial^2 \hat{T}_p^t}{\partial \tilde{z}^2}\right] u_t \\ &+ \left[\alpha^2 k_r \frac{\partial^2 \hat{T}_p^b}{\partial \tilde{r}^2} + \gamma \frac{\partial \hat{T}_p^b}{\partial \tilde{r}} + \beta^2 k_z \frac{\partial^2 \hat{T}_p^b}{\partial \tilde{z}^2}\right] u_b. \end{aligned} \tag{A.35}$$

Based on the results in Sections 3.3–3.4, a control-oriented 2D thermal model of cylindrical battery cells is achieved. We define X as the state variable, u the input, $\mathcal{Y}$ the output, and w the disturbance. According to (A.19), X is given by

$$X = \left[c_{11}, \cdots, c_{1N}, c_{21}, \cdots, c_{2N}, \cdots, c_{M1}, \cdots, c_{MN}\right]^T. \tag{A.36}$$

The input u is composed of the cooling power applied at the different places, leading to $u = \left[u_s,\ u_t,\ u_b\right]^T$. The disturbance w represents the heat generation rate, namely $w = q$. The outputs of interest, $\mathcal{Y}$, include the temperatures at the mid-points of the surface, core, top, and bottom sides, namely,

$$\mathcal{Y}(t) = \begin{bmatrix}\hat{T}(1,0,t) & \hat{T}(-1,0,t) & \hat{T}(0,1,t) & \hat{T}(0,-1,t)\end{bmatrix}^T, \tag{A.37}$$

where $\hat{T}$ is the approximate solution of T in (A.9). With these definitions, this dynamic thermal model can be presented in the state-space form

$$\mathcal{G}\dot{X}(t) = \mathcal{A}X(t) + \mathcal{B}u(t) + \mathcal{F}w(t), \tag{A.38a}$$

$$\mathcal{Y}(t) = \mathcal{C}X(t) + \mathcal{D}u(t), \tag{A.38b}$$

where with the additional term

$$\mathcal{D} = \begin{bmatrix} \hat{T}_p^s(1,0) & \hat{T}_p^t(1,0) & \hat{T}_p^b(1,0) \\ \hat{T}_p^s(-1,0) & \hat{T}_p^t(-1,0) & \hat{T}_p^b(-1,0) \\ \hat{T}_p^s(0,1) & \hat{T}_p^t(0,1) & \hat{T}_p^b(0,1) \\ \hat{T}_p^s(0,-1) & \hat{T}_p^t(0,-1) & \hat{T}_p^b(0,-1) \end{bmatrix} \tag{A.39}$$

the matrices $\mathcal{G}$, $\mathcal{A}$, $\mathcal{B}$, $\mathcal{C}$, and the vector $\mathcal{F}$ are derived from (A.9), (A.12), (A.19), and (A.33)–(A.34) to yield

$$\mathcal{G} = \rho c_p \left\langle r \phi_m^{\tilde{r}} \phi_n^{\tilde{z}}, \eta \right\rangle, \tag{A.40a}$$

$$\mathcal{A} = \left\langle r \left[\alpha^2 k_r \frac{\partial^2 \phi_m^{\tilde{r}} \phi_n^{\tilde{z}}}{\partial \tilde{r}^2} + \gamma \frac{\partial \phi_m^{\tilde{r}} \phi_n^{\tilde{z}}}{\partial \tilde{r}} + \beta^2 k_z \frac{\partial^2 \phi_m^{\tilde{r}} \phi_n^{\tilde{z}}}{\partial \tilde{z}^2} \right], \eta \right\rangle, \tag{A.40b}$$

$$\mathcal{B}(:,1) = \left\langle r \left[\alpha^2 k_r \frac{\partial^2 \hat{T}_p^s}{\partial \tilde{r}^2} + \gamma \frac{\partial \hat{T}_p^s}{\partial \tilde{r}} + \beta^2 k_z \frac{\partial^2 \hat{T}_p^s}{\partial \tilde{z}^2} \right], \eta \right\rangle, \tag{A.40c}$$

$$\mathcal{B}(:,2) = \left\langle r \left[\alpha^2 k_r \frac{\partial^2 \hat{T}_p^t}{\partial \tilde{r}^2} + \gamma \frac{\partial \hat{T}_p^t}{\partial \tilde{r}} + \beta^2 k_z \frac{\partial^2 \hat{T}_p^t}{\partial \tilde{z}^2} \right], \eta \right\rangle, \tag{A.40d}$$

$$\mathcal{B}(:,3) = \left\langle r \left[\alpha^2 k_r \frac{\partial^2 \hat{T}_p^b}{\partial \tilde{r}^2} + \gamma \frac{\partial \hat{T}_p^b}{\partial \tilde{r}} + \beta^2 k_z \frac{\partial^2 \hat{T}_p^b}{\partial \tilde{z}^2} \right], \eta \right\rangle, \tag{A.40e}$$

$$\mathcal{C}(1,:) = \phi_m^{\tilde{r}}(1) \phi_n^{\tilde{z}}(0), \;\; \mathcal{C}(2,:) = \phi_m^{\tilde{r}}(-1) \phi_n^{\tilde{z}}(0), \tag{A.40f}$$

$$\mathcal{C}(3,:) = \phi_m^{\tilde{r}}(0) \phi_n^{\tilde{z}}(1), \;\; \mathcal{C}(4,:) = \phi_m^{\tilde{r}}(0) \phi_n^{\tilde{z}}(-1), \tag{A.40g}$$

$$\mathcal{F} = \langle r \cdot 1, \eta \rangle. \tag{A.40h}$$

Up to this point, the initial PDE-based model for cylindrical battery cells governed by (A.1)–(A.2) has been reformulated and approximated by the above ODE-based model. The order of this model, which is also called the number of states, is $O = MN$.

3.6 Application of the modelling techniques to pouch cells

The modelling techniques introduced in Section 3.1–3.5 can also be applied to pouch battery cells. For brevity, only the key mathematical results are presented.

As a result of non-dimensionalisation, the initial thermal model for pouch cells in (A.3)–(A.4) can be presented in the form

$$\rho c_p \frac{\partial T}{\partial t} - \lambda^2 k_x \frac{\partial^2 T}{\partial \tilde{x}^2} - \zeta^2 k_y \frac{\partial^2 T}{\partial \tilde{y}^2} = q, \tag{A.41}$$

with scaling factors $\lambda = 2/D$ and $\zeta = 2/H$. The new spatial coordinates are $\tilde{x} = 2x/D - 1$ and $\tilde{y} = 2y/H - 1$. The above dynamic equation is subject to the boundary conditions

$$h_{fs} T + \lambda k_x \frac{\partial T}{\partial \tilde{x}} = u_{fs} \qquad \text{at } \tilde{x} = 1, \tag{A.42a}$$

$$-h_{bs} T + \lambda k_x \frac{\partial T}{\partial \tilde{x}} = u_{bs} \qquad \text{at } \tilde{x} = -1, \tag{A.42b}$$

$$h_t T + \zeta k_y \frac{\partial T}{\partial \tilde{y}} = u_t \qquad \text{at } \tilde{y} = 1, \tag{A.42c}$$

$$-h_b T + \zeta k_y \frac{\partial T}{\partial \tilde{y}} = u_b \qquad \text{at } \tilde{y} = -1. \tag{A.42d}$$

where u_{fs}, u_{bs}, u_t, and u_b represent the cooling power applied per unit area at the front, back, top, and bottom sides of the pouch cell, respectively, which are mathematically defined as

$$u_{fs}(t) = h_{fs}(t) T_{fs,\infty}(t), \;\; u_{bs}(t) = -h_{bs}(t) T_{bs,\infty}(t), \tag{A.43a}$$

$$u_t(t) = h_t(t) T_{t,\infty}(t), \;\; u_b(t) = -h_b(t) T_{b,\infty}(t). \tag{A.43b}$$

The outputs of interest $\mathcal{Y}$, for pouch battery cells are located at the middle of each side, with the mathematical definition being the same as (A.37).

Following similar steps from (A.9)–(A.35), the battery model in (A.41)–(A.42) can also be approximated by an ODE-based model in the state-space form of (A.38) with the state variable defined in (A.36) and the input $u = [u_{fs},\ u_{bs},\ u_t,\ u_b]^T$. The resulting matrices and vectors for this state-space

model are given by

$$\mathcal{G} = \rho c_p \big\langle \phi_m^{\tilde{x}} \phi_n^{\tilde{y}}, \eta \big\rangle, \tag{A.44a}$$

$$\mathcal{A} = \left\langle \lambda^2 k_x \frac{\partial^2 \phi_m^{\tilde{x}} \phi_n^{\tilde{y}}}{\partial \tilde{x}^2} + \zeta^2 k_y \frac{\partial^2 \phi_m^{\tilde{x}} \phi_n^{\tilde{y}}}{\partial \tilde{y}^2}, \eta \right\rangle, \tag{A.44b}$$

$$\mathcal{B}(:,1) = \left\langle \lambda^2 k_x \frac{\partial^2 \hat{T}_p^{fs}}{\partial \tilde{x}^2} + \zeta^2 k_y \frac{\partial^2 \hat{T}_p^{fs}}{\partial \tilde{y}^2}, \eta \right\rangle, \tag{A.44c}$$

$$\mathcal{B}(:,2) = \left\langle \lambda^2 k_x \frac{\partial^2 \hat{T}_p^{bs}}{\partial \tilde{x}^2} + \zeta^2 k_y \frac{\partial^2 \hat{T}_p^{bs}}{\partial \tilde{y}^2}, \eta \right\rangle, \tag{A.44d}$$

$$\mathcal{B}(:,3) = \left\langle \lambda^2 k_x \frac{\partial^2 \hat{T}_p^{t}}{\partial \tilde{x}^2} + \zeta^2 k_y \frac{\partial^2 \hat{T}_p^{t}}{\partial \tilde{y}^2}, \eta \right\rangle, \tag{A.44e}$$

$$\mathcal{B}(:,4) = \left\langle \lambda^2 k_x \frac{\partial^2 \hat{T}_p^{b}}{\partial \tilde{x}^2} + \zeta^2 k_y \frac{\partial^2 \hat{T}_p^{b}}{\partial \tilde{y}^2}, \eta \right\rangle, \tag{A.44f}$$

$$\mathcal{C}(1,:) = \phi_m^{\tilde{x}}(1)\phi_n^{\tilde{y}}(0), \;\; \mathcal{C}(2,:) = \phi_m^{\tilde{x}}(-1)\phi_n^{\tilde{y}}(0), \tag{A.44g}$$

$$\mathcal{C}(3,:) = \phi_m^{\tilde{x}}(0)\phi_n^{\tilde{y}}(1), \;\; \mathcal{C}(4,:) = \phi_m^{\tilde{x}}(0)\phi_n^{\tilde{y}}(-1), \tag{A.44h}$$

$$\mathcal{F} = \langle 1, \eta \rangle, \tag{A.44i}$$

where $\hat{T}_p^{fs}$ and $\hat{T}_p^{bs}$ represent the particular solutions of the front- and backside expansions, respectively.

At higher C-rates, localized heat generation at the battery tabs can significantly influence the internal temperature distribution of the cell [26]. Future work will explore modelling approaches that treat the tabs and cell body as thermally coupled regions. This will enable a more spatially resolved thermal analysis and better capture the impact of tab heating on overall cell performance.

Due to the internal structure of prismatic cells, significant thermal gradients can develop in all space dimensions (x, y, and z), requiring a 3D thermal model for better accuracy. However, the underlying solution processes for the prismatic cell are equivalent to those of the cylindrical and pouch cells provided in Sections 3.3–3.4.

4 Model validation and evaluation

In this section, we evaluate the developed spatially distributed thermal models for cylindrical battery cells in (A.38)–(A.40) and for pouch cells in (A.38)

and (A.44). In addition to computational efficiency, the evaluation criteria include model accuracy for the maximum temperature $T_{\max}$, mean temperature T_{mean}, maximum thermal gradient along the radius direction $\delta T^r_{\max}$, maximum thermal gradient along the axial direction $\delta T^z_{\max}$, and mean thermal gradient δT^r_{mean}. Then, we incorporate these models in a feedback controller for battery cooling and investigate the effectiveness of various tab and surface cooling scenarios. Additionally, using cylindrical cells as an example, we study the influence of varying dimensions on battery thermal performance.

4.1 Benchmarks

To comprehensively and systematically assess the proposed modelling approach, we introduce two types of benchmarks. The first type is the original PDE models, i.e., (A.1)–(A.4), solved by the FEM. Results from these models will serve as ground truth in comparative studies. The second type of benchmark is TEC models that are widely studied in the literature [12], [13], [27], [28] and have been utilised in commercial battery management systems (BMSs). This kind of model features lumped parameters and two states in the form

$$C_c\dot{T}_c(t) = q(t) + \frac{T_s(t) - T_c(t)}{R_c}, \tag{A.45a}$$

$$C_s\dot{T}_s(t) = \frac{T_\infty - T_s(t)}{R_u} + \frac{T_s(t) - T_c(t)}{R_c}, \tag{A.45b}$$

where T_c and T_s represent the core and surface temperatures, respectively. R_c is the heat conduction resistance, which quantifies the heat exchange between the core and surface. R_u is the convection resistance, which measures the convection cooling along the battery surface and is dependent on the pack geometry, coolant type, and flow rate. C_c and C_s indicate the heat capacity of the core and surface (casing), respectively.

To utilize the TEC model in (A.45), we first need to identify the unknown parameters, i.e., C_c, C_s, R_c and R_u, for given battery cells and cooling scenarios. We follow the parameterisation scheme proposed in [13], which was designed using the least squares algorithm and the measurable current I (or q), T_s, and T_∞. According to [13], only three of the four parameters can be identified from these measurements, and one parameter must be assumed. In fact, the heat capacity of the aluminium casing C_s can be calculated from its specific heat capacity and dimensions.

4.2 Specification and simulation setup

To evaluate the modelling for cylindrical cells, we choose a large format lithium-iron-phosphate (LFP) 45Ah cell and take the corresponding model parameters from [29]. Specifically, for dimension-related parameters, we have $L = 198$ mm, $R_{\text{out}} = 32$ mm, $R_{\text{in}} = 4$ mm [29], and for thermo-physical parameters, we use $\rho = 2118$ kgm^{-3}, $c_p = 795$ Jkg^{-1}K^{-1} $k_r = 0.67$ Wm^{-1}K^{-1}, and $k_z = 66.6$ Wm^{-1}K^{-1} [13], [30]. For the TEC model of this cylindrical cell, C_s is found to be 48.35 JK^{-1}. Regarding R_c and R_u, it is important to note that different sets of parameters will be obtained, depending on the employed cooling scenario. For brevity, only the identified parameters for surface cooling are presented here. After validation, we obtained the parameter values as $C_c = 1079.6$ JK^{-1}, $R_c = 0.65$ KW^{-1}, and $R_u = 0.08$ KW^{-1}.

In the simulation-based study of the proposed battery models and their benchmarks, the initial temperatures, i.e., $T(r, z, 0)$, $T(x, y, 0)$, $T_c(0)$, and $T_s(0)$, and the fluid free-stream (ambient) temperature T_∞ are all chosen to be 15 °C. To facilitate the study without losing generality, the volumetric heat generation rate q is given by the heat generation model proposed by Bernardi *et al.* [31]

$$q(t) = I(t)\left[V(t) - V_{\text{ocv}}(t)\right]/V_b, \tag{A.46}$$

where I, V, V_{ocv}, and V_b are the battery current, terminal voltage, open-circuit voltage, and volume, respectively. Here, the trajectories of I and V depend on the battery usage. We select the worldwide harmonised light vehicle test procedure (WLTP) [32], which represents a wide range of vehicle driving conditions on urban, suburban, and highway roads, and scale it so that the maximum current becomes 1 C-rate for heat generation. The modified WLTP (mWLTP) heat generation profile is shown in Fig. 2(a). V_{ocv} depends on the state of charge (SoC), and the relationship between V_{ocv} and SoC has been determined using a lookup table for a typical LFP cell in [33].

Depending on where the cooling action is imposed, five different cooling scenarios are investigated, including surface cooling (SC), bottom tab cooling (bTC), bottom tab and surface cooling (bTSC), bottom and top tabs cooling (btTC), and all-tabs and surface cooling (aTSC). aTSC has all sides of the battery cell exposed to cooling and is representative of immersion cooling [34] where the cell is immersed in a dielectric fluid. Forced convection liquid cooling (for instance, using a cooling plate) via water or glycol is considered to

apply on areas where tab or surface cooling occurs. This typically has a large convection coefficient h, which is set to be 400 $\mathrm{Wm^{-2}K^{-1}}$, according to [4]. For areas not exposed to active liquid cooling, mild forced air convection occurs, resulting in a small convection coefficient specified as $h = 30$ $\mathrm{Wm^{-2}K^{-1}}$.

All the simulations are conducted on a MacBook Pro with a 2.6GHz 6-Core Intel Core i7 processor and 16GB RAM. The first type of benchmark model is solved by the finite element method with an extremely fine mesh of 3610 triangular elements and is implemented in COMSOL Multiphysics v6.0 [35]. Our developed CSG models and their TEC alternatives are implemented in MATLAB R2023b. For simplicity in simulating the CSG models, the numbers of basis functions in the radial and axial directions, as introduced in (A.15), are chosen to be equal, i.e., $M = N$, and the model order $O = N^2$.

4.3 Results of validation against the FEM model

The examination of the developed modelling method and their underlying assumptions for control-oriented applications is conducted for the battery cell and cooling scenarios specified in Section 4.2. Different CSG model orders, including $O = 1$, 4, 9, 16, and 25, were evaluated. The results for cylindrical cell modelling compared with the ground truth represented by the FEM model are presented in Fig. 2.

According to Fig. 2(d), the developed CSG model with an order of 25 can almost perfectly capture the battery cell's distributed thermal dynamics under any of the considered cooling scenarios. The resulting modelling error is bounded by 0.03 °C. Reducing the model order from 25 to 9, the maximum temperature deviation between the model's output and its ground truth is 0.4 °C within the entire test process. Similar observations are obtained for the studied pouch cell, in which the prediction error for the model with an order of 9 is less than 0.6 °C. Remarkably, this high level of model fidelity is achieved in the presence of aggressive driving and cooling. The prediction error gradually increases when the model order is further lowered. This is consistent with *Remark 2* and the applied finite-sum approximations in (A.14) and (A.23). Indeed, the model with an order of 1 or 4 can still well follow the trajectories of the true surface and core temperatures under different cooling scenarios, as seen in Fig. 2(b) and (c). Specifically, for SC, which involves one-side cooling, $O = 1$ results in a maximum error of 1.26 °C. For aTSC, which involves resolution in multiple directions (multi-side cooling), the same

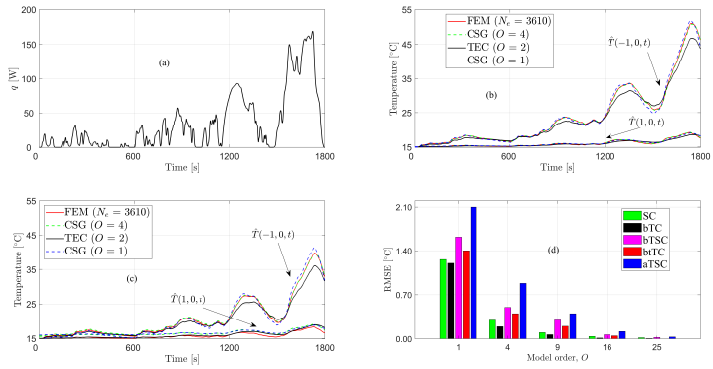

Figure 2: Validation of the developed CSG thermal model for cylindrical battery cells. (a) Heat generation profile of the mWLTP. (b) Comparison of the CSG model of orders $O = 1$ and 4, against the FEM model with the element number $N_e = 3610$ and TEC model, under surface cooling (SC). Comparisons are done for the surface and core mid-point temperatures, i.e., $\hat{T}(1, 0, t)$ and $\hat{T}(-1, 0, t)$. (c) Comparison of the CSG thermal model against FEM and TEC models under the scenario of all-tabs and surface cooling (aTSC). (d) Maximum error of the CSG model with different orders under all cooling cases.

order produces a maximum error of 2.10 °C.

The impact of current rates on model accuracy is also evaluated. To do so, the mWLTP profile is scaled by factors of two and three, resulting in maximum C-rates of 2C and 3C, respectively. Under these highly aggressive operating conditions and the SC scenario, the errors of the CSG model with different orders are summarized in Table 1. The results demonstrate a clear necessity to increase the model order for accurately capturing battery cell thermal dynamics at elevated current rates due to pronounced thermal gradients and intensified spatial interactions. Considering a maximum permissible error of 1.5 °C, model orders $O = 1$, 9, and 16 are sufficient to faithfully model the thermal behaviour in the mWLTP, mWLTP×2, and mWLTP×3 scenarios, respectively. It is worth mentioning that with this relatively small magnitude of local errors obtained under all considered cooling scenarios and extreme battery usage conditions, the CSG model proves suitable for a wide range of battery management applications.

Table 1: Effect of current rates on model accuracy.

Current profile	O=1	O=4	O=9	O=16	O=25
mWLTP	1.26	0.46	0.13	0.09	0.03
mWLTP×2	3.83	2.36	1.11	0.98	0.73
mWLTP×3	4.89	3.62	1.85	1.36	1.01

The above results consequently verify the effectiveness of the modelling approach proposed in Section 3. Specifically, this confirms that the developed adaptable approach for the particular solution T_p can efficiently handle any abrupt changes in the operating conditions and spatial gradients distributed across the battery cell; the CSG method can well reproduce the state dynamics governed by the PDE and its homogeneous solution. Furthermore, it numerically proves the conclusions drawn in *Remark 4*. Namely, for the particular solution problem in (A.13), the response to the cooling power applied at individual boundaries satisfies the superposition principle and can be expressed as a linear combination of the cooling power applied at these boundaries.

In addition to the above validation against the high-fidelity FEM model implemented in COMSOL, our future work will explore direct experiments to evaluate our model under various cooling configurations.

4.4 Results of comparison against the TEC model

To investigate the efficacy and applicability of the CSG model validated in Section 4.3, we compare it with the second type of benchmark defined in Section 4.1, i.e., the TEC model. In the TEC model, the average temperature is derived as $T_{\text{mean}}^{\text{TEC}} = (T_s+T_c)/2$, and the radial gradient is defined by $\delta T_{max}^{r} = (T_c - T_s)/(R_{\text{out}} - R_{\text{in}})$. To possibly fully reveal their predictive capability, particularly under highly dynamic operating conditions, we introduce a pulsed power load profile to generate the heat $q(t)$, as shown in Fig. 3(a). This heat trajectory is characterised by high heat generation rates but a relatively short duration. For completeness, the TEC model was also included in the WLTP validation scenario shown in Fig. 2(b) and (c). The comparative results between the FEM, CSG, and TEC models under the pulsed power profile with surface cooling are presented in Fig. 3.

It is found that the CSG model of $O = 1$ is capable of closely following the true trajectories of the mean temperature, max temperature, and max gradient in the entire test cycle. By contrast, in general, the TEC model with an order of 2 significantly underestimates all three considered indicators, resulting in temperature errors of several degrees. To observe more closely, the model's error for $T_{\max}$ is much larger than that for T_{mean}, and it increases when there are abrupt changes in the heat generation. This observation is consistent with the comparison under the WLTP profile in Fig. 2(b) and (c), where the TEC model similarly underestimates the core and surface temperatures. This means that compared to our developed model, the TEC model is unreliable for predicting the local thermal behaviour and may be limited to applications with only low heat generation in the battery cell. Particularly, this model cannot be used for battery safety management, where the maximum temperature must be accurately estimated and constrained, and for thermal uniformity control, where the information on the maximum gradients is critical.

Computational efficiency is of great importance for model-based applications, particularly real-time thermal optimization and control of large-scale battery packs. With this in mind, we examine the computational times of the developed model with various orders and compare them with that of the TEC and FEM models. To have a fair comparison, we run each model ten times and compute their average simulation times in milliseconds (ms), with the results listed in Table 2. Clearly, the CSG model with an order O=1 is more efficient than its TEC counterpart, leading to a reduction of 28.7% in the required time. The model with higher orders takes longer to simulate, but its time consumption remains within the same order of magnitude as that of the TEC model. As expected, the FEM model takes a longer computational time, about 8s to compute the solution, making it impractical for real-time control applications.

The above results indicate that the CSG model with O=1 offers distinct advantages in both accuracy and computational efficiency. Therefore, it can readily replace the TEC model in existing BMSs to improve battery thermal performance in the real world. In addition, it may be computationally feasible to use a high-order CSG model, e.g., the order of 4, within a BMS, allowing a significantly more advanced management of battery safety and lifetime.

Table 2: Average computational times of CSG model of different orders compared with the TEC model.

Model type	FEM	TEC	O=1	O=4	O=9	O=16	O=25
Time [ms]	8000	307	219	431	517	652	746

5 Model-based applications

The developed models for cylindrical and pouch battery cells, featuring high fidelity and efficiency, can be effectively applied to various model-based applications. These include rapid analysis of different cooling scenarios, real-time closed-loop temperature control, and thermal assessment of cell dimensions, as demonstrated in this section. Here, based on the model of 9th order, simulations are conducted under the mWLTP profile on the battery cell and cooling scenarios specified in Section IV-B.

5.1 Model-based analysis for different cooling scenarios

For a given battery system and usage profiles, selecting the most appropriate cooling scenario can make a big difference to all its key properties. We analyze thermal performance merits, T_{mean}, $T_{\max}$, $\delta T^r_{\max}$, $\delta T^z_{\max}$, and $\Delta T = T_{\max} - T_{\min}$, under the five cooling scenarios defined in Section 4.2, with the results for cylindrical cells presented in Table 3.

In terms of T_{mean}, bTC gives rise to the highest value at 22.77 °C, followed by SC with the second-highest value. This is because the average temperature is primarily influenced by the cooling area, and both bTC and SC have only one side exposed to cooling. With the largest cooling area, aTSC demonstrates a high capability of heat removal, leading to the lowest values of T_{mean} and $T_{\max}$ among all cases. However, aTSC does not achieve the lowest thermal gradient. In contrast, btTC provides the lowest $\delta T^r_{\max}$ at the local level and the lowest ΔT across the entire cell. This can be attributed to symmetric cooling from the cell's top and bottom sides, and that all electrodes get access to cooling at the same time under btTC because of the rolled electrode layer structure. On the contrary, in the case of SC, the outer layers get cooled before the internal layers, leading to inhomogeneous cooling and consequently

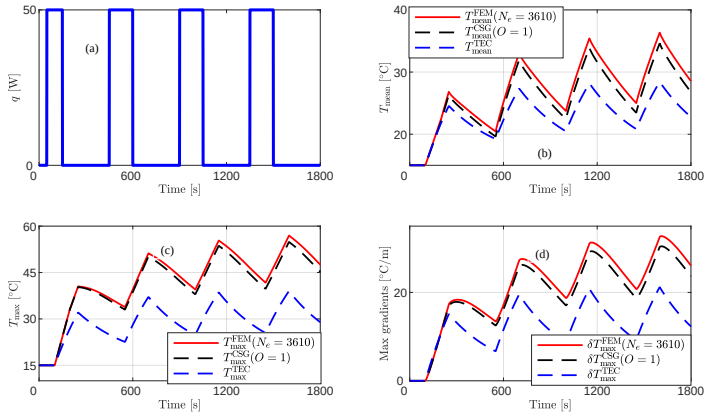

Figure 3: Comparison between the CSG and TEC models under surface cooling. (a) Pulse heat generation profile. (b) Evaluation of the mean temperature responses of the two models, i.e., $T^{\text{CSG}}_{\text{mean}}$ and $T^{\text{TEC}}_{\text{mean}}$, against their ground truth. (c) Comparison of their max temperature responses, which are $T^{\text{CSG}}_{\text{max}}$ and $T^{\text{TEC}}_{\text{max}}$, (d) Maximum thermal gradients of the CSG and TEC models, i.e., $\delta T^{\text{CSG}}_{\text{max}}$ and $\delta T^{\text{TEC}}_{\text{max}}$.

the highest gradient of 32.1 °C/mm.

In all scenarios, it is observed that the cylindrical cell experiences significantly higher thermal gradients along the radial direction compared to the axial direction. This is due to the radial thermal conductivity being two orders of magnitude lower than that of the axial direction, creating a significant bottleneck in heat transfer to the surface. This, combined with the absence of cooling applied to the core, explains why the core temperature is the highest. Increasing the cooling flow rate and using battery materials with higher thermal conductivity can enhance cooling efficiency and help reduce the core temperature.

Moreover, it is evident that no single cooling scenario consistently outperforms the others across all performance metrics. Cooling only one of the cell's tabs, as represented by bTC and bTSC, is unfavourable due to its low removed heat and the significant $\delta T^{z}_{\text{max}}$. When the cell's surface is involved in cooling, as in SC, bTSC, and aTSC, a considerable $\delta T^{r}_{\text{max}}$ is observed. Among SC, btTC, and aTSC, each scenario performs the best in one or two specific

Table 3: Thermal performance merits under the five scenarios.

Scenarios	T_{mean}	T_{max}	δT^r_{max}	δT^z_{max}	ΔT
SC	20.80	51.00	32.10	**0**.44	36.00
bTC	22.77	54.01	14.05	8.94	39.01
bTSC	19.38	46.18	27.55	6.70	31.22
btTC	20.29	43.36	**9**.32	3.40	**2**3.37
aTSC	**1**8.30	**3**9.68	21.42	3.17	24.70

metrics. For example, btTC offers the most homogeneous cooling effect but results in a slightly higher T_{mean}, whereas aTSC lowers T_{mean} at the cost of increasing δT^r_{max}. This indicates that different performance metrics may be in competition within each cooling scenario.

5.2 Closed-loop temperature control

The performance trade-offs revealed in Section 5.1 necessitate the selection of a scenario to satisfy application-specific requirements. In this regard, an interesting question is: which cooling scenario produces the lowest thermal gradient while maintaining a fixed average temperature?

To answer this question and demonstrate the application of our developed battery thermal models, we implement a simple feedback control strategy based on multiple proportional-integral (PI) controllers, as framed in Fig. 4. This control scheme is designed to maintain the average cell temperature, T_{mean}, at a predefined set-point, $T^{\text{ref}}_{\text{mean}}$, which is specified as 20 °C. For each cooling scenario, the control in the framework of Fig. 4 is tailored so that individual PI controllers regulate the cooling power directed to specific sides of the cell. For example, in the case of aTSC, all the top, bottom, and surface controllers are active, while for SC, only the surface controller is engaged, with $u_t(t) = u_b(t) = 0$. To simplify the control implementation, we assume a uniform distribution of cooling power across all active areas.

Since T_{mean} is not measurable during real-world battery operation, it is imperative to incorporate a state estimator within the designed control strategy. In this preliminary study, we design an open-loop estimator that mirrors the model (A.38)–(A.40) for cylindrical battery cells. However, the state-space model allows any standard state-estimator such as the Kalman filter. Beyond

its high accuracy and efficiency, this thermal model enables precise control of cooling applied to any side of the cell or combinations thereof, which is a capability not offered by existing models such as those in [17], [18]. This flexibility makes our model particularly suitable for battery temperature estimation and control in sophisticated cooling scenarios. For case studies, the

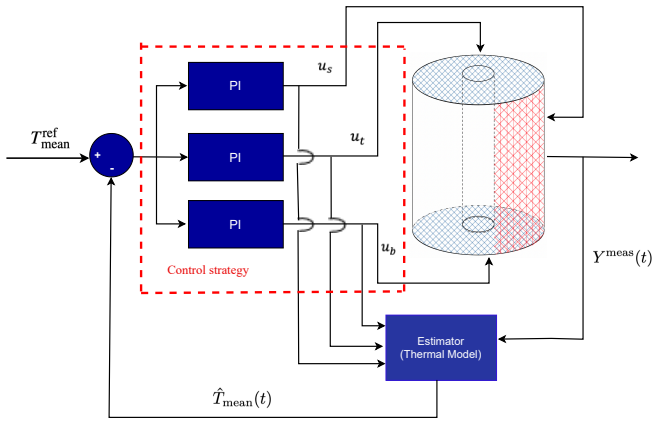

Figure 4: PI-based control scheme that tracks a given reference for the cylindrical cell's average temperature. $T^{\text{ref}}_{\text{mean}}$ and $\hat{T}_{\text{mean}}$ represent the set-point and estimate of the average temperature, respectively, and Y^{meas} is the cell's measurable outputs. The cooling power applied to the surface, top, and bottom, which are u_s, u_t, and u_b defined in (A.8), is regulated by their corresponding PI controllers. Their design parameters are set based on the specific cooling scenario.

original WLTP profile is scaled so that the maximum current becomes 1, 2, 3, and 4C, respectively. The control results are given in Fig. 5. To save space, in Fig. 5(a) and (b), only the trajectories obtained in aTSC and the 1C case are depicted.

Based on the results in Fig. 5(c) and (d), one can readily find the answer to the question raised. When the average cell temperature is maintained at 20 °C, btTC best balances the thermal gradients along the r- and z-directions and produces the most homogeneous thermal distribution across the cylindrical cell. Specifically, in btTC, the average thermal gradient along the r-direction, denoted by $\delta T^{r}_{\text{mean}}$, is substantially smaller than those observed in

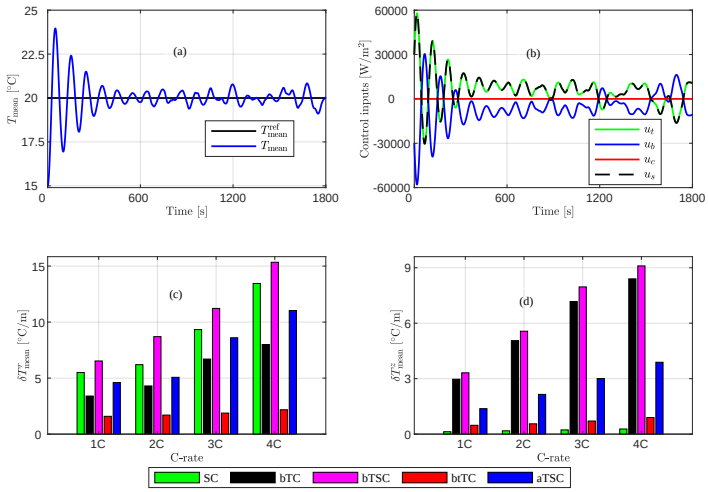

Figure 5: Results of battery thermal control of the cylindrical cell. (a) and (b) show the results obtained in aTSC and the 1C case of WLTP. (c) and (d) compare the mean values of the radial and axial thermal gradients, respectively, for different cooling scenarios and WLTP cases.

all other cooling scenarios. Regarding the average thermal gradient along the z-direction, btTC is comparable to the best-performing scenario, i.e., SC.

In terms of the control's behaviour, as demonstrated in Fig. 5(a) and (b), substantial control efforts are initially required to attenuate the large tracking errors. Once T_{mean} approaches the set-point, the magnitude of u_t, u_b, and u_s decreases significantly. Note that the temperature reference of 20 °C is used here purely for demonstration purposes—to showcase the responsiveness and precision of the proposed model under stringent conditions. Beyond the PI controller, future work will explore more sophisticated control strategies to systematically enhance the control performance, in which multiple control objectives and potential constraints will be explicitly addressed. Within the control strategy, the new models developed in Section 3 can serve as efficient tools for accurately simulating and predicting battery thermal states.

5.3 Model-based evaluation of cylindrical cell aspect ratio

A cell's geometry includes critical design parameters that dictate its intrinsic thermal resistance and interface with the thermal management system. We utilised our proposed model to investigate this relationship by performing a parameter sweep on the cell aspect ratio, i.e. L/R_{out}. During this study, the total cell volume V_b was held constant so that the heat generation trajectory $q(t)$ under the mWLTP remained identical across all geometries. Fig. 6 shows that the optimal aspect ratio is not universal; instead, it depends strongly on the cooling strategy employed.

The observed behaviours can be grouped into two physical regimes. The first regime includes the four scenarios involving tab cooling, namely, bTC, bTSC, btTC, and aTSC. In these cases, the thermal behaviour is conduction-limited. The trends are monotonic: short and bulky cells corresponding to low aspect ratios yield lower mean temperatures and smaller gradients. As the aspect ratio increases, performance degrades because tall, slim cells introduce a long, high-resistance axial conduction path. This impedes heat transfer from the cell's core to the actively cooled tabs, causing heat accumulation. Among these scenarios, aTSC produces the lowest mean temperatures across all aspect ratios due to its larger total cooling area, whereas btTC achieves the lowest radial gradients because of the symmetric cooling applied at both the top and bottom tabs. Axial gradients, although small compared to radial gradients, are non-zero in asymmetrical cooling cases (bTC and bTSC) where $h_t \neq h_b$, and these gradients grow with an increasing aspect ratio. In contrast, for symmetric cases (btTC, SC, and aTSC), where $h_t = h_b$, the axial gradients are identically zero.

The second regime consists of lateral surface-only cooling (SC), which exhibits a different trend. Here, the system is surface-area-limited. For a fixed volume $V_b = \pi(R_{\text{out}}^2 - R_{\text{in}}^2)L$, the available cooling surface area is $2\pi R_{\text{out}}L$. The corresponding surface-area-to-volume ratio, $2R_{\text{out}}/(R_{\text{out}}^2 - R_{\text{in}}^2)$, is maximised when radius R_{out} is small, i.e., for tall, slim cells. The simulation results confirm this: after a brief rise at very low L/R_{out}, both mean temperature and radial gradients decrease almost linearly with increasing aspect ratio. While performance improves for tall cells due to their larger heat-exchange area, the absolute magnitudes of temperature and radial gradients remain significantly higher than those achieved with tab-involved cooling. Notably, the mean temperature under SC becomes lower than in bTC and btTC beyond aspect-ratio

crossover points of roughly 4 and 5, respectively, but even at these crossovers the corresponding thermal gradients under SC remain substantially larger. Overall, SC benefits tall, slender geometries, but is generally insufficient as a standalone strategy when uniformity or minimal gradients are required.

These insights also allow us to evaluate commonly used cylindrical cell formats in today's market. The popular 18650, 21700, 26650, and 4680 cells have corresponding L/R_{out} ratios of 7.22, 6.67, 5.00, and 3.48, respectively. Based on our analysis, the 4680 format, thanks to its relatively low aspect ratio, offers improved thermal behaviour for all cooling strategies involving tab contact, since its short axial conduction path promotes efficient heat removal. Conversely, high-aspect-ratio formats such as the 18650 and 21700 benefit more from SC due to their larger heat-exchange area, but their performance still remains inferior to tab-based approaches, especially when spatial uniformity and gradient minimisation are critical.

Taken together, these findings indicate that no single cylindrical cell geometry is universally optimal. Instead, the ideal aspect ratio must be selected in conjunction with the intended cooling configuration: tab-based strategies favour short, bulky cells, while surface-only strategies favour tall, slim cells—though the latter still struggle to match the thermal uniformity achievable with tab cooling. Therefore, cylindrical cell design should be co-optimised with the thermal management strategy rather than determined in isolation.

6 Conclusion

This article has developed a new battery modelling framework for spatially-resolved thermal dynamics based on Chebyshev spectral-Galerkin approach and model component decomposition. The obtained library of models is well-suited for online thermal performance optimisation, thanks to its ability to independently control the tab and surface cooling channels of the battery. This framework has been presented for both cylindrical and pouch cells and thoroughly evaluated through various combinations of tab and surface cooling cases under real-world vehicle driving profiles.

By using four states, the model accurately predicts the local thermal behaviour, even under aggressive vehicle driving and cooling scenarios. The model with only one state, in particular, is both more accurate and computationally efficient than the widely utilised two-state thermal equivalent circuit

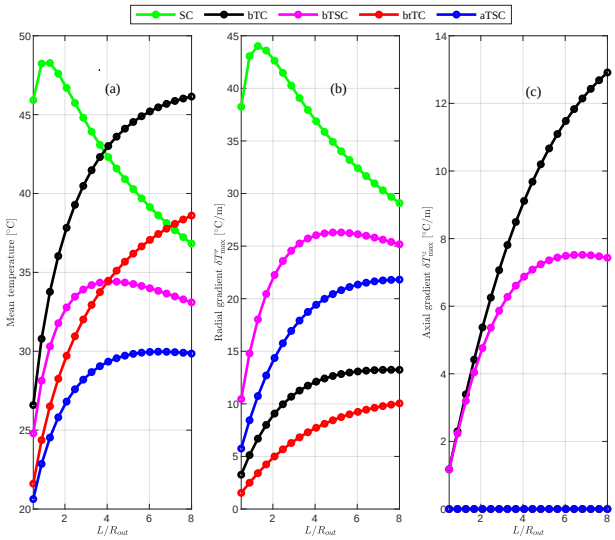

Figure 6: Evaluation of the thermal performance of cylindrical cells with varying dimensions under five cooling scenarios. (a) Evolution of the mean temperature T_{mean}. (b) Evolution of the maximal thermal gradients in the radial direction δT^r_{max}. (c) Evolution of the maximal thermal gradients in the axial direction δT^z_{max}.

(TEC) model, with a 28.7% reduction in computational time. Even as the number of states increases, it remains in the same computational order of magnitude as the TEC model, suggesting it could readily replace the TEC model in existing battery management systems. Closed-loop temperature control of cylindrical cells under different cooling scenarios demonstrated that exclusively cooling through the top and bottom tabs yields the lowest thermal gradients under various C-rates. Furthermore, the model's utility extends to the evaluation of cell geometry on thermal performance. Dimension-perturbation analysis showed that no single aspect ratio is universally optimal and cylindrical cells should be co-designed with the intended cooling strategy rather than optimised in isolation.

Future work will focus on extending the thermal model to pack-level assemblies. Beyond batteries, the versatile modelling framework can be extended to other systems governed by partial differential equations, making it a valuable

tool across multiple engineering disciplines.

References

[1] S. Cevik, "Climate change and energy security: The dilemma or opportunity of the century?" International Monetary Fund, IMF Working Paper, 2022.

[2] S. Ma et al., "Temperature effect and thermal impact in lithium-ion batteries: A review," *Prog. Nat. Sci.: Mater. Int.*, vol. 28, no. 6, pp. 653–666, 2018.

[3] J. Lin, X. Liu, S. Li, C. Zhang, and S. Yang, "A review on recent progress, challenges and perspective of battery thermal management system," *Int. J. Heat Mass Transf.*, vol. 167, p. 120 834, 2021.

[4] P. R. Tete, M. M. Gupta, and S. S. Joshi, "Developments in battery thermal management systems for electric vehicles: A technical review," *J. Energy Storage*, vol. 35, p. 102 255, 2021.

[5] Y. Zhao, L. B. Diaz, Y. Patel, T. Zhang, and G. J. Offer, "How to cool lithium ion batteries: Optimising cell design using a thermally coupled model," *J. Electrochem. Soc.*, vol. 166, no. 13, A2849–A2859, 2019.

[6] I. A. Hunt, Y. Zhao, Y. Patel, and G. Offer, "Surface cooling causes accelerated degradation compared to tab cooling for lithium-ion pouch cells," *J. Electrochem. Soc.*, vol. 163, no. 9, A1846, 2016.

[7] Y. Zhao, Y. Patel, T. Zhang, and G. J. Offer, "Modeling the effects of thermal gradients induced by tab and surface cooling on lithium ion cell performance," *J. Electrochem. Soc.*, vol. 165, no. 13, A3169–A3178, 2018.

[8] O. Dondelewski et al., "The role of cell geometry when selecting tab or surface cooling to minimise cell degradation," *eTransportation*, vol. 5, p. 100 073, 2020.

[9] S. Li et al., "Optimal cell tab design and cooling strategy for cylindrical lithium-ion batteries," *J. Power Sources*, vol. 492, p. 229 594, 2021.

[10] C. Bolsinger and K. P. Birke, "Effect of different cooling configurations on thermal gradients inside cylindrical battery cells," *J. Energy Storage*, vol. 21, pp. 222–230, 2019.

[11] T. Ahmad, A. Mishra, S. Ghosh, and C. S. Casari, "Identifying efficient cooling approach of cylindrical lithium-ion batteries," *Energy Technol.*, vol. 10, no. 2, p. 2 100 888, 2022.

[12] G. K. Peprah, F. Liberati, F. Altaf, G. Osei-Dadzie, A. Di Giorgio, and A. Pietrabissa, "Optimal load sharing in reconfigurable battery systems using an improved model predictive control method," in *Proc. 29th Mediterranean Conf. Control Autom. (MED)*, IEEE, 2021, pp. 979–984.

[13] X. Lin et al., "A lumped-parameter electro-thermal model for cylindrical batteries," *J. Power Sources*, vol. 257, pp. 1–11, 2014.

[14] V. Srinivasan and C. Wang, "Analysis of electrochemical and thermal behavior of Li-ion cells," *J. Electrochem. Soc.*, vol. 150, no. 1, A98–A106, 2003.

[15] J. Shen, T. Tang, and L.-L. Wang, *Spectral methods: algorithms, analysis and applications*. Springer Science & Business Media, 2011, vol. 41.

[16] L. N. Trefethen, *Spectral methods in MATLAB*. SIAM, 2000.

[17] R. R. Richardson, S. Zhao, and D. A. Howey, "On-board monitoring of 2-D spatially-resolved temperatures in cylindrical lithium-ion batteries: Part i. low-order thermal modelling," *J. Power Sources*, vol. 326, pp. 377–388, 2016.

[18] X. Hu, W. Liu, X. Lin, Y. Xie, A. M. Foley, and L. Hu, "A control-oriented electrothermal model for pouch-type electric vehicle batteries," *IEEE Trans. Power Electron.*, vol. 36, no. 5, pp. 5530–5544, 2020.

[19] D. W. Hahn and M. N. Özisik, *Heat conduction*. John Wiley & Sons, 2012.

[20] T. G. Tranter, R. Timms, P. R. Shearing, and D. Brett, "Communication—prediction of thermal issues for larger format 4680 cylindrical cells and their mitigation with enhanced current collection," *J. Electrochem. Soc.*, vol. 167, no. 16, p. 160 544, 2020.

[21] K.-J. Lee, G. Kim, and K. Smith, "Three-dimensional thermal and electrochemical model for spirally wound large-format lithium-ion batteries," in *218th ECS Meeting*, Abstract 1114, 2010.

[22] S. J. Farlow, *Partial differential equations for scientists and engineers*. Courier Corporation, 1993.

[23] R. K. Nagle, E. B. Saff, A. D. Snider, and B. West, *Fundamentals of differential equations and boundary value problems*. Addison-Wesley New York, 1996.

[24] E. H. Doha and A. H. Bhrawy, "An efficient direct solver for multidimensional elliptic robin boundary value problems using a legendre spectral-galerkin method," *Comput. & Math. Appl.*, vol. 64, no. 4, pp. 558–571, 2012.

[25] F. Auteri, N. Parolini, and L. Quartapelle, "Essential imposition of neumann condition in galerkin–legendre elliptic solvers," *J. Comput. Phys.*, vol. 185, no. 2, pp. 427–444, 2003.

[26] Y. Xie et al., "An improved resistance-based thermal model for a pouch lithium-ion battery considering heat generation of posts," *Appl. Therm. Eng.*, vol. 164, p. 114 455, 2020.

[27] X. Lin et al., "Online parameterization of lumped thermal dynamics in cylindrical lithium ion batteries for core temperature estimation and health monitoring," *IEEE Trans. Control Syst. Technol.*, vol. 21, no. 5, pp. 1745–1755, 2013.

[28] C. Zou, A. Klintberg, Z. Wei, B. Fridholm, T. Wik, and B. Egardt, "Power capability prediction for lithium-ion batteries using economic nonlinear model predictive control," *J. Power Sources*, vol. 396, pp. 580–589, 2018.

[29] V. Roscher et al., "Synchronisation using wireless trigger-broadcast for impedance spectroscopy of battery cells," in *IEEE Sensors Appl. Symp.*, IEEE, 2015, pp. 1–6.

[30] M. Fleckenstein, S. Fischer, O. Bohlen, and B. Bäker, "Thermal impedance spectroscopy–a method for the thermal characterization of high power battery cells," *J. Power Sources*, vol. 223, pp. 259–267, 2013.

[31] D. Bernardi, E. Pawlikowski, and J. Newman, "A general energy balance for battery systems," *J. Electrochem. Soc.*, vol. 132, no. 1, p. 5, 1985.

[32] P. Mock et al., "The WLTP: How a new test procedure for cars will affect fuel consumption values in the EU," *Intern. council on clean transp.*, vol. 9, no. 3547, 2014.

[33] J. Li, J. K. Barillas, C. Guenther, and M. A. Danzer, "A comparative study of state of charge estimation algorithms for LiFePO4 batteries used in electric vehicles," *J. Power Sources*, vol. 230, pp. 244–250, 2013.

[34] C. Roe et al., "Immersion cooling for lithium-ion batteries–a review," *J. Power Sources*, vol. 525, p. 231 094, 2022.

[35] COMSOL AB, *Introduction to COMSOL multiphysics*, COMSOL AB.

PAPER B

Model Predictive Coolant Allocation for Integrated Tab–Surface Cooling of Battery Cells

Godwin K. Peprah, Torsten Wik, Masood Tamadondar, Changfu Zou

Manuscript, 2026. Submitted to IEEE Transactions on Control Systems Technology for review. Available as *arXiv preprint arXiv: 2607.10872.* The version reproduced in this thesis is the manuscript dated 12 July 2026. A later revised or published version may differ.

© 2026 The Authors

The layout has been revised.

Abstract

Battery electrical tab cooling is effective at reducing internal thermal gradients by exploiting the high thermal conductivity of the current collectors, whereas surface cooling is effective at reducing temperature rise because of its large heat transfer area. Using either strategy alone, however, limits the achievable trade-off between thermal uniformity and temperature rise reduction. This work proposes an integrated tab-surface cooling (ITSC) system in which coolant is dynamically allocated among the lateral surface and tab channels. The allocation is formulated as an optimal control problem in which the battery temperature is regulated towards a desired reference and thermal gradients are minimised. To support this formulation, a first-principles coolant model is developed and coupled with battery and valve-actuation models. The resulting optimal coolant-allocation problem is solved using a computationally efficient real-time iteration model predictive control (RTI-MPC) scheme, with a nonlinear MPC serving as a closed-loop performance benchmark. Evaluation results under realistic driving conditions showed that RTI-MPC reproduces the nonlinear MPC thermal response with absolute errors below 0.0035 °C while reducing the computational cost from several seconds to 19.3 ms, indicating strong potential for real-time implementation. Additionally, evaluation of the proposed ITSC system against conventional cooling configurations demonstrates that ITSC achieves the best overall trade-off between temperature regulation and thermal gradient reduction.

1 Introduction

Thanks to their long extended cycle life, high energy and power density compared to other rechargeable batteries, Lithium (Li)-ion batteries are the dominant energy storage technology for numerous applications, including electric vehicles (EVs), stationary energy storage systems, and portable electronics [1]. Despite these advantages, their safety, power and energy capability, charge ac-

ceptance, and lifespan are strongly dependent on temperature [2]. Operating at elevated temperatures accelerates irreversible degradation and increases the risk of thermal runaway, while large temperature differences within a cell, known as thermal gradients, could lead to uneven ageing and localised mechanical stress within the electrodes [2], [3]. Consequently, a sophisticated battery thermal management system (BTMS) [4] is not only beneficial but essential for ensuring safe, long service life, and reliable operation of these batteries.

BTMS can be broadly categorised as passive or active. Passive systems rely on natural heat dissipation mechanisms, whereas active systems employ auxiliary devices such as fans or pumps to enhance heat transfer. Established BTMS include air cooling, liquid cooling, phase-change materials, and immersion cooling [4], [5], [6]. Air cooling is simple and inexpensive but offers limited heat removal capacity, limiting its suitability for high-power applications. Liquid cooling, typically using water-glycol mixtures, provides much higher heat transfer coefficients and is widely adopted in EVs. Phase-change materials passively buffer transient heat generation through latent heat absorption, but their low thermal conductivity and limited ability to dissipate accumulated heat can restrict long-term thermal regulation. Immersion cooling achieves a more uniform temperature distribution by directly submerging cells in dielectric fluids at the expense of increased system complexity, sealing requirements, and material compatibility constraints. Each of these approaches involves a trade-off in thermal performance, cost, weight, and complexity, driving continuous research efforts into hybrid and optimised strategies.

Within liquid-based BTMS, the design of cell-level cooling interfaces critically influences thermal behaviour. Two methods that have received particular attention are lateral surface cooling and electrical connection tab (terminals) cooling [7]. For Li-ion pouch cells, Hunt *et al.* [8] and Zhao *et al.* [9] showed that due to the larger available lateral heat transfer area, surface cooling is generally more effective at limiting the temperature rise of the cell, particularly under high current operation. In contrast, tab cooling removes heat directly through the current collectors, leveraging their high thermal conductivity to homogenise the temperature distribution across the electrode layers, thereby reducing internal thermal gradients. Dondelewski *et al.* [10] in a follow-up study on a large-format lithium-iron-phosphate (LFP) pouch cell, reported similar results. Under a 2C rate over 103 cycles, surface cooling resulted in a

lower temperature rise and slower degradation, whereas tab cooling yielded a higher temperature rise but improved thermal homogeneity and usable capacity. Similar trade-offs were reported for cylindrical cells in [11], [12] where tab cooling reduced internal temperature non-uniformities relative to surface cooling, albeit at the cost of higher temperature rise. Complementary findings in [13] further showed that surface cooling produced a faster overall temperature drop owing to the larger exposed area, while tab cooling delivered superior thermal uniformity, with peak-to-peak temperatures approximately half those observed under surface cooling in their test cell.

These studies collectively indicate that surface and tab cooling offer distinct but complementary benefits. However, each method has inherent limitations, i.e., neither strategy can simultaneously minimise temperature rise and internal gradients alone, particularly under demanding operating conditions such as fast-charging or elevated ambient temperatures. To date, there is no battery control framework in the state-of-the-art literature that systematically investigates the optimal integration of these two cooling strategies for advanced battery thermal management.

To bridge this gap, we propose a control framework for integrated cooling, with the main contributions summarised as follows:

1. A control-oriented dynamic coolant model for multi-channel battery cooling is developed and coupled with existing battery thermal and electrical models, enabling flow-based actuation of the battery cooling process. Additionally, a valve-actuation model is developed and incorporated to capture the dynamic routing of coolant across the channels, thereby reflecting actuator-level behaviour.

2. The integration of lateral surface and tab cooling is formulated as an optimal control problem, where the available coolant flow is dynamically allocated across multiple cooling channels to regulate the cell's temperature rise toward a desired reference while suppressing internal thermal gradients.

3. The resulting optimal coolant-allocation problem is solved for the first time within a real-time iteration model predictive control (RTI-MPC) framework to improve computational efficiency and support real-time implementation.

The proposed control framework is systematically evaluated against conventional cooling configurations under real-world driving conditions to assess the achievable trade-offs between temperature regulation and thermal uniformity. Although it is demonstrated here using a cylindrical cell for clarity, the proposed framework can equally be applicable to pouch and prismatic cell formats with their appropriate thermal and coolant models.

The remainder of this article is organised as follows. Section 2 presents the modelling of the battery and coolant system. Section 3 details the proposed controller framework. Section 4 discusses the results. Finally, Section 5 concludes the work.

2 Control-oriented battery cooling system modelling

This section presents the architecture for the proposed integrated battery cooling system and the models of the various subsystems therein. Fig. 1 shows a representative example of an EV BTMS for a pack containing cylindrical cells. The system comprises a coolant circulation loop in which a water-glycol mixture, driven by a pump, flows through channels embedded within the high-voltage battery pack. Cooling plates attached to the curved lateral surface or base of the cells facilitate heat absorption and removal. The absorbed heat is subsequently transferred to a heat exchanger, which dissipates it to the ambient air and/or a secondary cooling circuit, such as a heat pump, depicted in the green-dashed box in Fig. 1.

The proposed cooling architecture, shown in the red-dashed box of Fig. 1 and referred to as the Integrated Tab-Surface Cooling (ITSC) system, employs three dedicated cooling channels targeting the curved lateral surface, top, and bottom sides of the cell, respectively. It has been demonstrated in [12], [14] that cylindrical cell designs with tab areas much smaller than the top/bottom areas are not efficient for battery cooling merely from the tabs. Thus, this work considers the entire area of the top/bottom as the tab, forming the so-called all-tab or tab-less cell design. The coolant pump is rated to provide a maximum volumetric flow rate $\dot{\mathcal{V}}_{\max}$, which is distributed across the channels according to

$$\dot{\mathcal{V}}_s(t) + \dot{\mathcal{V}}_t(t) + \dot{\mathcal{V}}_b(t) \leq \dot{\mathcal{V}}_{\max}, \tag{B.1}$$

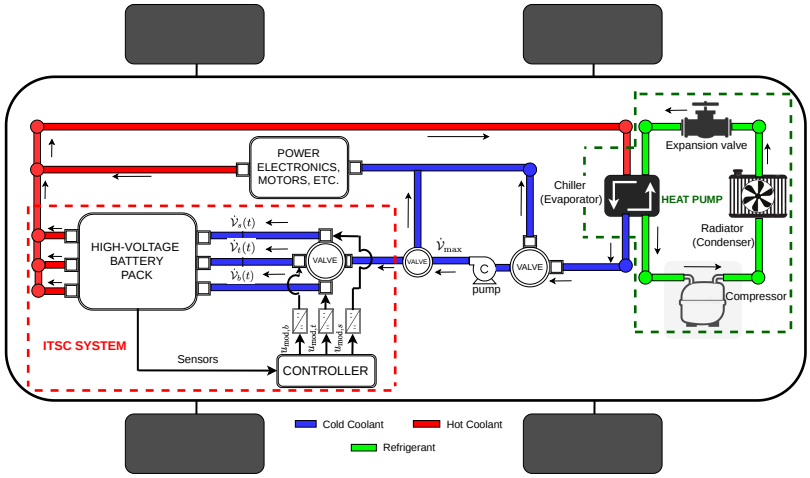

Figure 1: BTMS of an EV comprising a heat pump with its refrigeration cycle (green-dashed box) and the proposed ITSC system (red-dashed box). Centrifugal pumps circulate the water-glycol mixture through the cooling loop. The red-coloured channels represent hot coolant, the blue, the cooled coolant, and the green, the refrigerant. It is assumed that the control of the chiller manages to attenuate temperature variations in the return flow from the batteries.

where $\dot{\mathcal{V}}$ denotes the volumetric flow rates with subscripts s, t, and b denoting the cell's lateral surface, top, and bottom channels, respectively. This inequality allows partial utilisation of the pump capacity when thermal loads are low. A three-way solenoid valve [15], [16], actuated via DC–DC power converters [17], [18], regulates the distribution of coolant among the channels. The volumetric flow rates $\dot{\mathcal{V}}$ are treated as the manipulated inputs, and their dynamic allocation constitutes the primary control action. For modelling purposes, coolant transport delays between the pump and the cell are neglected, as the coolant transport dynamics are significantly faster than the dominant time constant of the battery thermal dynamics.

The proposed ITSC system integrates the battery, coolant, and valve actuation within a unified control-oriented framework, with each component modelled in the remainder of this section.

2.1 Battery modelling

The battery is modelled as a coupled electrothermal system, as described in the following subsections.

Battery thermal dynamics

In this subsection, we present the two-dimensional (2D) reduced-order thermal model for the cylindrical cell developed in our previous work [19], [20]. The governing heat conduction model and its reduced-order state-space representation via the Chebyshev spectral-Galerkin (CSG) method are briefly summarised here. They provide the cell thermal dynamics used in the subsequent coupled modelling framework.

The temperature distribution $T(r, z, t)$ within the cylindrical cell, defined over the radial coordinate r, axial coordinate z, and time t, is governed by the following 2D boundary value problem [21]:

$$\begin{aligned} &\rho_b c_{p,b} \frac{\partial T(r,z,t)}{\partial t} - k_{b,r} \frac{\partial^2 T(r,z,t)}{\partial r^2} \\ &-k_{b,r} \frac{\partial T(r,z,t)}{r \partial r} - k_{b,z} \frac{\partial^2 T(r,z,t)}{\partial z^2} = q(t), \end{aligned} \tag{B.2}$$

subject to the non-homogeneous convection boundary conditions given by

$$-k_{b,r} \frac{\partial T}{\partial r} = h_s \left(T - T_{s,\infty}\right) \quad \text{at } r = R_{\text{out}}, \tag{B.3a}$$

$$k_{b,r} \frac{\partial T}{\partial r} = h_c \left(T - T_{c,\infty}\right) \text{ at } r = R_{\text{in}}, \tag{B.3b}$$

$$-k_{b,z} \frac{\partial T}{\partial z} = h_t \left(T - T_{t,\infty}\right) \quad \text{at } z = L, \tag{B.3c}$$

$$k_{b,z} \frac{\partial T}{\partial z} = h_b \left(T - T_{b,\infty}\right) \text{ at } z = 0. \tag{B.3d}$$

where ρ_b, $c_{p,b}$, k_b, h, and T_∞ represent the volume-average density, volume-average specific heat capacity, thermal conductivity, convection coefficient, and fluid free-stream temperature, respectively. R_{in} and R_{out} are the cylindrical cell's inner and outer radius, L is the length, and $q(t)$ is the heat generation rate. The subscripts s, c, t, and b denote the cell's lateral surface, core, top, and bottom, respectively. For cylindrical cells, there is no cooling in the core region, i.e., $h_c = 0$; thus, the radial temperature gradient in (B.3b) is zero.

Applying the CSG method to (B.2)–(B.3) yields a finite-dimensional system of coupled ordinary differential equations (ODEs) of the form [19], [20]

$$\mathcal{G}_T \dot{X}_T(t) = \mathcal{A}_T X_T(t) + \mathcal{B}_T U_T(t) + \mathcal{F}_T W_T(t), \tag{B.4a}$$

$$\mathcal{Y}_T(t) = \mathcal{C}_T X_T(t) + \mathcal{D}_T U_T(t), \tag{B.4b}$$

where the state vector X_T is defined as

$$X_T = [c_{11},\cdots, c_{1N}, c_{21},\cdots, c_{2N},\cdots, c_{M1},\cdots, c_{MN}]^T, \tag{B.5}$$

with the states $c_{mn}(t)$ denoting time-varying coefficients of the employed Chebyshev basis functions, $m \in \{1,\dots,M\}$, and $n \in \{1,\dots,N\}$. The tuning parameters M and N denote the number of basis functions used in the r and z directions, respectively, and are chosen based on the desired approximation accuracy and computational budget. The resulting model order is $O = MN$. The input vector U_T is composed of the boundary inputs at the lateral surface, top, and bottom sides of the cell, leading to

$$U_T(t) = [u_s(t),\ u_t(t),\ u_b(t)]^T, \tag{B.6}$$

with

$$u_s(t) = h_s(t)\, T_{s,\infty}(t), \quad u_t(t) = h_t(t)\, T_{t,\infty}(t), \quad u_b(t) = -h_b(t)\, T_{b,\infty}(t). \tag{B.7}$$

The cooling inputs in (B.6)–(B.7) capture the thermal interaction between the battery cell and the surrounding coolant and will be explicitly related to coolant model variables in subsection 2.2. The output vector $\mathcal{Y}_T$ consists of representative temperatures within and on the cell, including the quantities required for electrothermal and battery–coolant coupling. In particular, the volume-averaged cell temperature is defined as

$$T_v(t) = \frac{1}{V_b} \int_{\Omega_b} T(r,z,t)\, d\Omega_b, \tag{B.8}$$

where V_b and Ω_b are the cell volume and battery physical domain, respectively. The area-averaged battery–coolant interface temperatures at the lateral sur-

face, top, and bottom sides are defined as

$$T_{bc,s}(t) = \frac{1}{A_{ht,s}} \int_{A_{ht,s}} T(R_{\text{out}}, z, t)\, dA_{ht,s}, \tag{B.9a}$$

$$T_{bc,t}(t) = \frac{1}{A_{ht,t}} \int_{A_{ht,t}} T(r, L, t)\, dA_{ht,t}, \tag{B.9b}$$

$$T_{bc,b}(t) = \frac{1}{A_{ht,b}} \int_{A_{ht,b}} T(r, 0, t)\, dA_{ht,b}, \tag{B.9c}$$

respectively, where A_{ht} denotes the effective heat transfer area. The system matrices $\mathcal{G}_t$, $\mathcal{A}_t$, $\mathcal{B}_t$, $\mathcal{C}_t$, $\mathcal{D}_t$, and $\mathcal{F}_t$ are obtained via spectral-Galerkin projections of the governing heat-conduction partial differential equation (PDE) onto a finite set of basis functions, with the detailed derivations provided in [19], [20]. Finally, the disturbance input vector $W_t(t)$ in (B.4a) denotes the volumetric heat generation rate $q(t)$.

Battery electrical dynamics

The battery's electrical behaviour is modelled using a lumped-parameter equivalent circuit model (ECM) [22], which is widely adopted for control-oriented battery modelling due to its balance between accuracy and computational efficiency. The ECM captures the dynamic current-voltage relationship of the cell. A schematic of the ECM topology is shown in Fig. 2a. The model consists of an open-circuit voltage source V_{ocv}, which is a nonlinear function of the state of charge (SoC), a series internal resistance R_0, and a series connection of p paralleled resistance-capacitance branches. Each branch is characterised by a resistance R_i and capacitance C_i to account for the cell's polarisation dynamics. The ECM is expressed in continuous-time state-space form as

$$\dot{X}_e(t) = \mathcal{A}_e X_e(t) + \mathcal{B}_e U_e(t), \tag{B.10a}$$

$$\mathcal{Y}_e(t) = \mathcal{C}_e X_e(t) + \mathcal{D}_e U_e(t) + V_{\text{ocv}}\left(\text{SoC}(t)\right), \tag{B.10b}$$

where the state vector is defined as

$$X_e(t) = \left[\text{SoC}(t),\, V_1(t),\, \cdots,\, V_p(t)\right]^T, \tag{B.11}$$

with V_i being the voltage across the ith RC branch. The input $U_e = I_b$ is the applied current, with $U_e > 0$ implying discharging and $U_e < 0$ charging. The output $\mathcal{Y}_e$ is the terminal voltage, which depends nonlinearly on SoC

through V_{ocv}. The explicit expressions of $\mathcal{A}_e$, $\mathcal{B}_e$, $\mathcal{C}_e$ and $\mathcal{D}_e$ are provided in Appendix Appendix A.

Battery electrothermal coupling

The heat generation rate W_t that drives the thermal dynamics in (B.4) is calculated from the thermal model output, and the electrical model states and inputs. Assuming spatially uniform heat generation, the heat generated by the cell comprises reversible component (q_{rev}) and irreversible component (q_{irr}), expressed as [23]

$$q(t) = \frac{q_{\text{irr}}(t) + q_{\text{rev}}(t)}{V_b}, \tag{B.12a}$$

$$q_{\text{rev}}(t) = U_e(t) T_v(t) \frac{dV_{\text{ocv}}}{dT_v}, \tag{B.12b}$$

$$q_{\text{irr}}(t) = U_e^2(t) R_0 + \sum_{i=1}^{p} \frac{V_i^2(t)}{R_i}, \tag{B.12c}$$

where T_v is volume-averaged cell temperature defined in (B.8) and dV_{ocv}/dT_v is the entropic temperature coefficient of the open-circuit voltage. q_{rev} accounts for the entropy change associated with electrochemical reactions and q_{irr} represents ohmic and charge-transfer heat generation. The above equations give the disturbance input in the thermal model as $W_t(t) = q(t)$.

2.2 Coolant modelling

This subsection develops a dynamic model for the bulk coolant temperature in each cooling channel, capturing the effect of coolant flow rate on convective heat transfer and advective energy transport. The model is derived from a first-principles energy balance and is formulated in a manner suitable for control design. For notational simplicity, the derivation is presented for a generic cooling channel and then applied to the lateral surface, top, and bottom channels.

Heat transfer between the battery cell and the coolant is modelled as internal forced convection, where a water-glycol mixture is circulated by a pump through cooling plates or jackets attached to the cell. This configuration is commonly used in liquid-based BTMS because it provides efficient and controllable heat removal [24], [25]. The convective heat transfer coefficient is

characterised by [26]

$$h = \frac{\mathrm{Nu} \cdot k_c}{D_h}, \tag{B.13}$$

where k_c is the thermal conductivity of the coolant, D_h is the hydraulic diameter of the flow passage, and Nu is the Nusselt number.

To derive the governing equation for the coolant temperature, we apply the first law of thermodynamics to a control volume comprising the coolant flowing through a channel in thermal contact with the battery. The following modelling simplifications are adopted:

i) *Incompressible flow:* The coolant is assumed incompressible with constant density ρ_c and specific heat capacity $c_{p,c}$, which can be justified by the relatively small coolant temperature variations typically encountered in BTMS operation.

ii) *Sensible energy storage:* Only sensible thermal energy is considered, as the coolant operates below its boiling point and no phase change occurs.

iii) *Negligible mechanical energy effects:* Kinetic and potential energy changes, as well as pressure-drop effects, are negligible relative to thermal energy transport.

iv) *Mean temperature representation:* The coolant is characterised by a bulk mean temperature, assuming dominant streamwise heat transfer and sufficient mixing across the channel cross-section.

Applying these simplifications to the energy balance yields the following first-order dynamical model:

$$\frac{d\bar{T}_{cl}(t)}{dt} = \frac{h(t)A_{ht}(T_{bc}(t) - \bar{T}_{cl}(t))}{\rho_c V_c c_{p,c}} + \frac{\dot{\mathcal{V}}(t)(T_{cl,\mathrm{in}}(t) - \bar{T}_{cl}(t))}{V_c}, \tag{B.14}$$

where $\bar{T}_{cl}$ is the mean coolant temperature, T_{bc} is the battery–coolant interface temperature defined in (B.9), $T_{cl,\mathrm{in}}$ is the coolant inlet temperature, and V_c is the coolant volume in the channel. A detailed derivation of (B.14) is provided in Appendix Appendix B. The state variable is $\bar{T}_{cl}$, and the volumetric flow rate $\dot{\mathcal{V}}$ is selected as the manipulated input, as it can be directly actuated using pumps and valves in practical systems. Under this selection, system (B.14) becomes nonlinear due to the coupling between the state $\bar{T}_{cl}$ and the control input $\dot{\mathcal{V}}$. The treatment of this nonlinearity within the control framework is detailed in Section 3.

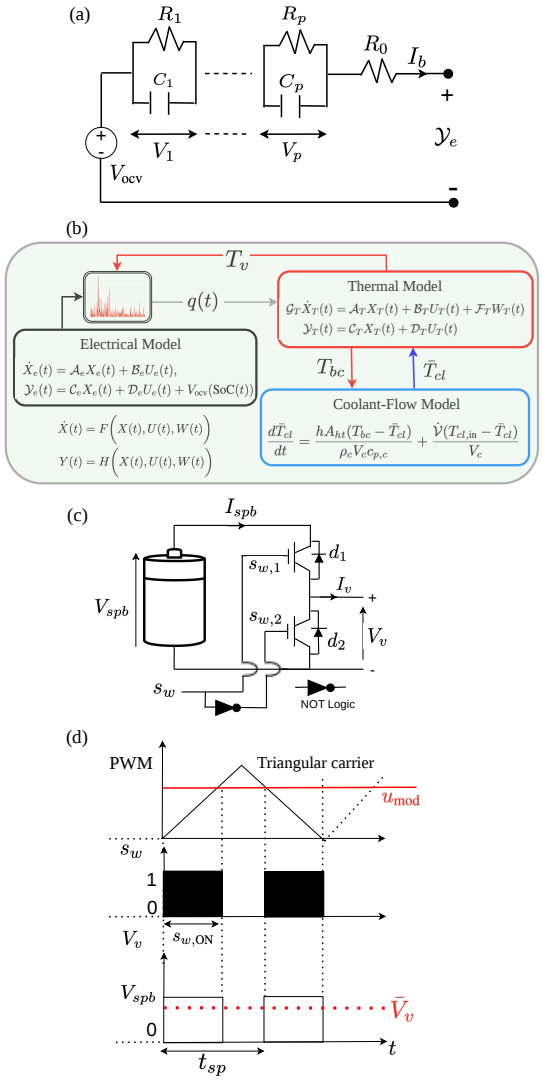

Figure 2: Battery–coolant model and valve operation. (a) ECM topology (b) Interconnections among battery–coolant models. (c) A simplified schematic of a half-bridge DC-DC power converter. It consists of the switching function s_w, transistor switches $s_{w,1}$ and $s_{w,2}$, each with diodes d_1 and d_2 in antiparallel. (c) PWM operation showing triangular wave carrier and modulating signal. The comparison of these two signals produce the average valve voltages, $\bar{V}_v$ that determine the fraction of coolant flow to each channel.

2.3 Battery–coolant coupling

The battery is in thermal contact with the coolant circulating through the ITSC system. Since heat exchange occurs at the cell boundaries, the battery–coolant interaction is introduced through the thermal subsystem (B.4), via the boundary inputs (B.7) that govern convective heat transfer at each cooled surface. Specifically, replacing the free-stream temperatures T_∞ in (B.7) with the corresponding mean coolant temperatures $\bar{T}_{cl}$ yields the modified thermal boundary inputs

$$u_s(t) = h_s(t)\bar{T}_{cl,s}(t), \quad u_t(t) = h_t(t)\bar{T}_{cl,t}(t), \quad u_b(t) = -h_b(t)\bar{T}_{cl,b}(t). \tag{B.15a}$$

Equations (B.15) establish a bidirectional coupling between the battery thermal model and coolant model. The coolant model (B.14) provides the channel mean temperatures as dynamic boundary inputs to the thermal model, while the battery boundary temperatures T_{bc} computed by the thermal model drive the coolant dynamics through the convective heat transfer term in (B.14). The battery–coolant subsystem interconnections is illustrated in Fig. 2b.

Assembling the electrical (B.10), thermal (B.4), and coolant (B.14) subsystems yields the complete nonlinear state-space model

$$\dot{X}(t) = F\left(X(t), U(t), W(t)\right), \tag{B.16a}$$

$$Y(t) = H\left(X(t), U(t), W(t)\right), \tag{B.16b}$$

where $F(\cdot)$ and $H(\cdot)$ denote the nonlinear state and output mappings capturing the coupled battery–coolant dynamics. The state vector is defined as

$$X(t) = \left[X_e(t), \quad X_T(t), \quad X_{cl}(t)\right]^T, \tag{B.17}$$

where $X_{cl}(t)$ denotes the coolant state vector with

$$X_{cl}(t) = \left[\bar{T}_{cl,s}(t), \quad \bar{T}_{cl,t}(t), \quad \bar{T}_{cl,b}(t)\right]^T. \tag{B.18}$$

The control input vector is given by

$$U(t) = \left[\dot{\mathcal{V}}_s(t), \quad \dot{\mathcal{V}}_t(t), \quad \dot{\mathcal{V}}_b(t)\right]^T, \tag{B.19}$$

and the exogenous disturbance input vector is defined as

$$W(t) = \left[U_e(t), \quad T_{cl,\mathrm{in},s}(t), \quad T_{cl,\mathrm{in},t}(t), \quad T_{cl,\mathrm{in},b}(t)\right]^T. \tag{B.20}$$

The output vector is given by

$$Y(t) = \left[\mathcal{Y}_e(t), \quad \mathcal{Y}_T(t), \quad X_{cl}(t)\right]^T, \tag{B.21}$$

with X_{cl} assumed measurable. The mapping $F(\cdot)$ is given by

$$\dot{X}_e(t) = \mathcal{A}_e X_e(t) + \mathcal{B}_e U_e(t), \tag{B.22a}$$
$$\dot{X}_T(t) = \mathcal{G}_T^{-1}\left(\mathcal{A}_T X_T(t) + \mathcal{B}_T U_T\big(X_{cl}(t)\big) + \mathcal{F}_T W_T(t)\right), \tag{B.22b}$$
$$\dot{X}_{cl}(t) = f\left(X_{cl}(t), X_T(t), U(t), T_{cl,\mathrm{in}}(t)\right), \tag{B.22c}$$

where $f(\cdot)$ denotes the nonlinear coolant dynamics (B.14). The volumetric flow rates $\dot{\mathcal{V}}$ in (B.19) constitute the effective control inputs to the coupled battery–coolant system. In practice, however, these flows are not directly actuated and are instead realised through a valve–converter actuation mechanism, described in the following subsection.

2.4 Valve actuation modelling

Remark 5. *Many existing BTMS control studies treat the coolant flow rate as a directly controlled input [27], abstracting away the underlying actuator-level dynamics through which flow distribution is realised. As established in [28], this oversimplifies the coolant circuit, since the flow is an outcome of controlling physical actuators including the valves.*

In response to *Remark 5*, we explicitly model the valve actuation mechanism governing coolant distribution across the channels. For notational simplicity, the derivation is presented for a generic valve–converter pair; the resulting model applies independently to each of the three channels.

The channel solenoid valves are assumed to be actuated using half-bridge DC–DC power converters, whose topology is illustrated in Fig. 2c. The converters are powered by a supply battery providing voltage V_{spb} and current I_{spb}. The coil voltage and current associated with each valve are denoted by V_v and I_v, respectively. Each converter consists of two ideal power transistors with switching states $s_{w,1}, s_{w,2} \in \{0, 1\}$. To prevent short-circuiting across the DC bus, the transistors are constrained by the orthogonality condition

$$\int_{t-t_{sp}}^{t} s_{w,1}(\tau) s_{w,2}(\tau)\, d\tau = 0, \tag{B.23}$$

where t_{sp} is the switching period.

The switching signal $s_w(t) \in \{0, 1\}$ is generated via a pulse-width modulation (PWM) process as illustrated in Fig. 2d. Specifically, a modulating signal $u_{\mathrm{mod}}(t) \in [0, 1]$, computed by the controller, is compared with a high-frequency waveform to produce the binary switching signal. The resulting valve voltage is given by

$$V_v(t) = \begin{cases} V_{spb}(t), & \text{if } s_w = 1, \\ 0, & \text{if } s_w = 0. \end{cases} \tag{B.24}$$

To avoid explicit treatment of the high-frequency binary switching dynamics, we apply the switching-period averaging framework [17], [29], [30], which replaces the switching signals with their time-averaged equivalents:

$$u_{\mathrm{mod}}(t) := \frac{1}{t_{sp}} \int_{t-t_{sp}}^{t} s_w(\tau)\, d\tau = \frac{s_{w,\mathrm{ON}}}{t_{sp}}(t), \tag{B.25}$$

where $s_{w,\mathrm{ON}}$ is the total ON duration of s_w within one PWM period. The resulting continuous-valued signals u_{mod} represent the valve duty-cycles, corresponding to the fraction of time each valve remains open per switching period. The effective volumetric flow rates entering the lateral surface, top, and bottom channels are thus

$$\dot{\mathcal{V}}_s = \dot{\mathcal{V}}_{\max} u_{\mathrm{mod},s}, \quad \dot{\mathcal{V}}_t = \dot{\mathcal{V}}_{\max} u_{\mathrm{mod},t}, \quad \dot{\mathcal{V}}_b = \dot{\mathcal{V}}_{\max} u_{\mathrm{mod},b}, \tag{B.26}$$

subject to the physical duty-cycle bounds

$$0 \le u_{\mathrm{mod}}(t) \le 1, \tag{B.27}$$

and the conservation constraint

$$u_{\mathrm{mod},s}(t) + u_{\mathrm{mod},t}(t) + u_{\mathrm{mod},b}(t) \le 1, \tag{B.28}$$

which allows the controller to utilise less than the full pump capacity when thermal loads are low. Equations (B.26)–(B.28) complete the actuator model, closing the link between the controller outputs u_{mod} and the manipulated inputs $\dot{\mathcal{V}}$ of model (B.16).

3 MPC formulation

Building on the battery–coolant and valve-actuation models developed in Section 2, this section presents the control formulation used to optimally allocate coolant across the cooling channels. The control architecture, cost function, constraints, and controller implementation are introduced in turn.

3.1 Control architecture and problem statement

The integration of lateral surface and tab cooling is formulated as the following optimal control problem: *The optimal coolant-allocation control problem (OCAP) is to determine, over the prediction horizon, the allocation of coolant flow among the lateral surface, top tab and bottom tab cooling channels such that the battery temperature is regulated toward a desired reference while internal thermal gradients are minimised, without violating battery safety, health-related, and actuator constraints.*

The regulation objective in the OCAP is motivated by the thermal degradation mechanisms that define an optimal operating window for Li-ion cells [28], [31], [32], [33]. Elevated temperatures accelerate side reactions such as solid electrolyte interphase (SEI) growth and electrolyte decomposition, leading to capacity fade and increased internal resistance. Conversely, lower temperatures increase the risk of Li plating during charging due to reduced ionic conductivity. These effects jointly define an optimal temperature window, motivating regulation toward a reference within this range. In addition, minimising internal thermal gradients mitigates localised ageing and thermally induced mechanical stress within the cell [8], [9].

The OCAP exhibits three key features that motivate the use of MPC. First, the battery-thermal and coolant-flow dynamics are coupled, making the coolant-allocation problem inherently multivariable. Second, the system is subject to hard constraints on battery operation and actuator limits, which must be satisfied at all times. Third, predictive optimisation enables anticipation of future thermal loads and coolant-temperature variations over the horizon.

MPC-based BTMS have previously been demonstrated to improve temperature regulation, energy efficiency, and battery lifespan [27], [28]. In EV battery thermal management, RTI-MPC [34] has been adopted as a computationally efficient alternative to nonlinear (N)MPC [35], [36]. These studies regulate

a single coolant flow rate, single coolant temperature, and/or auxiliary cooling power. However, they do not address the dynamic allocation of a fixed coolant budget among complementary tab and surface cooling channels with actuator-level routing constraints. This work addresses this gap by solving, for the first time, the OCAP within an RTI-MPC framework.

The proposed MPC-based ITSC feedback control strategy is illustrated in Fig. 3 for a single cell. The control architecture includes state estimation blocks for reconstructing unmeasurable battery states, particularly the SoC and internal temperatures. Based on these estimates, the controller computes optimal valve duty-cycles to achieve the desired thermal objectives.

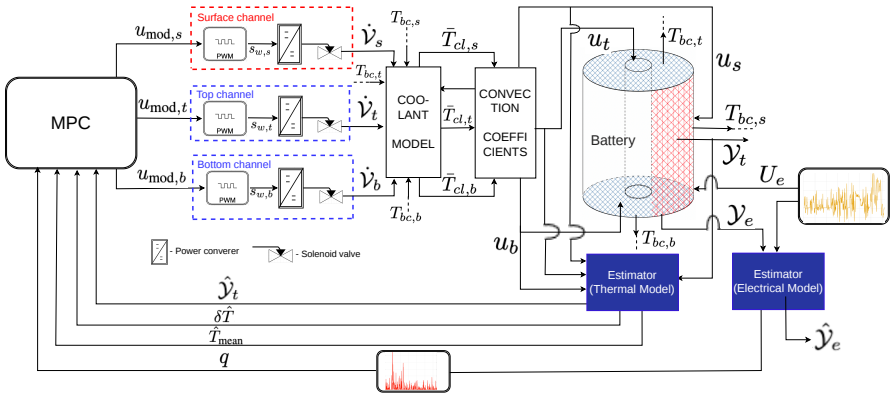

Figure 3: Feedback control architecture for implementing the OCAP of a single battery cell under the proposed ITSC system. The setup comprises an MPC controller, temperature and electrical sensors, state estimation blocks, and three cooling channels for the cell's tabs and lateral surface, each governed by a valve-switching mechanism. The tabs fully cover the top and bottom areas of the cell, with their corresponding cooling channels shaded in blue. The surface cooling channel, shaded in red, covers the lateral side of the cell.

3.2 Cost function

At each sampling instant k, the MPC solves the OCAP by determining the sequence of valve duty-cycles $u_{\text{mod}}(k)$ that minimise the finite-horizon objec-

tive

$$J(k) := \sum_{i=0}^{N_p-1} \Big[\lambda_1 \Big\| T^r(k+i) - \hat{\mathcal{Y}}_T(k+i) \Big\|_{Q_e}^2 + \lambda_2 \Big\| \delta\hat{T}(k+i) \Big\|_{Q_\delta}^2 + \lambda_3 \Big\| \Delta u_{\text{mod}}(k+i) \Big\|_R^2 \Big], \tag{B.29}$$

where N_p is the prediction horizon, $\|x\|_D^2 := x^\top D x$ denotes the weighted quadratic norm of vector x, and $\{\lambda_i \geq 0\,; i = 1, 2;\ \lambda_3 > 0\}$ are trade-off weights that determine the relative significance of each objective. The weighting matrices satisfy $Q_e \succeq 0$, $Q_\delta \succeq 0$, and $R \succ 0$. The three terms in (B.29) are described below.

Temperature regulation

The first term penalises the deviation of the predicted temperature output $\hat{\mathcal{Y}}_T$ from the desired reference T^r, which may be specified as a constant setpoint or a time-varying trajectory. This term promotes battery operation within a thermally safe and performance-optimal temperature range, necessary for mitigating temperature-induced degradation, as discussed in Section 3.1. In practice, since the duty-cycle bound (B.27) restricts the controller to cooling only, zero coolant flow is applied when the battery temperature $\hat{\mathcal{Y}}_T \leq T^r$.

Thermal gradient suppression

The second term penalises internal temperature non-uniformities through the gradient-based measure $\delta\hat{T}$, defined as the root-mean-square (RMS) spatial temperature gradient of the predicted temperature field:

$$\delta\hat{T}(k+i) = \left(\frac{1}{|\Omega|} \int_\Omega \Big\| \nabla\hat{T}(r, z, k+i) \Big\|_2^2 \, d\Omega \right)^{1/2}, \tag{B.30}$$

where $|\Omega|$ denotes the measure of the spatial domain Ω, and $\nabla\hat{T} = \left[\frac{\partial\hat{T}}{\partial r}, \quad \frac{\partial\hat{T}}{\partial z} \right]^\top$ is the spatial temperature gradient with respect to coordinates (r, z). Minimising (B.30) promotes spatially uniform temperature distributions, thereby reducing localised cell ageing and thermomechanical stress.

Control variation penalty

The third term penalises changes in the duty-cycle input according to

$$\Delta u_{\mathrm{mod}}(k+i) = u_{\mathrm{mod}}(k+i) - u_{\mathrm{mod}}(k+i-1), \tag{B.31}$$

thereby discouraging aggressive changes in coolant-routing commands. This regularises the control action and reduces excessive chattering of the valve–converter system.

3.3 Constraints

The objective in (B.29) is minimised subject to the following constraints $\forall i \in \{0, 1, \cdots, N_p - 1\}$;

$$\hat{X}_e(k+1+i) = \mathcal{A}_{e,d}\hat{X}_e(k+i) + \mathcal{B}_{e,d}U_e(k+i), \tag{B.32a}$$

$$\mathcal{G}_T\hat{X}_T(k+1+i) = \mathcal{A}_{T,d}\hat{X}_T(k+i) + \mathcal{B}_{T,d}U_T(k+i) + \mathcal{F}_{T,d}W_T(k+i), \tag{B.32b}$$

$$\hat{\mathcal{Y}}_T(k+i) = \mathcal{C}_T\hat{X}_T(k+i) + \mathcal{D}_T U_T(k+i), \tag{B.32c}$$

$$X_{cl}(k+1+i) = f_d\Big(X_{cl}(k+i), \hat{\mathcal{Y}}_T(k+i), U(k+i), T_{cl,\mathrm{in}}(k+i)\Big), \tag{B.32d}$$

$$0 \leq \hat{\mathrm{SoC}}(k+i) \leq 1, \tag{B.32e}$$

$$\hat{T}_c(k+i) \leq \hat{T}_{c,\max}, \tag{B.32f}$$

$$U_{e,\min} \leq U_e(k+i) \leq U_{e,\max}, \tag{B.32g}$$

$$0 \leq u_{\mathrm{mod}}(k+i) \leq 1, \tag{B.32h}$$

$$u_{\mathrm{mod},s}(k+i) + u_{\mathrm{mod},t}(k+i) + u_{\mathrm{mod},b}(k+i) \leq 1. \tag{B.32i}$$

Equations (B.32a)–(B.32d) define the discrete-time prediction model used by the MPC, obtained by zero-order hold discretisation of the continuous-time system in (B.22). The matrices with subscript d denote the discrete-time battery subsystem, while $f_d(\cdot)$ denotes the discrete-time coolant subsystem, which retains the nonlinear structure of (B.14). Here, $(\hat{\cdot})$ denotes variables that are either estimated by the observer or propagated forward through the prediction model.

The remaining constraints enforce admissible bounds on the system states and inputs. Constraint (B.32e) restricts the battery SoC to the physically ad-

missible interval $[0, 1]$, while (B.32f) and (B.32g) impose safety limits on the estimated core temperature $\hat{T}_c$ and battery current U_e, respectively. Restricting $\hat{T}_c$, which often corresponds to the hottest region of the cell, mitigates the risk of accelerated degradation and thermal runaway. Finally, (B.32h) and (B.32i) enforce the physical limits of the valve–converter actuation and ensure that the total commanded coolant distribution does not exceed the available pump capacity.

3.4 Controller Implementation

At each sampling instant k, the constrained optimisation problem (B.29)–(B.32) is solved in a receding-horizon fashion. The optimal duty-cycle sequence $\{u^*_{\mathrm{mod}}(k+i)\}_{i=0}^{N_p-1}$ is computed, only the first element $u^*_{\mathrm{mod}}(k)$ is applied to the valve–converter actuation system, and the optimisation is repeated at the next sampling instant using updated state information.

In the NMPC formulation, the prediction model (B.32a)–(B.32d) is retained in its original form, including the nonlinear coolant dynamics. The resulting optimisation problem is therefore a nonlinear programme, which is solved directly at each sampling instant k. Solving OCAP (B.29)–(B.32) to convergence at every sampling instant via NMPC can be computationally demanding, which may hinder real-time implementation. To reduce this burden, the proposed controller is implemented within the RTI-MPC scheme [34], [37]. The NMPC is used as a benchmark for assessing the closed-loop thermal performance attainable under the proposed ITSC system.

RTI-MPC for the coolant-allocation problem

At each sampling instant k, the RTI-MPC implementation consists of a warm-start step, a preparation phase, and a feedback phase.

Warm-start step Let

$$u^{*,\mathrm{prev}}_{\mathrm{mod}}(i), \qquad i = 0, \ldots, N_p - 1, \tag{B.33}$$

denote the optimal duty-cycle sequence computed at the previous sampling instant $k-1$. At k, this sequence is shifted forward by one step to initialise the nominal duty-cycle input sequence $\bar{u}_{\mathrm{mod}}$ in the current prediction horizon:

$$\bar{u}_{\mathrm{mod}}(k+i) = u^{*,\mathrm{prev}}_{\mathrm{mod}}(i+1), \qquad i = 0, \ldots, N_p - 2. \tag{B.34}$$

The terminal input is held constant as

$$\bar{u}_{\mathrm{mod}}(k+N_p-1) = u_{\mathrm{mod}}^{*,\mathrm{prev}}(N_p-1). \tag{B.35}$$

The nominal volumetric flow trajectory $\bar{U}(k+i)$ is then obtained from (B.19) through (B.26). The corresponding nominal mean coolant temperature trajectory is obtained by forward propagation of nonlinear model (B.32d) as

$$\bar{X}_{cl}(k+1+i) = f_d\big(\bar{X}_{cl}(k+i), \bar{\mathcal{Y}}_T(k+i), \bar{U}(k+i), T_{cl,\mathrm{in}}(k+i)\big). \tag{B.36}$$

for $i = 0, \ldots, N_p-1$, where $\bar{\mathcal{Y}}_T$ denotes the nominal battery temperature output trajectory.

Preparation phase The coolant dynamics (B.32d) are linearised about the nominal coolant trajectory (B.36). Define the perturbation variables

$$\Delta X_{cl}(k+i) = X_{cl}(k+i) - \bar{X}_{cl}(k+i), \tag{B.37a}$$
$$\Delta U(k+i) = U(k+i) - \bar{U}(k+i), \tag{B.37b}$$
$$\Delta \hat{\mathcal{Y}}_T(k+i) = \hat{\mathcal{Y}}_T(k+i) - \bar{\mathcal{Y}}_T(k+i). \tag{B.37c}$$

A first-order Taylor approximation of (B.32d) gives

$$\begin{aligned}\Delta X_{cl}(k+1+i) = {} & A_{cl,i}\Delta X_{cl}(k+i) + B_{cl,i}\Delta U(k+i) \\ & + E_{cl,i}\Delta \hat{\mathcal{Y}}_T(k+i),\end{aligned} \tag{B.38}$$

where the Jacobians are evaluated along the nominal trajectory as

$$A_{cl,i} = \left.\frac{\partial f_d}{\partial X_{cl}}\right|_{\bar{\xi}_i}, \quad B_{cl,i} = \left.\frac{\partial f_d}{\partial U}\right|_{\bar{\xi}_i}, \quad E_{cl,i} = \left.\frac{\partial f_d}{\partial \hat{\mathcal{Y}}_T}\right|_{\bar{\xi}_i}, \tag{B.39a}$$

with $\bar{\xi}_i = \big(\bar{X}_{cl}(k+i), \bar{\mathcal{Y}}_T(k+i), \bar{U}(k+i), T_{cl,\mathrm{in}}(k+i)\big)$. Replacing (B.32d) with (B.38) yields a local linear approximation of the coolant model, which transforms the original OCAP into the following local quadratic programme

(QP) at each k:

$$\begin{aligned}
\min_{\Delta u_{\mathrm{mod},k}} \quad & J(k) \\
\text{s.t.} \quad & \text{(B.32a)–(B.32c), (B.38), (B.32e)–(B.32g)}, \\
& 0 \le \bar{u}_{\mathrm{mod}}(k+i) + \Delta u_{\mathrm{mod}}(k+i) \le 1, \\
& \bar{u}_{\mathrm{mod},s}(k+i) + \Delta u_{\mathrm{mod},s}(k+i) + \\
& \bar{u}_{\mathrm{mod},t}(k+i) + \Delta u_{\mathrm{mod},t}(k+i) + \\
& \bar{u}_{\mathrm{mod},b}(k+i) + \Delta u_{\mathrm{mod},b}(k+i) \le 1, \\
& i = 0, \ldots, N_p - 1.
\end{aligned} \tag{B.40}$$

The QP decision variable is the input correction Δu_{mod}, so that the updated duty-cycle sequence is

$$u_{\mathrm{mod}}(k+i) = \bar{u}_{\mathrm{mod}}(k+i) + \Delta u_{\mathrm{mod}}(k+i). \tag{B.41}$$

Feedback phase Once the current states $\hat{X}_e(k)$, $\hat{X}_T(k)$, and $X_{cl}(k)$ become available, QP (B.40) is solved once. The first duty-cycle command is then updated as

$$u^*_{\mathrm{mod}}(k) = \bar{u}_{\mathrm{mod}}(k) + \Delta u^*_{\mathrm{mod}}(k), \tag{B.42}$$

and applied through (B.26). The remaining elements of the updated input sequence are stored and used to warm-start the optimisation at the next sampling instant $k+1$.

This single-QP-per-step structure is the key computational advantage of RTI-MPC over the NMPC. Rather than repeatedly solving the nonlinear programme to convergence, RTI-MPC performs a single linearisation of the coolant dynamics per step and solves one structured QP, retaining the receding-horizon and constraint-handling structure of the OCAP. The nonlinear coolant dynamics are preserved in the nominal trajectory propagation (B.36), while the QP solves a local correction about this trajectory. The accuracy of this approximation improves as the warm-start converges toward the true optimum.

4 Results and discussion

4.1 Specification and simulation setup

The proposed ITSC system is simulated on an A123 2.3 Ah LFP cylindrical cell [38]. The thermo-physical parameters used in the thermal model are adopted from [19], namely $\rho = 2118$ kgm^{-3}, $c_{p,b} = 795$ Jkg^{-1}K^{-1}, $k_r = 0.67$ Wm^{-1}K^{-1}, and $k_z = 66.6$ Wm^{-1}K^{-1}. For the ECM, $R_0 = 0.0106$ Ω, and a single RC branch is employed with $R_1 = 0.0169$ Ω and $C_1 = 2249$ F [39]. Fig 4(a) shows the nonlinear V_{ocv}–SoC relationship, obtained from experimental data and implemented as a lookup table for the considered LFP cell.

In this study, only the irreversible heat generation component q_{irr} is considered. This assumption is justified by the fact that for LFP chemistry under moderate-to-high current operation, irreversible losses dominate the total heat generation, while the reversible contribution is comparatively small [23]. Accordingly, the heat generation term in (B.12a) is approximated as $q \approx q_{\text{irr}}/V_b$. The applied battery current I_b depends on the vehicle operating conditions. We select the Worldwide Harmonised Light Vehicle Test Procedure (WLTP) [40], which captures a wide range of driving conditions on urban, suburban, and highway roads, to generate a realistic, dynamically varying current load. To obtain a thermally demanding scenario within the simulation window, the WLTP current trajectory is scaled by a factor chosen to maximise the mean C-rate while avoiding cell undervoltage and excessive depletion. The modified WLTP (mWLTP) has a mean C-rate of 2. The SoC evolution and applied current trajectory are shown in Fig. 4(b)–(c).

For the coolant, a typical 50:50 ethylene-glycol-water mixture at 20 °C is adopted, with properties $\rho_c = 1069$ kgm^{-3}, $c_{p,c} = 3323$ Jkg^{-1}K^{-1}, and $k_c = 0.3892$ Wm^{-1}K^{-1} [41]. The convection coefficients h used in equations (B.14) and (B.15) are derived in Appendix Appendix C. Since a high heat transfer area-to-volume ratio is desirable, the channel thicknesses are chosen to be very thin: $\epsilon_s = 2$ mm for the lateral surface channel and, for simplicity, $\epsilon_t = \epsilon_b = 2$ mm for the top and bottom channels. Substituting these dimensions into (B.57) and assuming laminar fully developed flow with Nu = 4.86 [26] yields the convection coefficients $h_s \approx h_t \approx h_b \approx 480$ Wm^{-2}K^{-1}. These values are consistent with those typically reported for forced convection liquid cooling in BTMS [5], [42]. The total coolant flow rate for the cell, $\dot{\mathcal{V}}_{\max}$, is sized from

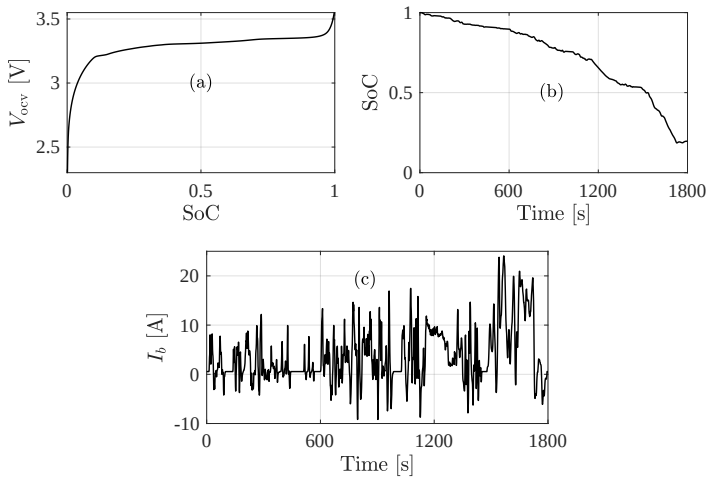

Figure 4: Electrical model variables. (a) V_{ocv}–SoC relationship of the considered LFP cell. (b) SoC trajectory under the applied mWLTP current input. (c) Applied mWLTP current input.

an energy balance based on the cell's expected thermal load. To dissipate a representative peak cell heat generation $q \approx 100$ W with an allowable coolant temperature rise of $\Delta T_{cl} \approx 5$ K, the required flow rate is

$$\dot{\mathcal{V}}_{\max} = \frac{q}{\rho_c c_{p,c} \Delta T_{cl}} \approx 5.63 \times 10^{-6}\ \mathrm{m^3 s^{-1}}, \tag{B.43}$$

which corresponds to a nominal flow rate of approximately 0.34 $\mathrm{L min^{-1}}$.

For the CSG thermal model (B.4), the number of basis functions in the radial and axial directions is chosen equal, i.e., $M = N$, yielding a model order of $O = N^2$. The MPC prediction model uses $N = 2$ (four states) for both the RTI-MPC and NMPC implementations, whereas the plant model uses $N = 10$ (100 states). The initial cell temperature distribution $T(r, z, 0)$, initial mean coolant temperatures $X_{cl}(0)$, and inlet coolant temperatures $T_{cl,\text{in}}$ are all set to 30 °C. All simulations are conducted on a MacBook Pro with a 2.6GHz 6-core Intel Core i7 processor and 16GB RAM. The optimisation problem is modelled with Yalmip [43] in MATLAB R2025b and solved using Gurobi for the RTI-MPC and fmincon for the NMPC. The MPC sampling time is set to 1 s, with a prediction horizon of $N_p = 20$ steps. In the objective function, the temperature

reference is set to $T^r = 35$ °C, corresponding approximately to the midpoint of the safe and degradation-conscious battery operating temperature range highlighted in [28], [31], [32], [33]. The trade-off weights are $\lambda_1 = \lambda_2 = 1$ and $\lambda_3 = 0.5$, and the core temperature constraint in (B.32f) is set to $T_{c,\max} = 50$ °C.

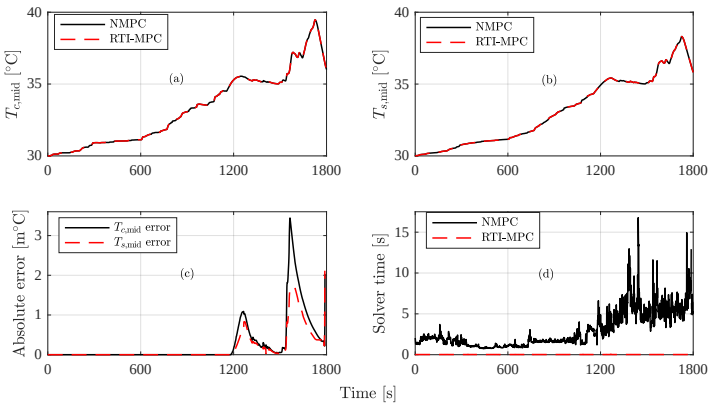

Figure 5: Comparison of the NMPC and RTI-MPC implementations. (a) and (b) show the battery core mid-point and surface mid-point temperature evolutions. (c) Absolute error in core and surface mid-point temperatures between NMPC and RTI-MPC. (d) Online solver time per sampling instant.

4.2 RTI-MPC vs. NMPC

The closed-loop temperature responses of the NMPC and RTI-MPC implementations are compared in Fig. 5. As shown in Fig. 5(a)–(b), the RTI-MPC reproduces the core and surface temperature trajectories of the NMPC almost exactly over the entire mWLTP. The corresponding absolute temperature errors in Fig. 5(c) remain below 0.0035 °C for the core and 0.002 °C for the surface throughout the entire cycle, with only minor deviations appearing during the most thermally demanding part of the drive cycle. These errors are negligible relative to the temperature scales of practical interest, confirming that the successive linearisation employed in RTI-MPC introduces no meaningful loss of control performance under this operating profile.

The computational advantage of the RTI-MPC is evident from Fig. 5(d). While the NMPC exhibits substantially larger online solve times and occasional spikes, the RTI-MPC maintains consistently low solver times throughout the cycle. For the present case, the RTI-MPC achieved a maximum and mean online solve times of 0.41 s and 0.0193 s, respectively, on a laptop-class processor in MATLAB, which is well below the 1 s sampling interval and indicates strong potential for real-time embedded controller implementation.

4.3 Benchmarks for the ITSC system

To evaluate the proposed ITSC-based control strategy, its performance is compared against five baseline cooling configurations. Conventional cooling configurations for cylindrical cells apply cooling along the lateral surface only, through a single tab only, through a combination of lateral surface and single tab, or through both the top and bottom tabs [44], [45]. These four configurations are denoted here as surface cooling (SC), bottom tab cooling (bTC), bottom tab and surface cooling (bTSC), and bottom and top tabs cooling (btTC). A fifth benchmark, denoted as equal-split (ES), corresponds to a fixed uniform distribution of coolant across all available channels. This benchmark represents a predefined coolant manifold without active feedback control or optimisation, reflecting rule-based designs commonly used for their simplicity and low implementation cost.

Performance is assessed using metrics that quantify the extent to which the cell temperature overshoots the desired reference, temperature rise, and thermal non-uniformity. The reference overshoot is evaluated using the local tracking error

$$e(r, z, t) = \max\left(0,\, T(r, z, t) - T^{r}\right), \tag{B.44}$$

with the corresponding maximum and mean deviation errors defined as

$$e_{\max}(t) = \max_{(r,z)\in\Omega} e(r, z, t), \tag{B.45a}$$

$$e_{\text{mean}}(t) = \frac{1}{|\Omega|}\int_{\Omega} e(r, z, t)\, d\Omega. \tag{B.45b}$$

Since cooling action is not activated when $T(r, z, t) \leq T^{r}$, temperatures below the reference do not necessarily indicate poor cooling performance; the relevant tracking error is the local temperature overshoot above T^{r}, as reflected

in the one-sided definition in (B.44). The temperature rise metrics are the maximum and mean cell temperatures,

$$T_{\max}(t) = \max_{(r,z)\,\in\,\Omega} T(r,z,t), \tag{B.46a}$$

$$T_{\text{mean}}(t) = \frac{1}{|\Omega|}\int_{\Omega} T(r,z,t)\,d\Omega, \tag{B.46b}$$

and the thermal-uniformity metrics are the maximum and RMS magnitudes of the spatial temperature gradient,

$$\delta T_{\max}(t) = \max_{(r,z)\,\in\,\Omega} \left\|\nabla T(r,z,t)\right\|_2, \tag{B.47a}$$

$$\delta T_{\text{rms}}(t) = \left(\frac{1}{|\Omega|}\int_{\Omega} \left\|\nabla T(r,z,t)\right\|_2^2\, d\Omega\right)^{1/2}. \tag{B.47b}$$

With the exception of ES, the proposed ITSC system and benchmark configurations are implemented using the RTI-MPC formulation described in Section 3.4. The same prediction models (B.32a)–(B.32c) and (B.38), prediction horizon, objective function, constraints, and observer are employed across all configurations; they differ only in the admissible set of control inputs. For each benchmark, only the cooling channels corresponding to a given configuration are assumed physically available, while the total coolant flow rate remains identical across all configurations. Accordingly, for SC and bTC, only the lateral surface duty-cycle input $u_{\text{mod},s}$, and the bottom tab input $u_{\text{mod},b}$ are available, respectively. For bTSC, the admissible inputs are $u_{\text{mod},b}$ and $u_{\text{mod},s}$, while for btTC, $u_{\text{mod},t}$ and $u_{\text{mod},b}$ are available. For both the proposed ITSC strategy and ES benchmark, all three inputs $u_{\text{mod},s}, u_{\text{mod},t}$ and $u_{\text{mod},b}$ are available. In the ES configuration, the inputs are fixed according to $u_{\text{mod},s} = u_{\text{mod},t} = u_{\text{mod},b} = \frac{1}{3}$, resulting in an equal allocation of the total volumetric coolant flow rate among the three channels.

4.4 ITSC against conventional cooling configurations

The thermal performance of the proposed ITSC system relative to the benchmark cooling configurations are shown in Fig. 6. From Figs. 6(a)–(b), the ES configuration yields the lowest $T_{\max}$ and T_{mean} values. This behaviour is expected because ES continuously distributes the full available coolant flow across all three channels, thereby maximising the effective heat transfer area

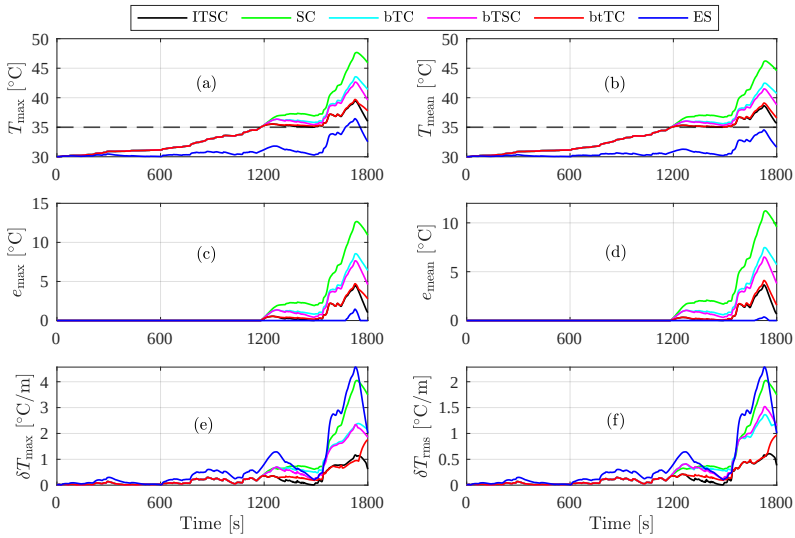

Figure 6: Comparison of the proposed ITSC system and benchmark cooling configurations under the mWLTP. (a) and (b) present the maximum and mean temperature responses; (c) and (d), the maximum and mean tracking errors; and (e) and (f), the maximum and RMS thermal gradients.

and total heat removal. However, this does not necessarily imply superior feedback regulation. ES operates in open-loop and cannot track a desired temperature reference. The deviation errors in Figs. 6(c)–(d) provide a more direct assessment of temperature regulation relative to T^r. Although the temperature trajectory under ES remains below the reference (zero error) for most of the cycle, it does not track T^r in any meaningful sense. It operates without adaptation to the instantaneous thermal demand and continuously uses the full pump capacity. Therefore, its lower temperature rise and tracking error are achieved at the expense of unnecessary coolant usage under moderate thermal loading.

Among the feedback-controlled configurations, the proposed ITSC system achieves the smallest reference overshoot. This indicates that the ITSC controller is able to exploit all available cooling areas to keep the cell temperature closer to the desired operating range. In contrast, SC exhibits the largest deviation error, indicating that lateral surface cooling alone is insufficient to

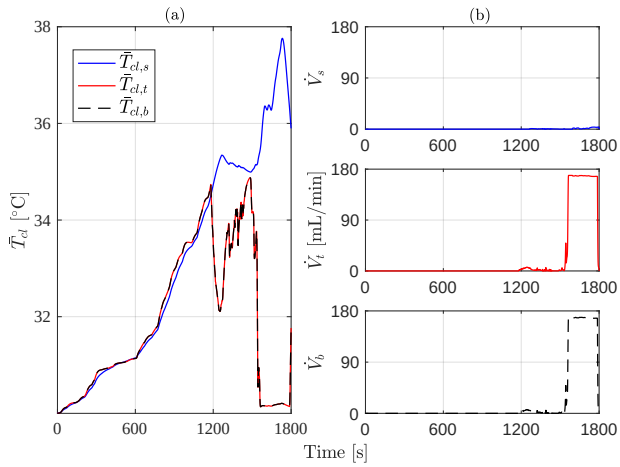

Figure 7: ITSC system channel responses. (a) Mean coolant temperature evolution in the surface, top, and bottom channels. (b) Volumetric flow rate allocation.

regulate the thermally dominant internal region of the cell. Although bTC uses a smaller heat-transfer area than SC, it interestingly achieves lower temperature rise and smaller deviation above the reference. This is attributed to the homogenising effect of tab cooling: by reducing internal temperature gradients, the optimiser drives spatial temperatures closer together, which indirectly constrains the peak temperature. The addition of both tab and surface cooling in bTSC leads to a further temperature reduction, while btTC improves the response further still. These trends highlight the importance of tab cooling for extracting heat from the cell more effectively than surface cooling alone.

Figs. 6(e)–(f) show that the proposed ITSC system yields the lowest thermal gradients among all configurations. The improvement over SC, bTC, and bTSC is substantial, while btTC is the closest conventional competitor. The flow allocation in Fig. 7(b) provides useful insight into this behaviour: the ITSC controller predominantly routes coolant through the top and bottom channels during periods of high thermal loading. Under the chosen cost weights ($\lambda_1 = \lambda_2$), the optimiser assigns comparable importance to temperature tracking and gradient minimisation. Since tab cooling is intrinsically

more effective at homogenising the temperature field, the controller naturally favours the tab channels under this tuning. The btTC configuration exclusively utilises the top and bottom axial paths, thus similar gradient response arises. However, ITSC retains the surface channel as an additional controllable degree of freedom, allowing the optimiser to balance gradient suppression and temperature tracking more flexibly.

Across all feedback-controlled configurations, substantial coolant-flow actuation commences around $t = 1500$ s, coinciding with the onset of sustained high heat generation in the latter phase of the mWLTP. Prior to this period, the optimiser assigns little or no coolant flow because the cell temperature remains close to or below the reference. This is consistent with the one-sided deviation metric and the cooling-only nature of the actuation. Coolant flow is primarily used to suppress overshoot above T^r, rather than to force unnecessary cooling when the cell is already within the desired operating range.

Taken together, these results show that the proposed ITSC system provides the best overall trade-off among the considered configurations. While ES attains the lowest temperature rise through continuous full-flow operation, it lacks feedback adaptation and does not optimise coolant usage relative to the reference. ITSC, by contrast, actively balances temperature regulation and thermal gradient minimisation under the same flow budget. A quantitative summary of the performance metrics is reported in Table I.

Table 1: Maximum values of thermal performance metrics under the mWLTP for the proposed ITSC system and benchmarks.

Configurations	$T_{\max}$	T_{mean}	$e_{\max}$	e_{mean}	$\delta T_{\max}$	δT_{mean}
ITSC	39.47	38.65	4.47	3.65	**1.17**	**0.61**
SC	47.66	46.22	12.66	11.22	4.05	2.02
bTC	43.55	42.47	8.55	7.47	2.38	1.36
bTSC	42.65	41.50	7.65	6.50	2.34	1.52
btTC	39.71	39.10	4.71	4.10	1.76	0.96
ES	**36.45**	**34.54**	**1.45**	**0.36**	4.57	2.28

The mean coolant temperature evolution in the three ITSC channels is shown in Fig. 7(a). The surface channel exhibits the highest coolant temperature, compared to the top and bottom channels. This follows directly from

the flow allocation in Fig. 7(b), where the controller assigns only a small fraction of the total coolant flow to the surface channel and substantially larger flow rates to the top and bottom channels. This trend agrees with the coolant energy balance in (B.14). Defining the characteristic coolant residence time as $\tau_{\mathrm{res}} = V_c/\dot{\mathcal{V}}$, the advective exchange with the inlet scales as $\dot{\mathcal{V}}/V_c = 1/\tau_{\mathrm{res}}$. At lower flow rates, the residence time increases and the advective replacement of warmed coolant by colder inlet coolant weakens. As a result, the coolant temperature rises toward the local battery temperature T_{bc} through the convective heating term. Conversely, at higher flow rates, advective transport dominates and constrains the coolant temperature more closely to $T_{cl,\mathrm{in}}$. This increases the battery-to-coolant temperature difference and therefore enhances heat removal.

5 Conclusion

This article has proposed an integrated tab-surface cooling (ITSC) system for battery thermal management, formulated as a receding-horizon optimal control problem in which coolant is dynamically allocated among the lateral surface and tab channels. To enable this formulation, a first-principles coolant model with volumetric flow rate as the manipulated input was derived and coupled with battery thermal and electrical models. A valve-actuation model was also incorporated to capture actuator-level coolant routing. The resulting battery–coolant model served as the predictive model within an MPC framework targeting temperature regulation and thermal gradient minimisation.

To reduce the prohibitive computational expense, a real-time iteration (RTI)-MPC was developed for the proposed coolant-allocation problem. By locally linearising the coolant dynamics and solving a single quadratic programme per sampling instant, the RTI-MPC closely reproduced the NMPC thermal response while reducing the online computation time well below the sampling interval. This demonstrates that the proposed RTI-MPC can retain near-NMPC thermal performance at a computational cost compatible with real-time implementation.

Evaluation results under a real-world driving profile showed that the ITSC system provides the lowest thermal gradients among all considered conventional cooling configurations. Also, it achieved the lowest temperature tracking error among the feedback-controlled cooling configurations. The rule-

based equal-split configuration attained the lowest tracking errors due to the continuous use of the full coolant flow capacity across all channels. This was, however, achieved without feedback adaptation and at the expense of larger thermal gradients. The ITSC system, by contrast, actively allocated coolant according to the instantaneous thermal demand, balancing temperature tracking and gradient minimisation under the same total flow budget.

Future work will address joint pump-flow and coolant-allocation control, where the total coolant flow rate is treated as an additional decision variable. Further extensions will incorporate explicit degradation- and energy-aware cost formulations to quantify long-term benefits, and will apply the proposed framework to module- and pack-level assemblies where inter-cell thermal interactions and coolant-network constraints become significant.

Appendix A Equivalent circuit model matrices

This appendix provides the explicit expressions of the system matrices associated with ECM (B.10).

$$\mathcal{A}_e = \begin{bmatrix} 0 & 0 & \cdots & 0 \\ 0 & -\frac{1}{R_1 C_1} & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & -\frac{1}{R_p C_p} \end{bmatrix}, \; \mathcal{B}_e = \begin{bmatrix} -\frac{1}{3600 C_{nom}} \\ \frac{1}{C_1} \\ \vdots \\ \frac{1}{C_p} \end{bmatrix}, \tag{B.48a}$$

$$\mathcal{C}_e = \begin{bmatrix} 0 & -1 & \cdots & -1 \end{bmatrix}, \qquad D_e = -R_0, \tag{B.48b}$$

where C_{nom} is the nominal charge capacity of the cell.

Appendix B Derivation of the mean coolant temperature dynamics

This appendix provides the detailed derivation of the coolant temperature dynamics presented in (B.14). We apply the first law of thermodynamics to a control volume consisting of the coolant in thermal contact with the battery surface. The overall energy balance is given by

$$\frac{dE}{dt} = Q_{\text{conv}} + \dot{E}_g + \dot{E}_{\text{in}} - \dot{E}_{\text{out}} - \dot{W}_f, \tag{B.49}$$

where E is the total energy within the control volume, Q_{conv} is the convective heat transfer from the battery surface to the coolant, $\dot{E}_g$ is the heat generated within the coolant, $\dot{E}_{\text{in}}$ and $\dot{E}_{\text{out}}$ are the energy fluxes at the inlet and outlet, respectively, and $\dot{W}_f$ is the flow work rate.

Under the stated simplifications, the stored energy reduces to sensible thermal energy, yielding

$$\frac{dE}{dt} = m_c c_{p,c} \frac{d\bar{T}_{cl}}{dt}, \tag{B.50}$$

where m_c is the coolant mass in the control volume.

The convective heat transfer is given by Newton's law of cooling:

$$Q_{\text{conv}} = hA_{ht}(T_s - \bar{T}_{cl}). \tag{B.51}$$

Since the coolant is passive, no internal heat generation occurs, i.e., $\dot{E}_g = 0$.

The energy transport due to fluid motion is expressed using the specific enthalpy formulation. For a mass flow rate $\dot{m}_c$, the net energy flow becomes

$$\dot{E}_{\text{in}} - \dot{E}_{\text{out}} = \dot{m}_c c_{p,c}(T_{cl,\text{in}} - T_{cl,\text{out}}). \tag{B.52}$$

Assuming a lumped control volume, the outlet temperature is approximated by the mean coolant temperature, i.e., $T_{cl,\text{out}} \approx \bar{T}_{cl}$. Neglecting mechanical energy contributions and pressure work, substituting (B.50)–(B.52) into (B.49) yields

$$m_c c_{p,c} \frac{d\bar{T}_{cl}}{dt} = hA_{ht}(T_s - \bar{T}_{cl}) + \dot{m}_c c_{p,c}(T_{cl,\text{in}} - \bar{T}_{cl}). \tag{B.53}$$

Finally, expressing the coolant mass and mass flow rate in terms of density and volumetric flow rate, $m_c = \rho_c V_c$ and $\dot{m}_c = \rho_c \dot{\mathcal{V}}$, yields

$$\frac{d\bar{T}_{cl}}{dt} = \frac{hA_{ht}(T_s - \bar{T}_{cl})}{\rho_c V_c c_{p,c}} + \frac{\dot{\mathcal{V}}(T_{cl,\text{in}} - \bar{T}_{cl})}{V_c}, \tag{B.54}$$

which is identical to (B.14).

Appendix C Determining convection coefficients

To determine the appropriate convection coefficients associated with each cooling channel, we utilise the general relation in (B.13). The hydraulic diameter D_h depends on the specific geometry of each channel.

For the lateral surface, the coolant flows through a narrow annular channel of height L and thickness ϵ_s. Approximating the local geometry as a rectangular duct, the hydraulic diameter is given by

$$D_{h,s} = \frac{4A_{cf,s}}{P_s}, \tag{B.55}$$

where $A_{cf,s} = L\epsilon_s$ is the flow cross-sectional area and $P_s = 2(L + \epsilon_s)$ is the wetted perimeter.

For the top and bottom channels, the coolant flows between parallel plates of thickness ϵ_t and ϵ_b, respectively. In this case, the hydraulic diameters are given by [26]

$$D_{h,t} = 2\epsilon_t, \quad D_{h,b} = 2\epsilon_b. \tag{B.56}$$

Substituting (B.55) and (B.56) into (B.13), the convection coefficients become

$$h_s = \frac{\mathrm{Nu} \cdot k_c(L + \epsilon_s)}{2L\epsilon_s}, \quad h_t = \frac{\mathrm{Nu} \cdot k_c}{2\epsilon_t}, \quad h_b = \frac{\mathrm{Nu} \cdot k_c}{2\epsilon_b}. \tag{B.57}$$

While ϵ_t and ϵ_b are often equal, the known higher heat generation from the positive tab areas compared to the negative [46], [47] may necessitate designing $\epsilon_t \neq \epsilon_b$ to accommodate the resulting thermal asymmetry. D_h can also serve as a design parameter to guide the optimal sizing of the cooling channels.

References

[1] S. Choi and G. Wang, "Advanced lithium-ion batteries for practical applications: Technology, development, and future perspectives," *Adv. Mater. Technol.*, vol. 3, no. 9, p. 1 700 376, 2018.

[2] S. Ma et al., "Temperature effect and thermal impact in lithium-ion batteries: A review," *Prog. Nat. Sci.: Mater. Int.*, vol. 28, no. 6, pp. 653–666, 2018.

[3] C. Bibin, M. Vijayaram, V. Suriya, R. S. Ganesh, and S. Soundarraj, "A review on thermal issues in Li-ion battery and recent advancements in battery thermal management system," *Mater. Today Proc.*, vol. 33, pp. 116–128, 2020.

[4] J. Lin, X. Liu, S. Li, C. Zhang, and S. Yang, "A review on recent progress, challenges and perspective of battery thermal management system," *Int. J. Heat Mass Transf.*, vol. 167, p. 120 834, 2021.

[5] P. R. Tete, M. M. Gupta, and S. S. Joshi, "Developments in battery thermal management systems for electric vehicles: A technical review," *J. Energy Storage*, vol. 35, p. 102 255, 2021.

[6] C. Roe et al., "Immersion cooling for lithium-ion batteries–a review," *J. Power Sources*, vol. 525, p. 231 094, 2022.

[7] Y. Zhao, L. B. Diaz, Y. Patel, T. Zhang, and G. J. Offer, "How to cool lithium ion batteries: Optimising cell design using a thermally coupled model," *J. Electrochem. Soc.*, vol. 166, no. 13, A2849–A2859, 2019.

[8] I. A. Hunt, Y. Zhao, Y. Patel, and G. Offer, "Surface cooling causes accelerated degradation compared to tab cooling for lithium-ion pouch cells," *J. Electrochem. Soc.*, vol. 163, no. 9, A1846, 2016.

[9] Y. Zhao, Y. Patel, T. Zhang, and G. J. Offer, "Modeling the effects of thermal gradients induced by tab and surface cooling on lithium ion cell performance," *J. Electrochem. Soc.*, vol. 165, no. 13, A3169–A3178, 2018.

[10] O. Dondelewski et al., "The role of cell geometry when selecting tab or surface cooling to minimise cell degradation," *eTransportation*, vol. 5, p. 100 073, 2020.

[11] C. Bolsinger and K. P. Birke, "Effect of different cooling configurations on thermal gradients inside cylindrical battery cells," *J. Energy Storage*, vol. 21, pp. 222–230, 2019.

[12] S. Li et al., "Optimal cell tab design and cooling strategy for cylindrical lithium-ion batteries," *J. Power Sources*, vol. 492, p. 229 594, 2021.

[13] T. Ahmad, A. Mishra, S. Ghosh, and C. S. Casari, "Identifying efficient cooling approach of cylindrical lithium-ion batteries," *Energy Technol.*, vol. 10, no. 2, p. 2 100 888, 2022.

[14] T. G. Tranter, R. Timms, P. R. Shearing, and D. Brett, "Communication—prediction of thermal issues for larger format 4680 cylindrical cells and their mitigation with enhanced current collection," *J. Electrochem. Soc.*, vol. 167, no. 16, p. 160 544, 2020.

[15] M. Taghizadeh, A. Ghaffari, and F. Najafi, "Modeling and identification of a solenoid valve for PWM control applications," *C. R. Mecanique*, vol. 337, no. 3, pp. 131–140, 2009.

[16] S. V. Angadi and R. L. Jackson, "A critical review on the solenoid valve reliability, performance and remaining useful life including its industrial applications," *Eng. Fail. Anal.*, vol. 136, p. 106 231, 2022.

[17] M. K. Kazimierczuk, *Pulse-Width Modulated DC-DC Power Converters*. John Wiley & Sons, 2015.

[18] R. W. Erickson, "DC–DC power converters," *Wiley Encycl. Electr. Electron. Eng.*, 2001.

[19] G. K. Peprah, Y. Huang, T. Wik, F. Altaf, and C. Zou, "Thermal modeling of battery cells for optimal tab and surface cooling control," *IEEE Trans. Control Syst. Technol.*, vol. 34, no. 4, pp. 1831–1844, 2026.

[20] G. K. Peprah, T. Wik, Y. Huang, F. Altaf, and C. Zou, "Control-oriented 2D thermal modelling of cylindrical battery cells for optimal tab and surface cooling," in *2024 IEEE American Control Conf. (ACC)*, pp. 4651–4656.

[21] D. W. Hahn and M. N. Özisik, *Heat conduction*. John Wiley & Sons, 2012.

[22] X. Hu, S. Li, and H. Peng, "A comparative study of equivalent circuit models for Li-ion batteries," *J. Power Sources*, vol. 198, pp. 359–367, 2012.

[23] Y. Xie, S. Shi, J. Tang, H. Wu, and J. Yu, "Experimental and analytical study on heat generation characteristics of a lithium-ion power battery," *Int. J. Heat Mass Transf.*, vol. 122, pp. 884–894, 2018.

[24] S. Ghasemi, M. Khoshvaght-Aliabadi, M. J. Asadollahzadeh, and N. Salehi, "Improving heat transfer in indirect liquid-based battery thermal management systems through turbulator-equipped conical flow paths," *Int. J. Heat Mass Transf.*, vol. 229, p. 125 701, 2024.

[25] W. Zeng, C. Ma, S. Hu, S. Li, and Y. Zhang, "The performance investigation and optimization of reciprocating flow applied for liquid-cooling-based battery thermal management system," *Energy Convers. Manage.*, vol. 292, p. 117 378, 2023.

[26] F. P. Incropera and D. P. DeWitt, *Fundamentals of Heat and Mass Transfer*. New York, NY: Wiley, 1985.

[27] Y. Xie, C. Wang, X. Hu, X. Lin, Y. Zhang, and W. Li, "An MPC-based control strategy for electric vehicle battery cooling considering energy saving and battery lifespan," *IEEE Trans. Veh. Technol.*, vol. 69, no. 12, pp. 14 657–14 673, 2020.

[28] L. Acker, P. Hofmann, and J. Konrad, "Predictive battery thermal management for fast charging of electric vehicles using nonlinear model predictive control and dynamic programming," *Automot. Engine Technol.*, vol. 11, no. 1, 2026, Art. no. 1.

[29] G. K. Peprah, F. Liberati, F. Altaf, G. Osei-Dadzie, A. Di Giorgio, and A. Pietrabissa, "Optimal load sharing in reconfigurable battery systems using an improved model predictive control method," in *Proc. 29th Mediterranean Conf. Control Autom. (MED)*, IEEE, 2021, pp. 979–984.

[30] G. K. Peprah, T. Wik, Y. Huang, F. Altaf, and C. Zou, "Model predictive cooling control of cylindrical battery cells through tab and surface channels," in *Proc. IEEE Amer. Control Conf. (ACC)*, 2025, pp. 4071–4076.

[31] X.-G. Yang and C.-Y. Wang, "Understanding the trilemma of fast charging, energy density and cycle life of lithium-ion batteries," *J. Power Sources*, vol. 402, pp. 489–498, 2018.

[32] M. Schimpe, M. E. von Kuepach, M. Naumann, H. C. Hesse, K. Smith, and A. Jossen, "Comprehensive modeling of temperature-dependent degradation mechanisms in lithium iron phosphate batteries," *J. Electrochem. Soc.*, vol. 165, no. 2, A181–A193, 2018.

[33] T. Waldmann, M. Wilka, M. Kasper, M. Fleischhammer, and M. Wohlfahrt-Mehrens, "Temperature dependent ageing mechanisms in lithium-ion batteries: A post-mortem study," *J. Power Sources*, vol. 262, pp. 129–135, 2014.

[34] M. Diehl, H. G. Bock, and J. P. Schlöder, "A real-time iteration scheme for nonlinear optimization in optimal feedback control," *SIAM J. Control Optim.*, vol. 43, no. 5, pp. 1714–1736, 2005.

[35] J. Lopez-Sanz et al., "Thermal management in plug-in hybrid electric vehicles: A real-time nonlinear model predictive control implementation," *IEEE Trans. Veh. Technol.*, vol. 66, no. 9, pp. 7751–7760, 2017.

[36] T. Fischer, T. Kraus, C. Kirches, and F. Gauterin, "Demonstration of a nonlinear model predictive control of a thermal management system for electric vehicles in real-time," in *Proc. IEEE Conf. Control Technol. Appl. (CCTA)*, 2018, pp. 676–682.

[37] S. Gros, M. Zanon, R. Quirynen, A. Bemporad, and M. Diehl, "From linear to nonlinear MPC: Bridging the gap via the real-time iteration," *Int. J. Control*, vol. 93, no. 1, pp. 62–80, 2020.

[38] Y.-S. Duh, J.-H. Theng, C.-C. Chen, and C.-S. Kao, "Comparative study on thermal runaway of commercial 14500, 18650 and 26650 LiFePO4 batteries used in electric vehicles," *J. Energy Storage*, vol. 31, p. 101 580, 2020.

[39] S. Mohan, Y. Kim, and A. G. Stefanopoulou, "Estimating the power capability of Li-ion batteries using informationally partitioned estimators," *IEEE Trans. Control Syst. Technol.*, vol. 24, no. 5, pp. 1643–1654, 2016.

[40] P. Mock et al., "The WLTP: How a new test procedure for cars will affect fuel consumption values in the EU," *Intern. council on clean transp.*, vol. 9, no. 3547, 2014.

[41] G.-H. Kim and A. Pesaran, "Battery thermal management system design modeling," National Renewable Energy Lab. (NREL), Golden, CO, Tech. Rep., 2006.

[42] W. Wu, S. Wang, W. Wu, K. Chen, S. Hong, and Y. Lai, "A critical review of battery thermal performance and liquid based battery thermal management," *Energy Convers. Manage.*, vol. 182, pp. 262–281, 2019.

[43] J. Lofberg, "Yalmip: A toolbox for modeling and optimization in matlab," in *2004 IEEE Int. Conf. on Robotics and Automation (IEEE Cat. No. 04CH37508)*, pp. 284–289.

[44] D. Worwood, R. Algoo, R. J. McGlen, J. Marco, and D. Greenwood, "A study into different cell-level cooling strategies for cylindrical lithium-ion cells in automotive applications," *Int. J. Powertrains*, vol. 7, no. 1-3, pp. 199–226, 2018.

[45] L. I. Uwalaka, Q. Yao, P. Kollmeyer, and A. Emadi, "Review of production electric vehicle battery thermal management systems and experimental testing of a production battery module," *SAE Int. J. Adv. Curr. Pract. Mobil.*, vol. 7, no. 2, pp. 798–809, 2025.

[46] F.-q. An, W.-n. Zhou, and P. Li, "Modeling and validation on the respective contributions of positive and negative tabs for 18650 cell," *J. Energy Storage*, vol. 20, pp. 439–450, 2018.

[47] P. Lyu, X. Liu, C. Liu, and Z. Rao, "Experimental and modeling investigation on thermal risk evaluation of tabs for pouch-type lithium-ion battery and the relevant heat rejection strategies," *Int. J. Heat Mass Transf.*, vol. 202, p. 123 770, 2023.

PAPER C

Degradation-Aware Model Predictive Coolant Allocation for Integrated Tab-Surface Cooling of Battery Cells

Godwin K. Peprah, Phil Kollmeyer, Torsten Wik, Changfu Zou

Manuscript. To be submitted to IEEE Transactions on Control Systems Technology.
The version reproduced in this thesis is the manuscript dated 20 November 2026.
A later submitted or published version may differ.

© 2026 The Authors

The layout has been revised.

Abstract

This work proposes a degradation-aware coolant allocation framework for an integrated tab–surface cooling (ITSC) system, in which coolant is dynamically allocated among the lateral surface and tab channels to minimise battery degradation. A coupled distributed electrothermal–degradation model is developed to capture the spatial distributions of temperature, current density, state of charge, and capacity loss. Model evaluations reveal that intra-cell electrothermal and degradation non-uniformities are considerably more pronounced in larger cells, despite relatively small differences in spatially averaged capacity loss. The resulting nonlinear coolant allocation problem is solved using real-time iteration model predictive control (RTI-MPC), reducing the mean online execution time from 4.6 s to 0.1 s relative to nonlinear MPC, with negligible deterioration in closed-loop performance. Evaluations under repeated drive–charge sequences against conventional cooling configurations demonstrate that ITSC achieves the best overall trade-off between accumulated capacity loss and coolant utilisation. Furthermore, under comparable coolant utilisation, direct degradation minimisation achieves lower accumulated capacity loss than thermal-proxy control, despite the latter maintaining a lower mean temperature and smaller spatial temperature gradients.

1 Introduction

Lithium-ion battery packs constitute a substantial share of the cost of an electric vehicle (EV) drivetrain [1], [2]. Their performance, safety, and lifespan are strongly dependent on operating conditions, particularly temperature, current, and state of charge (SoC). Elevated temperatures accelerate parasitic reactions and associated capacity fade, whereas charging at lower temperatures can increase transport limitations and the susceptibility to lithium plating [3]. Consequently, battery thermal management systems (BTMS)s [4] play an important role not only in maintaining safe operating temperatures, but

also in preserving battery health and extending service life.

The thermal state of a battery cannot, in general, be fully characterised by a single temperature. Owing to anisotropic heat conduction, non-uniform heat rejection, cell geometry, and the locations of cooling interfaces, large temperature differences or thermal gradients can develop within a cell [3], [5]. These gradients are particularly relevant to battery ageing because electrochemical reaction rates, internal resistance, and degradation rates are temperature dependent [3]. Spatial temperature differences can give rise to non-uniform current and SoC distributions, which in turn influence local heat generation and degradation. Experimental and modelling studies have increasingly demonstrated the importance of these effects. Carter *et al.* [6] experimentally showed that the direction of an imposed intra-cell thermal gradient can alter the dominant degradation mode and substantially affect the resulting capacity loss. Du *et al.* [7] similarly observed increased capacity fade under spatially non-uniform thermal conditions in a large-format pouch cell. More recently, Li *et al.* [8] developed a distributed electrothermal–degradation model and demonstrated, for the investigated cell and operating conditions, that an intra-cell temperature difference of approximately 3 °C can trigger positive feedback between local current and resistance inhomogeneities, resulting in an approximately threefold increase in the rate of available capacity loss.

The cooling interface plays an important role in determining this spatial thermal behaviour [9]. Prior studies [10], [11], [12], [13], [14] have shown that surface cooling provides a comparatively large heat-transfer area and can effectively suppress the mean temperature and overall cell temperature rise, whereas cooling through the electrical tabs can exploit the high thermal conductivity of the current collectors to achieve a more uniform internal temperature distribution. These studies have also reported different degradation behaviour and trade-off between surface- and tab-cooled cells, motivating battery thermal management strategies that exploit the complementary benefits of both cooling approaches [15], [16].

Capturing these effects within a control framework requires a thermal model that retains the relevant spatial information while remaining computationally tractable. High-fidelity thermal models based on partial differential equations and numerical discretisation methods such as the finite element method can accurately resolve the temperature field [17], but their computational burden limits their suitability for online predictive control. Lumped thermal mod-

els [18], [19], [20] are attractive for battery management because of their low order, but provide only a small number of representative temperatures and cannot resolve the complete intra-cell temperature distribution. Similarly, conventional lumped electrical models represent the electrical behaviour of the cell using spatially uniform quantities and therefore cannot capture current redistribution and local SoC differences arising from non-uniform thermal conditions. These limitations become increasingly important under aggressive electrical loading and spatially non-uniform cooling. In our previous work [21], [22], a computationally efficient two-dimensional (2D) thermal model was therefore developed to reconstruct the spatial temperature distribution of battery cells and facilitate integrated tab–surface cooling.

A corresponding challenge exists in degradation modelling. Numerous battery ageing models have been developed to predict capacity and power fade under different operating conditions [23], [24], [25]. Among these, Schimpe *et al.* [26] developed a semi-empirical capacity loss model for lithium-iron-phosphate (LFP)/graphite cells in which calendar ageing and several temperature-dependent cycling-ageing mechanisms are represented separately. The model provides a useful balance between physical interpretability and computational simplicity. However, its degradation stress factors are evaluated using cell-level temperature, SoC, and current. It therefore cannot distinguish between cells having the same mean operating conditions but different internal electrothermal distributions. Moreover, owing to the nonlinear dependence of degradation rates on these quantities, evaluating degradation using spatially averaged operating conditions is not generally equivalent to spatially averaging the local degradation rates. A distributed formulation is therefore needed to investigate both the spatially averaged capacity loss and its distribution within the cell.

In parallel with developments in battery degradation modelling, battery ageing and lifetime have been incorporated into predictive thermal and energy management strategies [27], [28], [29], [30]. These studies demonstrate that explicitly accounting for battery health can alter optimal thermal and energy management decisions. However, existing health-conscious BTMS strategies generally employ lumped electrothermal and degradation models and therefore do not explicitly account for the coupled effects of intra-cell temperature non-uniformity, current redistribution, and spatially varying degradation. Furthermore, although distributed electrothermal–degradation models have been

developed [8], their integration into online predictive control frameworks for battery thermal management remains limited.

Using the spatially resolved thermal model from [21], our previous work [16] addressed the optimal coolant allocation problem through an integrated tab–surface cooling (ITSC) architecture, in which a finite coolant supply was dynamically distributed among the available battery cooling channels. The allocation problem was implemented using real-time iteration model predictive control (RTI-MPC) [31], [32], with temperature tracking and thermal gradient penalties used as control objectives. Although these quantities are physically motivated proxies for battery health, minimising them does not necessarily minimise degradation because the underlying ageing mechanisms exhibit different and, under some operating conditions, competing dependencies on temperature, SoC, and current. Direct degradation minimisation instead enables the available coolant resource to be allocated spatially and temporally according to its predicted influence on the active degradation mechanisms.

To address these gaps, this work develops a degradation-aware predictive control framework for integrated battery cooling. The main contributions are summarised as follows:

1. A coupled distributed electrothermal–degradation model is developed to capture the spatial distributions of temperature, current density, SoC, and capacity loss within battery cells. The conventional equivalent circuit model is extended to a spatially distributed formulation, while the degradation mechanisms in [26] are reformulated as dynamic state equations and evaluated locally. The resulting model accounts for the coupled effects of thermal non-uniformity, current redistribution, and spatially varying degradation, and is integrated with existing coolant and valve-actuation models.

2. A degradation-aware optimal coolant allocation problem is proposed, where the available coolant resource is dynamically distributed across the surface and tab channels to minimise the incremental capacity loss at a reduced cooling effort. Unlike conventional thermal-proxy control, the proposed formulation explicitly accounts for the competing dependencies of the individual degradation mechanisms on the local electrothermal operating conditions.

3. The resulting nonlinear control problem incorporating the distributed

electrothermal–degradation dynamics is solved for the first time within an RTI-MPC framework to reduce its online computational burden.

The proposed modelling and control framework is systematically evaluated for two cylindrical LFP cells of different capacities under repeated discharging and charging conditions to investigate the influence of cell size and intra-cell electrothermal non-uniformity on battery degradation. The control framework is subsequently evaluated against conventional cooling configurations and a thermal-proxy controller to assess the achievable battery health and coolant utilisation trade-offs. Although demonstrated here using cylindrical cells, the modelling and control framework can be extended to pouch and prismatic cell formats by employing appropriate spatial thermal, electrical, and coolant models.

The remainder of this article is organised as follows. Section 2 presents the modelling of the battery and coolant system. Section 3 evaluates the proposed distributed degradation model. Section 4 details the proposed control framework. Section 5 discusses the controller evaluation results. Finally, Section 6 concludes the work.

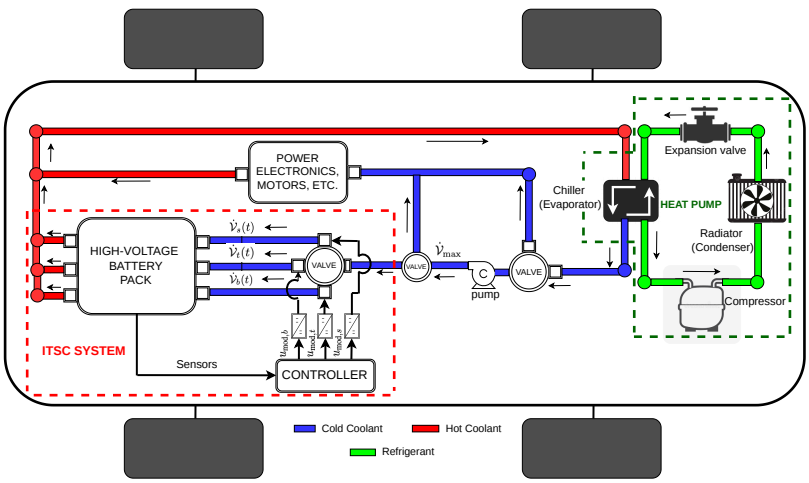

Figure 1: Representative EV BTMS incorporating the ITSC architecture [16]. The ITSC system distributes the available coolant resource among the lateral surface, top, and bottom cooling channels of the battery cell.

2 Integrated Tab–Surface Cooling System Modelling

This work adopts the ITSC architecture developed in our previous work [16]. As illustrated in Fig. 1, the ITSC system employs three independently actuated coolant channels in thermal contact with the lateral surface, top, and bottom areas of the cylindrical cell. A finite coolant supply is distributed among these channels, allowing the cooling effort to be dynamically allocated across the different cell interfaces. The corresponding channel flow rates satisfy

$$\dot{\mathcal{V}}_s(t) + \dot{\mathcal{V}}_t(t) + \dot{\mathcal{V}}_b(t) \leq \dot{\mathcal{V}}_{\max}, \tag{C.1}$$

where $\dot{\mathcal{V}}_{\max}$ denotes the maximum available volumetric coolant flow rate and the subscripts s, t, and b refer to the surface, top, and bottom cooling channels, respectively. The valve-actuation system is also retained from [16]. The following subsections present the battery and coolant models required for the subsequent control design.

2.1 Battery Modelling

The battery is represented by coupled electrical, thermal, and degradation subsystems. The electrical model describes the spatial distribution of current and state of charge (SoC), while the thermal model captures the temperature field within the cell. These distributions are subsequently incorporated into the degradation model to account for spatially non-uniform ageing.

Battery electrical dynamics

To ensure a balance between accuracy and computational efficiency, we adopt an equivalent circuit model (ECM) [16], [33] to model the battery's electrical behaviour. The ECM comprises an open-circuit voltage source V_{ocv}, a series ohmic resistance R_0, and two parallel resistance–capacitance (R_jC_j, $j \in \{1, 2\}$) branches connected in series to capture the cell's polarisation dynamics. The ECM is extended to a spatially distributed formulation over the physical cylindrical cell domain Ω_b to capture intra-cell variations in current and SoC.

Assuming ideal electrical current collectors, all spatial locations share a common terminal voltage V_t. However, spatial variations in temperature,

SoC, polarisation voltages, and degradation states give rise to a non-uniform current distribution within the cell. Let $i(r, z, t)$ denote the local volumetric current density, defined over radial r and axial z coordinates at time t, with $i > 0$ corresponding to discharging and $i < 0$ to charging. The cell-equivalent local current is defined as

$$I_\ell(r, z, t) = i(r, z, t)V_b, \tag{C.2}$$

where V_b denotes the cell volume. The spatially distributed electrical dynamics are then expressed as

$$\dot{Z}(r, z, t) = -\frac{I_\ell(r, z, t)}{3600 C_{\text{avail}}(r, z, t)}, \tag{C.3a}$$

$$\dot{V}_1(r, z, t) = -\frac{V_1(r, z, t)}{R_1(r, z, t)C_1(r, z, t)} + \frac{I_\ell(r, z, t)}{C_1(r, z, t)}, \tag{C.3b}$$

$$\dot{V}_2(r, z, t) = -\frac{V_2(r, z, t)}{R_2(r, z, t)C_2(r, z, t)} + \frac{I_\ell(r, z, t)}{C_2(r, z, t)}, \tag{C.3c}$$

$$V_t(t) = V_{\text{ocv}}(Z(r, z, t)) - \sum_{j=1}^{2} V_j(r, z, t) - R_0(r, z, t)I_\ell(r, z, t), \tag{C.3d}$$

where $Z(r, z, t)$, V_j, and C_{avail} denote the local SoC, polarisation voltages, and available cell-equivalent local capacity, respectively. The ECM parameters R_0, R_j, and C_j depend on the local temperature, SoC, and operating mode (charging or discharging). Substituting (C.2) into (C.3d), the local current density is obtained as

$$i(r, z, t) = \frac{V_{\text{ocv}}(Z) - V_1 - V_2 - V_t}{R_0 V_b}. \tag{C.4}$$

The spatial current distribution satisfies the global current conservation condition

$$I_b(t) = \int_{\Omega_b} i(r, z, t)\, d\Omega_b = \frac{1}{V_b} \int_{\Omega_b} I_\ell(r, z, t) d\Omega_b, \tag{C.5}$$

where I_b denotes the applied battery current. Substituting (C.4) into (C.5) and rearranging yields the battery terminal voltage

$$V_t(t) = \frac{\displaystyle\int_{\Omega_b} \frac{V_{\text{ocv}}(Z) - V_1 - V_2}{R_0}\, \mathrm{d}\Omega_b - V_b I_b(t)}{\displaystyle\int_{\Omega_b} \frac{1}{R_0}\, \mathrm{d}\Omega_b}. \tag{C.6}$$

Defining the local electrical state vector as

$$X_e(r,z,t) = [Z(r,z,t), V_1(r,z,t), V_2(r,z,t)]^T, \tag{C.7}$$

the distributed electrical model can be expressed in state-space form as

$$\dot{X}_e(r,z,t) = \mathcal{A}_e X_e(r,z,t) + \mathcal{B}_e I_\ell(r,z,t), \tag{C.8a}$$

$$\mathcal{Y}_e(t) = V_t, \tag{C.8b}$$

where the system matrices A_e and B_e are spatially and temporally varying owing to their dependence on the local operating conditions. Their explicit expressions are provided in Appendix Appendix A. The electrical model output $\mathcal{Y}_e(t)$ is the battery terminal voltage (C.6).

For numerical tractability, the spatial integrals in (C.5) and (C.6) are evaluated using a tensor-product Gauss–Legendre quadrature in the radial and axial directions. Let N_e denote the total number of quadrature nodes and w_i the corresponding normalised volume weights, accounting for the cylindrical geometry and satisfying $\sum_{i=1}^{N_e} w_i = 1$. The volume average of a spatially varying quantity $f(r,z,t)$ is approximated as

$$\frac{1}{V_b}\int_{\Omega_b} f(r,z,t)\,\mathrm{d}\Omega_b \approx \sum_{i=1}^{N_e} w_i f_i(t), \tag{C.9}$$

where the subscript i denotes the quantity evaluated at the ith quadrature node. Applying (C.9) to the condition (C.5) gives

$$I_b(t) \approx \sum_{i=1}^{N_e} w_i I_{\ell,i}(t). \tag{C.10}$$

Similarly, the terminal voltage in (C.6) is evaluated as

$$V_t(t) = \frac{\displaystyle\sum_{i=1}^{N_e} w_i \frac{V_{\mathrm{ocv}}(Z_i) - V_{1,i} - V_{2,i}}{R_{0,i}} - I_b(t)}{\displaystyle\sum_{i=1}^{N_e} \frac{w_i}{R_{0,i}}}. \tag{C.11}$$

The resulting terminal voltage is used to determine the local currents from (C.2) and (C.4), yielding a coupled system of $3N_e$ ordinary differential equations describing the spatial electrical dynamics. Under spatially uniform operating conditions where $I_\ell(r,z,t) = I_b(t)$ throughout the cell domain, the distributed electrical model reduces to the conventional lumped ECM.

Battery thermal dynamics

The transient thermal dynamics of the cylindrical cell are captured using the two-dimensional (2D) thermal model developed in [21], [22]. Unlike in our previous work, where the volumetric heat generation was assumed to be spatially uniform, the thermal model is extended here to account for spatially distributed heat generation. The reduced-order state-space representation is given by

$$\mathcal{G}_T \dot{X}_T(t) = \mathcal{A}_T X_T(t) + \mathcal{B}_T U_T(t) + \mathcal{F}_T q(r, z, t) \tag{C.12a}$$

$$\mathcal{Y}_T(t) = \mathcal{C}_T X_T(t) + \mathcal{D}_T U_T(t) \tag{C.12b}$$

where

$$X_T = [c_{11}, \cdots, c_{1N}, c_{21}, \cdots, c_{2N}, \cdots, c_{M1}, \cdots, c_{MN}]^T \tag{C.13}$$

is the state vector, with the states $c_{mn}(t)$ denoting time-varying coefficients of the employed basis functions, $m \in \{1, \ldots, M\}$ and $n \in \{1, \ldots, N\}$. The tuning parameters M and N denote the number of basis functions used in the radial r and axial z directions of the cylindrical cell, respectively. The resulting model order is $O = MN$. The input vector U_T comprises the boundary cooling inputs at the lateral surface, top, and bottom sides of the cell,

$$U_T(t) = \begin{bmatrix} u_s(t) & u_t(t) & u_b(t) \end{bmatrix}^\top, \tag{C.14}$$

where

$$u_s(t) = h_s \bar{T}_{cl,s}(t), \quad u_t(t) = h_t \bar{T}_{cl,t}(t), \quad u_b(t) = -h_b \bar{T}_{cl,b}(t). \tag{C.15}$$

The variable h represents the convection coefficient, and the subscripts s, t, and b denote the cell's lateral surface, top, and bottom, respectively. $\bar{T}_{cl}$ is the mean coolant temperature, whose dynamics are briefly recapped in subsection 2.2. The output vector $\mathcal{Y}_T$ consists of temperatures at key locations within and on the cell. The output vector $\mathcal{Y}_T$ consists of temperatures at key locations within and on the cell. In particular, the volume-averaged cell temperature T_v and the root-mean-square (RMS) of the spatial temperature

gradient δT are defined as [16]

$$T_v(t) = \frac{1}{V_b} \int_{\Omega_b} T(r, z, t) \, \mathrm{d}\Omega_b, \tag{C.16a}$$

$$\delta T(t) = \left(\frac{1}{V_b} \int_{\Omega_b} \|\nabla T(r, z, t)\|_2^2 \, \mathrm{d}\Omega_b \right)^{1/2}, \tag{C.16b}$$

$$\nabla T(r, z, t) = \left[\frac{\partial T(r, z, t)}{\partial r}, \frac{\partial T(r, z, t)}{\partial z} \right]^T, \tag{C.16c}$$

where $T(r, z, t)$ is the temperature distribution of the cylindrical cell, which also provides the local temperatures required to evaluate the spatially varying ECM parameters in (C.3). The disturbance input $q(r, z, t)$ in (C.12a) denotes the spatially distributed volumetric heat generation rate within the cell. To incorporate this heat source into the thermal dynamics, the spectral-Galerkin projection is expressed as

$$\mathcal{F}_T q(r, z, t) = \int_{R_{\text{in}}}^{R_{\text{out}}} \int_0^L \boldsymbol{\phi}(r, z) q(r, z, t) \, r \, dz \, dr, \tag{C.17}$$

where $\boldsymbol{\phi}(\tilde{r}, \tilde{z}) \in \mathbb{R}^{MN}$ denotes the vector of Galerkin test functions evaluated at the normalised spatial coordinates $\tilde{r}$ and $\tilde{z}$ defined in [21]. R_{in}, R_{out}, and L denote the inner radius, outer radius, and length of the cylindrical cell, respectively. The factor r accounts for the cylindrical geometry of the cell. Under spatially uniform heat generation, $q(r, z, t) = q(t)$, (C.17) reduces to the heat source formulation in [21]. The remaining system matrices $\mathcal{G}_T$, $\mathcal{A}_T$, $\mathcal{B}_T$, $\mathcal{C}_T$, and $\mathcal{D}_T$ are retained from [21], [22].

Battery degradation dynamics

We adopt the semi-empirical degradation model of [26]. The model describes capacity loss due to calendar and cycling ageing through stress factors that depend on temperature, SoC, and current. We first reformulate its rate-based integral representation into a system of ordinary differential equations (ODEs) to facilitate its integration with the electrothermal model. In the original model, the stress factors were evaluated using lumped cell quantities. In the original formulation, the stress factors are evaluated using lumped cell quantities. To account for intra-cell variations in ageing, we extend the model to incorporate the spatial distributions of temperature, SoC, and current.

The total normalised capacity loss is represented as the superposition of calendar ageing Q_C, and three cycling-ageing mechanisms [26],

$$\begin{aligned} Q_{\text{loss}} &= Q_C + Q_H + Q_L + Q_{LS} && \text{(C.18a)} \\ &= k_C(\bar{T}, Z) \cdot \sqrt{\tau} \quad + \; k_H(\bar{T}) \cdot \sqrt{\varphi_{\text{tot}}} \\ &\quad + k_L(\bar{T}, I_{\text{ch}}) \cdot \sqrt{\varphi_{\text{ch}}} \; + \; k_{LS}(\bar{T}, I_{\text{ch}}, \text{SoC}) \cdot \varphi_{\text{ch}}, && \text{(C.18b)} \end{aligned}$$

where the cycling contributions, Q_H, Q_L, and Q_{LS} account for degradation associated with high-temperature cycling, low-temperature charging, and low-temperature charging at high SoC, respectively. φ_{tot} denotes the accumulated total charge and discharge throughput, φ_{ch} the charge-only throughput, and $\tau = t/3600$ the elapsed time in hours. $\bar{T}$ denotes a representative cell temperature in Kelvin, taken as the volumetric mean temperature T_v in this work, and I_{ch} is the charge current. The corresponding stress factors are detailed in Appendix B.a.

Degradation model state equations The relations in (C.18) assume constant operating conditions. However, in practice, the cell temperature, SoC, and current vary continuously. To address this, [26] evaluated the incremental capacity losses at the instantaneous operating point using a rate-based integral formulation

$$\begin{aligned} Q_{\text{loss}} = &\int k_C(\bar{T}, Z) \frac{1}{2\sqrt{\tau}} \, d\tau + \int k_H(\bar{T}) \frac{1}{2\sqrt{\varphi_{\text{tot}}}} \, d\varphi_{\text{tot}} \\ &+ \int k_L(\bar{T}, I_{\text{ch}}) \frac{1}{2\sqrt{\varphi_{\text{ch}}}} \, d\varphi_{\text{ch}} + \int k_{LS}(\bar{T}, I_{\text{ch}}, Z) \, d\varphi_{\text{ch}}, \end{aligned} \tag{C.19}$$

where τ and φ are time and throughput integration variables, respectively. To enable the application of the degradation model to cells of different nominal capacities under equivalent C-rates, the current supplied to the degradation model is scaled as

$$\tilde{I}_b(t) = \frac{C_{\text{ref}}}{C_0} I_b(t), \tag{C.20}$$

where C_{ref} denotes the nominal capacity of the reference cell used for degradation parameterisation and C_0 denotes that of the considered cell. Accordingly, $I_{\text{ch}} = \max(-\tilde{I}_b, 0)$.

To integrate the degradation model with the electrothermal battery model (C.8) and (C.12), we express (C.19) as a regularised system of degradation

state ODEs. The accumulated total and charging throughputs are defined as $\varphi_{\text{tot}}(t) = \frac{1}{3600}\int_0^t |\bar{I}_b(s)|\,ds$, $\varphi_{\text{ch}}(t) = \frac{1}{3600}\int_0^t I_{\text{ch}}(s)\,ds$, with the corresponding dynamics

$$\dot{\varphi}_{\text{tot}}(t) = \frac{|\bar{I}_b(t)|}{3600}, \qquad \dot{\varphi}_{\text{ch}}(t) = \frac{I_{\text{ch}}(t)}{3600}. \tag{C.21}$$

Using these throughput states, the differential form of (C.19) yields the individual capacity loss rates

$$\dot{Q}_C = \frac{k_C(\bar{T}, Z)}{7200\sqrt{\tau + \epsilon_\tau}}, \tag{C.22a}$$

$$\dot{Q}_H = k_H(\bar{T})\frac{|\bar{I}_b|}{7200\sqrt{\varphi_{\text{tot}} + \epsilon_{\text{tot}}}}, \tag{C.22b}$$

$$\dot{Q}_L = k_L(\bar{T}, I_{\text{ch}})\frac{I_{\text{ch}}}{7200\sqrt{\varphi_{\text{ch}} + \epsilon_{\text{ch}}}}, \tag{C.22c}$$

$$\dot{Q}_{LS} = k_{LS}(\bar{T}, I_{\text{ch}}, Z)\frac{I_{\text{ch}}}{3600}, \tag{C.22d}$$

where ϵ are small positive regularisation constants introduced to avoid singularities at zero elapsed time and accumulated throughput.

The degradation state vector is defined as

$$X_{\text{deg}}(t) = [\varphi_{\text{tot}},\ \varphi_{\text{ch}}, Q_C,\ Q_H,\ Q_L,\ Q_{LS}]^\top, \tag{C.23}$$

and the degradation model can be expressed in state-space form as

$$\dot{X}_{\text{deg}}(t) = f_{\text{deg}}(X_{\text{deg}}, \bar{T}, Z, \tilde{I}_b, \tau), \tag{C.24a}$$

$$\mathcal{Y}_{\text{deg}}(t) = \mathcal{C}_{\text{deg}} X_{\text{deg}}(t), \tag{C.24b}$$

where $\mathcal{C}_{\text{deg}} = [0,\ 0,\ 1,\ 1,\ 1,\ 1]$, and the output $\mathcal{Y}_{\text{deg}}$ represents the total normalised capacity loss Q_{loss} defined in (C.18). Accordingly, the capacity-based state of health (SoH) is defined as

$$\text{SoH}(t) = 1 - \mathcal{Y}_{\text{deg}}(t). \tag{C.25}$$

Distributed degradation model The lumped degradation model (C.22) neglects spatial variations in temperature, SoC, and current, which may cause different regions of the cell to age at different rates. To capture these variations, the degradation model is extended to the spatial domain Ω_b by evaluating the ageing mechanisms locally using the temperature distribution from (C.12) and the SoC and current distributions from (C.8).

Using the cell-equivalent local current defined in (C.2), the corresponding reference-capacity-scaled local current and charging current are given by

$$\tilde{I}_\ell(r,z,t) = \frac{C_{\mathrm{ref}}}{C_0} I_\ell(r,z,t), \tag{C.26a}$$

$$I_{\mathrm{ch}}(r,z,t) = \max(-\tilde{I}_\ell(r,z,t), 0). \tag{C.26b}$$

The local total and charging throughput dynamics are accordingly expressed as

$$\dot{\varphi}_{\mathrm{tot}}(r,z,t) = \frac{|\tilde{I}_\ell(r,z,t)|}{3600}, \qquad \dot{\varphi}_{\mathrm{ch}}(r,z,t) = \frac{I_{\mathrm{ch}}(r,z,t)}{3600}. \tag{C.27}$$

The local capacity loss rates are obtained by evaluating the stress factors in (C.22) using the corresponding local operating conditions. The total local capacity loss rate is thus given by

$$\dot{Q}_{\mathrm{loss}}(r,z,t) = \dot{Q}_C(r,z,t) + \dot{Q}_H(r,z,t) + \dot{Q}_L(r,z,t) + \dot{Q}_{LS}(r,z,t), \tag{C.28}$$

where the expressions of the individual local capacity loss rates are provided in Appendix B.b. The distributed degradation model retains the state and model structures defined in (C.23) and (C.24), respectively, with the states evaluated locally over the cell domain.

Assuming a uniform initial capacity density, the cell-level capacity loss rate is obtained by spatially averaging the local degradation rates,

$$\dot{Q}^{\mathrm{dist}}_{\mathrm{loss}}(t) = \frac{1}{V_b} \int_{\Omega_b} \dot{Q}_{\mathrm{loss}}(r,z,t)\, \mathrm{d}\Omega_b. \tag{C.29}$$

The accumulated distributed capacity loss is consequently given by

$$\begin{aligned} Q^{\mathrm{dist}}_{\mathrm{loss}}(t) &= Q^{\mathrm{dist}}_{\mathrm{loss}}(0) + \int_0^t \dot{Q}^{\mathrm{dist}}_{\mathrm{loss}}(s)\, \mathrm{d}s \\ &= \frac{1}{V_b} \int_{\Omega_b} Q_{\mathrm{loss}}(r,z,t)\, \mathrm{d}\Omega_b. \end{aligned} \tag{C.30}$$

For numerical evaluation, the local degradation states are propagated at each quadrature node using the corresponding local operating conditions. Applying (C.9) to (C.30), the cell-level capacity loss is approximated as

$$Q^{\mathrm{dist}}_{\mathrm{loss}}(t) \approx \sum_{i=1}^{N_e} w_i Q_{\mathrm{loss},i}(t). \tag{C.31}$$

The resulting distributed capacity degradation model comprises $6N_e$ states, with the corresponding degradation output and SoH given by

$$\mathcal{Y}_{\text{deg}}^{\text{dist}}(t) = Q_{\text{loss}}^{\text{dist}}(t), \tag{C.32a}$$

$$\text{SoH}^{\text{dist}}(t) = 1 - \mathcal{Y}_{\text{deg}}^{\text{dist}}(t). \tag{C.32b}$$

Under spatially uniform operating conditions and identical initial degradation states, the distributed degradation model reduces to the lumped degradation model in (C.24).

Battery electrothermal–degradation coupling

The electrical (C.8), thermal (C.12), and degradation model (C.32) are coupled through the spatial distributions of temperature, SoC, current, heat generation, and degradation. The applied battery current I_b drives the electrical dynamics and the resulting spatial variations in the electrical states and parameters give rise to non-uniform heat generation within the cell.

The spatially distributed volumetric heat generation comprises reversible and irreversible components, expressed as [16], [34]

$$q(r,z,t) = q_{\text{rev}}(r,z,t) + q_{\text{irr}}(r,z,t), \tag{C.33a}$$

$$q_{\text{rev}}(r,z,t) = -\frac{T}{V_b}\left(\frac{V_{\text{ocv}} - V_t - \sum_{j=1}^{2} V_j}{R_0}\right)\frac{\partial V_{\text{ocv}}}{\partial T}, \tag{C.33b}$$

$$q_{\text{irr}}(r,z,t) = \frac{1}{V_b}\left[\frac{\left(V_{\text{ocv}} - V_t - \sum_{j=1}^{2} V_j\right)^2}{R_0} + \sum_{j=1}^{2}\frac{V_j^2}{R_j}\right]. \tag{C.33c}$$

The reversible component accounts for the entropic heat associated with the electrochemical reactions, where $\frac{\partial V_{\text{ocv}}}{\partial T}$ denotes the entropic temperature coefficient. The irreversible component accounts for the ohmic and polarisation losses. For notational simplicity, the spatial and temporal arguments of the variables and parameters on the right-hand side of (C.33) are omitted. The resulting spatial heat generation provides the disturbance input to the thermal model through the projection in (C.17). Conversely, the obtained temperature distribution $T(r,z,t)$ influences the spatially varying ECM parameters and consequently the current distribution and heat generation. The electrothermal coupling is therefore bidirectional.

The local temperature, SoC, current, and accumulated throughputs determine the local degradation rates. The accumulated capacity loss reduces the available cell-equivalent local capacity according to

$$C_{\text{avail}}(r,z,t) = C_0\left(1 - Q_{\text{loss}}(r,z,t)\right), \tag{C.34}$$

where C_0 denotes the nominal cell capacity at the beginning of life. In addition to capacity fade, battery ageing leads to an increase in internal resistance due to both calendar and cycling ageing mechanisms [35], [36]. Such ageing-induced resistance variations can be incorporated into the electrical model as

$$R_j(r,z,t) = R_j^{\text{BOL}}(Z,T)\left(1 + \mu_{R,j}(r,z,t)\right), \;\; j \in \{0,1,2\}, \tag{C.35}$$

where R_j^{BOL} denotes the beginning-of-life resistance evaluated at the local SoC and temperature, while $\mu_{R,j}$ denotes the corresponding accumulated local relative resistance increase due to ageing, which depends on the operating history of the respective cell region. The resulting changes in available capacity and internal resistance influence the local electrical dynamics, current distribution, and heat generation, thereby affecting the subsequent temperature distribution and degradation rates. The electrothermal and degradation models are thus bidirectionally coupled.

2.2 Coolant modelling

In [16], the evolution of the bulk mean coolant temperature $\bar{T}_{cl}$ used in (C.15), was derived by applying an energy balance to a control volume comprising the coolant flowing through a channel in thermal contact with the battery in the ITSC system. The resulting dynamical system is expressed as [16]

$$\frac{d\bar{T}_{cl}(t)}{dt} = \frac{h(t)A_{ht}(T_{bc}(t) - \bar{T}_{cl}(t))}{\rho_c V_c c_{p,c}} + \frac{\dot{\mathcal{V}}(t)(T_{cl,\text{in}}(t) - \bar{T}_{cl}(t))}{V_c}, \tag{C.36}$$

where h, A_{ht}, T_{bc}, $T_{cl,\text{in}}$, ρ_c, V_c, and $c_{p,c}$ represent the convective heat transfer coefficient, effective heat transfer area, battery temperature at the battery–coolant interface, coolant inlet temperature, coolant density, coolant volume in the channel, and specific heat capacity, respectively. Eq. (C.36) is applied independently to the lateral surface, top, and bottom cooling channels. The corresponding effective volumetric flow rates are given by [16]

$$\dot{\mathcal{V}}_s(t) = \dot{\mathcal{V}}_{\max} u_{\text{mod},s}, \quad \dot{\mathcal{V}}_t(t) = \dot{\mathcal{V}}_{\max} u_{\text{mod},t}, \quad \dot{\mathcal{V}}_b(t) = \dot{\mathcal{V}}_{\max} u_{\text{mod},b}, \tag{C.37}$$

where u_{mod} is the valve duty-cycle generated through pulse-width modulation. The duty-cycles satisfy

$$0 \leq u_{\mathrm{mod}_j}(t) \leq 1, \quad j \in \{s, t, b\} \tag{C.38a}$$

$$u_{\mathrm{mod},s}(t) + u_{\mathrm{mod},t}(t) + u_{\mathrm{mod},b}(t) \leq 1, \tag{C.38b}$$

which constrain the total channel flow to the available pump capacity while permitting partial utilisation when the cooling demand is low. The coolant dynamics in (C.36) are nonlinear owing to the product between the coolant temperature state $\bar{T}_{cl}$ and the channel flow rate input $\dot{\mathcal{V}}$.

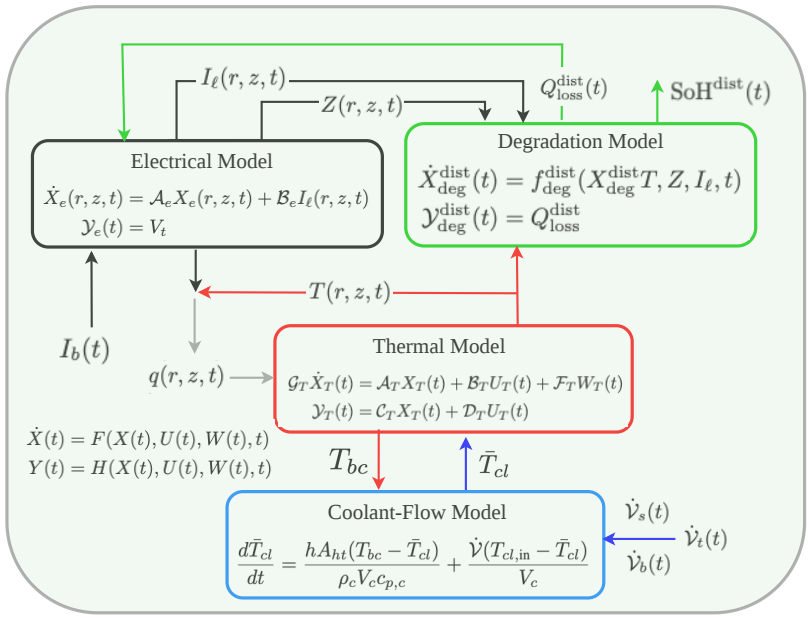

Figure 2: Interconnections among battery–coolant models.

2.3 Battery–coolant coupling

Heat exchange between the battery and coolant occurs bidirectionally through the battery–coolant contact interfaces. The mean coolant temperatures $\bar{T}_{cl}$ enter the battery thermal model (C.12) through the dynamic boundary inputs

defined in (C.15). Conversely, the thermal model provides the battery-coolant interface temperature T_{bc}, which govern the convective heat transfer term in the coolant dynamics (C.36). The interconnections among the electrical, thermal, degradation, and coolant subsystems are illustrated in Fig. 2.

Assembling the electrical (C.8), thermal (C.12), distributed degradation (C.28), and coolant (C.36) models yields the complete nonlinear system

$$\dot{X}(t) = F(X(t), U(t), W(t), t) \tag{C.39a}$$
$$Y(t) = H(X(t), U(t), W(t), t) \tag{C.39b}$$

where $F(\cdot)$ and $H(\cdot)$ denote the nonlinear state and output mappings, respectively. The complete state vector is defined as

$$X(t) = \left[X_e(t)^T, \quad X_T(t)^T, \quad (X_{\text{deg}}^{\text{dist}})^T(t), \quad X_{\text{cl}}^T(t)\right]^T. \tag{C.40}$$

The distributed degradation state vector can additionally accommodate resistance-growth states when resistance ageing is considered. The coolant state vector is

$$X_{\text{cl}}(t) = \left[\bar{T}_{\text{cl},s}(t), \quad \bar{T}_{\text{cl},t}(t), \quad \bar{T}_{\text{cl},b}(t)\right]^T. \tag{C.41}$$

The control input vector comprises the effective volumetric flow rates allocated to the three cooling channels,

$$U(t) = \left[\dot{\mathcal{V}}_s(t), \quad \dot{\mathcal{V}}_t(t), \quad \dot{\mathcal{V}}_b(t)\right]^T, \tag{C.42}$$

which are determined by the valve duty-cycles through (C.37). The exogenous disturbance vector is

$$W(t) = \left[I_b(t), \quad T_{\text{cl,in},s}(t), \quad T_{\text{cl,in},t}(t), \quad T_{\text{cl,in},b}(t)\right]^T. \tag{C.43}$$

The system outputs are assembled as

$$Y(t) = \left[\mathcal{Y}_e^T(t) \quad \mathcal{Y}_T^T(t) \quad \mathcal{Y}_{\text{deg}}^{\text{dist}}(t) \quad X_{\text{cl}}^T(t)\right]^T. \tag{C.44}$$

The mapping $F(\cdot)$ is given by

$$\dot{X}_e(t) = \mathcal{A}_e X_e(t) + \mathcal{B}_e \mathbf{I}_\ell(t), \tag{C.45a}$$
$$\dot{X}_T(t) = \mathcal{G}_T^{-1}\left(\mathcal{A}_T X_T(t) + \mathcal{B}_T U_T\big(X_{cl}(t)\big) + \mathcal{F}_T q(r,z,t)\right), \tag{C.45b}$$
$$\dot{X}_{\text{deg}}^{\text{dist}} = f_{\text{deg}}^{\text{dist}}(X_{\text{deg}}^{\text{dist}}(t), \mathbf{T}(t), \mathbf{Z}(t), \mathbf{I}_\ell(t)), \tag{C.45c}$$
$$\dot{X}_{cl}(t) = f_{cl}\left(X_{cl}(t), X_T(t), U(t), T_{cl,\text{in}}(t)\right), \tag{C.45d}$$

where $\mathcal{A}_e$ and $\mathcal{B}_e$ are the block-diagonal matrices assembled from the local electrical model matrices, and are dependent on the local operating conditions and degradation states. The vectors $\mathbf{T}$, $\mathbf{Z}$, and $\mathbf{I}_\ell$ contain the local temperatures, SoCs, and cell-equivalent currents evaluated at the quadrature nodes, respectively. The functions $f_{\text{deg}}^{\text{dist}}(\cdot)$ and $f_{\text{cl}}(\cdot)$ describe the nonlinear distributed degradation and coolant dynamics.

3 Degradation model evaluation

This section evaluates the distributed degradation model developed in Section 2.1 to investigate the influence of intra-cell electrothermal non-uniformity on battery ageing. The degradation predictions obtained using lumped and distributed models are compared to quantify the effects of spatial variations in temperature, SoC, and current on capacity loss. The sensitivity of the four ageing mechanisms to the principal operating conditions is also examined. Furthermore, the influence of cell geometry, electrical loading, and cooling configuration on the spatial electrothermal dynamics and the resulting degradation is investigated.

3.1 Simulation and analysis setup

The degradation analysis is conducted on two cylindrical LFP cells with nominal capacities of 2.3 Ah and 45 Ah. The former corresponds to an A123 ANR26650M1-A cell [37], while the latter represents the large-format cylindrical cell considered in [21]. To investigate the influence of cell geometry under identical material properties, both cells employ the thermophysical parameters adopted from [21], namely $\rho = 2118$ kgm^{-3}, $c_{p,b} = 795$ Jkg^{-1}K^{-1}, $k_r = 0.67$ Wm^{-1}K^{-1}, and $k_z = 66.6$ Wm^{-1}K^{-1}. The electrical model employs two RC branches, with resistances and capacitances dependent on SoC, temperature, and the charging or discharging operating mode, according to the parameterisation of the 2.3 Ah cell in [33]. The nonlinear V_{ocv}–SoC relationship shown in Fig. 3(a) is obtained from experimental data and implemented as a lookup table. For the 45 Ah cell, the electrical parameters are scaled

according to

$$\begin{aligned} R_j^{45} &= \frac{2.3}{45} R_j^{2.3}, \quad j \in \{0,1,2\}, \\ C_j^{45} &= \frac{45}{2.3} C_j^{2.3}, \quad j \in \{1,2\}, \end{aligned} \tag{C.46}$$

which preserves the RC time constants and voltage responses under equivalent C-rates and identical operating conditions. The electrical model of the 45 Ah cell is obtained through capacity-based parameter scaling rather than independent experimental parameterisation.

Following a spatial-resolution study considering 5, 8, and 10 Gauss–Legendre quadrature nodes per spatial direction, eight nodes were selected, yielding a total of $N_e = 64$ nodes within the cell domain. This resolution provides a suitable compromise between computational complexity and spatial accuracy, with negligible differences in the electrothermal and degradation predictions compared with the 100-node model. The corresponding normalised volume weights account for the cylindrical geometry and are used for spatial integration according to (C.9).

For LFP cells operating under moderate-to-high currents, irreversible losses generally dominate the generated heat [34]. Accordingly, the reversible heat contribution is neglected in the present study, and the spatially distributed volumetric heat generation in (C.33) is approximated as

$$q(r,z,t) \approx q_{\mathrm{irr}}(r,z,t). \tag{C.47}$$

To investigate the influence of electrical loading on spatial electrothermal non-uniformity and degradation, both cells are subjected to constant-current-discharge–rest–charge–rest sequences at discharge rates of 1C, 2C, and 3C, while the charging rate is maintained at 1C. For each discharge rate, the operating sequence is repeated ten times to accumulate degradation arising from both discharging and charging.

The degradation model parameters are adopted from [26], where they were identified for a 3 Ah LFP/graphite cell under the reference conditions $T_{\mathrm{ref}} = 298.15$ K, $I_{\mathrm{ch,ref}} = 1$C, and $Z_{\mathrm{ref}} = 0.82$. Since the current sensitivity of the degradation model is expressed in terms of C-rate (see (C.75c) and (C.75d)), the reference capacity and charging current are set to $C_{\mathrm{ref}} = 2.3$ Ah and $I_{\mathrm{ch,ref}} = 2.3$ A, respectively. The currents supplied to the degradation model are scaled according to (C.20) and (C.26), ensuring equivalent normalised

charge throughputs under equivalent C-rates for both cells. The remaining fitted temperature, SoC, and current sensitivity parameters are retained from [26], assuming comparable relative capacity loss across cell formats under equivalent local operating conditions and C-rates. An independent reparameterisation of the degradation model for each cell is beyond the scope of this work.

For the coolant, a 50:50 ethylene-glycol–water mixture is assumed, with thermophysical properties evaluated at 20 °C, namely $\rho_c = 1069$ kgm^{-3}, $c_{p,c} = 3323$ Jkg^{-1}K^{-1}, and a coolant thermal conductivity of 0.3892 Wm^{-1}K^{-1} [38]. The total available coolant flow rate, $\dot{\mathcal{V}}_{\max} = 5.63 \times 10^{-6}$ m^3 s^{-1} ($\approx$ 0.34 Lmin^{-1}) with the convection coefficients for the lateral surface, top and bottom channels set to $h_s \approx h_t \approx h_b \approx 480$ Wm^{-2}K^{-1} [16]. Five representative cooling configurations are considered: lateral surface cooling (SC), bottom-tab cooling (bTC), combined bottom-tab and surface cooling (bTSC), combined bottom- and top-tab cooling (btTC), and equal-split cooling (ES). Their corresponding valve duty-cycle allocations are

$$\begin{aligned} &\text{SC} : \{1, 0, 0\}, \quad \text{bTC} : \{0, 0, 1\}, \quad \text{bTSC} : \{0.5, 0, 0.5\}, \\ &\text{btTC} : \{0, 0.5, 0.5\}, \quad \text{ES} : \{1/3, 1/3, 1/3\}, \end{aligned} \tag{C.48}$$

where each allocation represents $\{u_{\text{mod},s}, u_{\text{mod},t}, u_{\text{mod},b}\}$. The total available coolant flow rate is identical for the considered cells and cooling configurations.

For the degradation analysis, the thermal model employs $N = M = 5$ basis functions per direction, resulting in 25 thermal states. Unless otherwise stated, the initial cell temperature distribution $T(r, z, 0)$, initial coolant mean temperatures $X_{cl}(0)$, and coolant inlet temperatures $T_{cl,\text{in}}$ are set to 30 °C. All simulations are conducted in MATLAB R2025b on a MacBook Pro with a 2.6 GHz 6-core Intel Core i7 processor and 16 GB RAM, with a time step of 1 s for the battery–coolant states propagation.

Three degradation evaluations are considered: distributed degradation $Q_{\text{loss}}^{\text{dist}}$, same-history lumped degradation $Q_{\text{loss}}^{\text{lump,sh}}$, and independently simulated lumped degradation $Q_{\text{loss}}^{\text{lump,ind}}$. Their respective capacity loss rates are ex-

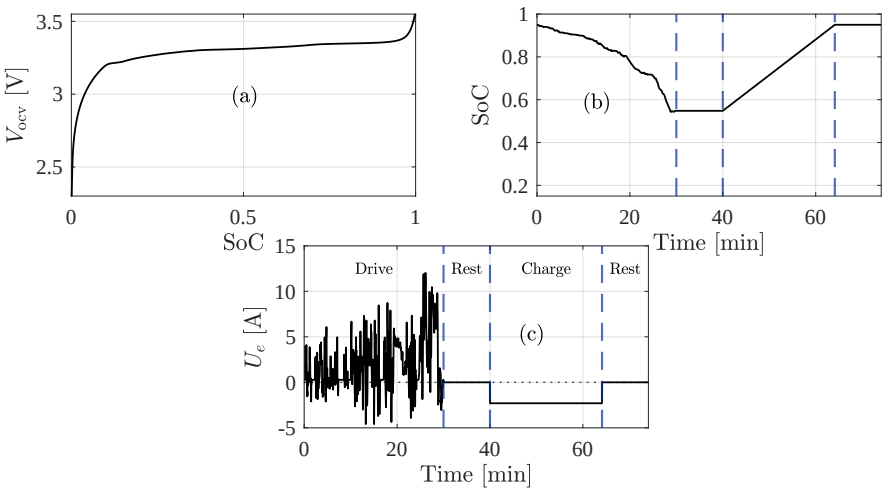

Figure 3: Electrical operating variables. (a) V_{ocv}–SoC relationship of the considered LFP cells. (b) SoC evolution over one drive–rest–charge–rest operating sequence. (c) Corresponding battery current profile, comprising the mWLTP driving phase, rest periods, and constant-current charging phase.

pressed as

$$\dot{Q}_{\mathrm{loss}}^{\mathrm{dist}} = \sum_{i=1}^{N_e} w_i f_Q\big(T_i, Z_i, \tilde{I}_{\ell,i}, \varphi_{\mathrm{tot},i}, \varphi_{\mathrm{ch},i}, \tau\big), \tag{C.49a}$$

$$\dot{Q}_{\mathrm{loss}}^{\mathrm{lump,sh}} = f_Q\big(T_v, Z_v, \tilde{I}_b, \varphi_{\mathrm{tot}}, \varphi_{\mathrm{ch}}, \tau\big), \tag{C.49b}$$

$$\dot{Q}_{\mathrm{loss}}^{\mathrm{lump,ind}} = f_Q\big(T_v^{\mathrm{ind}}, Z^{\mathrm{ind}}, \tilde{I}_b^{\mathrm{ind}}, \varphi_{\mathrm{tot}}^{\mathrm{ind}}, \varphi_{\mathrm{ch}}^{\mathrm{ind}}, \tau\big), \tag{C.49c}$$

where $f_Q(\cdot)$ denotes the combined capacity loss rate of the four ageing mechanisms, and Z_v is the volume-averaged SoC. The first two evaluations share the same distributed electrothermal operating history but differ in how the degradation rates are evaluated, i.e. either locally or using cell-averaged quantities. The third employs an independently simulated lumped electrical model with spatially uniform heat generation, while retaining the same thermal and coolant models.

3.2 Influence of electrothermal non-uniformity on degradation

Figs. 4–5 present the resulting model evaluation. As shown in Fig. 4(a), the distributed degradation model predicts approximately 0.42–0.46% accumulated capacity loss after ten operating sequences across the considered cell geometries, discharge rates, and cooling configurations. For most configurations, the accumulated capacity loss decreases with increasing discharge rate. This can partly be attributed to the shorter operating duration at higher discharge rates, which reduces the contribution of calendar ageing for the same number of operating sequences. At a given discharge rate, the differences in accumulated capacity loss among the cooling configurations are generally small under the investigated operating conditions. A notable exception is the 45 Ah cell under bTC at 3C, for which the accumulated capacity loss reaches approximately 0.44%, compared with approximately 0.42% under the other cooling configurations.

Fig. 4(b) compares the distributed and lumped degradation predictions, where

$$\Delta Q_{\mathrm{loss}}^{\mathrm{sh}} = Q_{\mathrm{loss}}^{\mathrm{dist}} - Q_{\mathrm{loss}}^{\mathrm{lump,sh}}, \tag{C.50a}$$

$$\Delta Q_{\mathrm{loss}}^{\mathrm{ind}} = Q_{\mathrm{loss}}^{\mathrm{dist}} - Q_{\mathrm{loss}}^{\mathrm{lump,ind}}. \tag{C.50b}$$

Despite sharing the same electrothermal operating history, differences arise between the distributed and same-history lumped degradation predictions due to the nonlinear dependence of the degradation rates on the local operating conditions. In general,

$$\sum_{i=1}^{N_e} w_i f_Q(T_i, \dots) \neq f_Q\left(\sum_{i=1}^{N_e} w_i T_i, \dots\right). \tag{C.51}$$

As observed in Fig. 4(b)(i), the lumped degradation evaluation can either overestimate or underestimate the distributed capacity loss, with the discrepancies generally becoming more pronounced at higher discharge rates. Nevertheless, the maximum absolute difference remains below approximately 4.3×10^{-4} percentage points (pp) for the investigated operating conditions. Larger discrepancies are observed when the distributed model is compared with the independently simulated lumped model in Fig. 4(b)(ii), particularly for the 45 Ah cell at higher discharge rates. The independently simulated

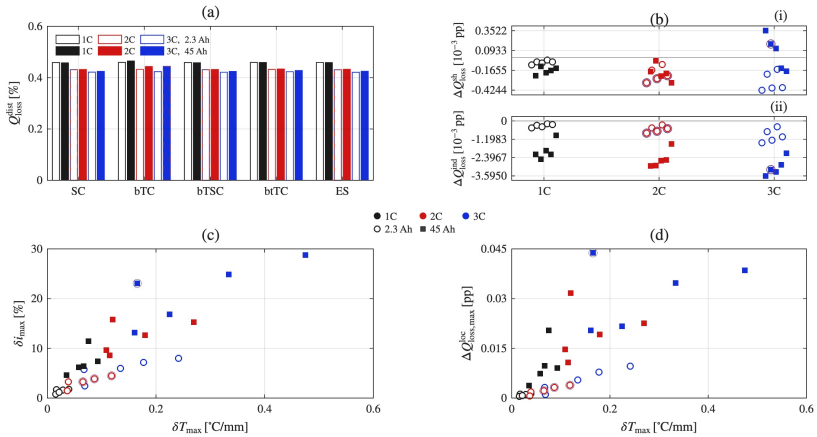

Figure 4: Influence of cell geometry, discharge rate, and cooling configuration on electrothermal non-uniformity and degradation after ten operating sequences. (a) Distributed capacity loss. (b) Differences in capacity loss between the distributed and (i) same-history lumped and (ii) independently simulated lumped degradation evaluations. (c) Maximum relative local current density deviation versus maximum RMS thermal gradient. (d) Maximum spatial variation in local capacity loss versus maximum RMS thermal gradient.

lumped model consistently predicts greater capacity loss, with the maximum difference reaching approximately 3.6×10^{-3} pp. The considerably larger discrepancies compared with the same-history evaluation indicate that spatial non-uniformity influences degradation not only through the nonlinear evaluation of the degradation rates but also through its influence on the underlying electrothermal dynamics.

To further investigate the influence of thermal non-uniformity on current redistribution and localised ageing, the maximum RMS thermal gradient, maximum relative local current density deviation, and maximum spatial variation

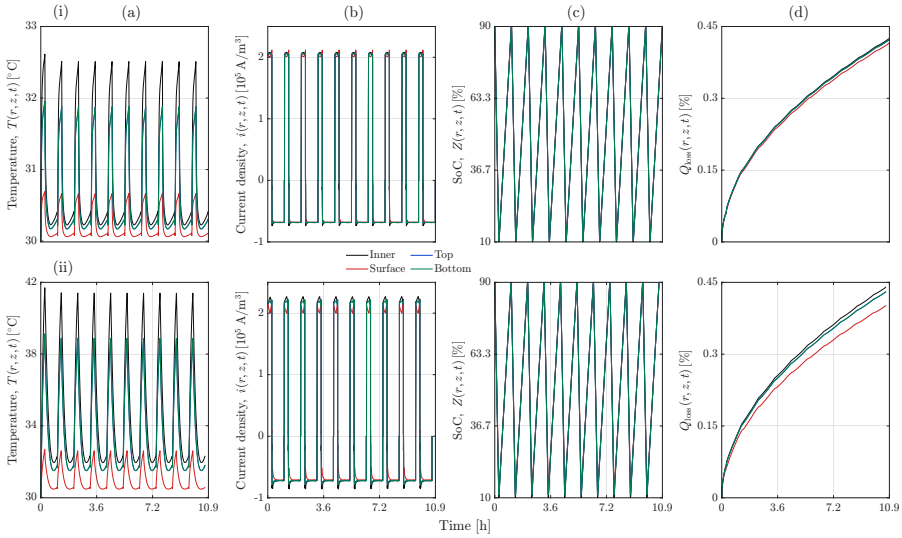

Figure 5: Evolution of local electrothermal variables and capacity loss during ten discharge–rest–charge–rest sequences under lateral surface cooling at a discharge rate of 3C for (i) the 2.3 Ah and (ii) the 45 Ah cells. (a) Temperature, (b) volumetric current density, (c) SoC, and (d) accumulated capacity loss at four representative locations within the cell.

in local capacity loss are defined as

$$\delta T_{\max} = \max_{t} \delta T(t), \tag{C.52a}$$

$$\delta i_{\max} = \max_{t,i} \left| \frac{I_{\ell,i}(t) - I_b(t)}{I_b(t)} \right| \times 100, \tag{C.52b}$$

$$\Delta Q^{\text{loc}}_{\text{loss},\max} = \max_{t} \left(\max_{i} Q_{\text{loss},i}(t) - \min_{i} Q_{\text{loss},i}(t) \right). \tag{C.52c}$$

The relative current density deviation is evaluated during non-zero applied current operation. Fig. 4(c) reveals a general increase in local current density deviations with increasing thermal gradients. This effect is more pronounced for the 45 Ah cell due to its larger characteristic heat conduction distances. At 3C, the maximum RMS thermal gradient reaches approximately 0.24 °C/mm for the 2.3 Ah cell and 0.48 °C/mm for the 45 Ah cell, with corresponding maximum local current density deviations of approximately 8% and 29%, re-

spectively. These results demonstrate that the assumption of spatially uniform current distribution becomes increasingly restrictive for large-format cells under high electrical loading.

A similar trend is observed in Fig. 4(d), where larger thermal gradients are generally associated with greater spatial variations in capacity loss. The maximum local capacity loss difference increases from approximately 0.01 pp for the 2.3 Ah cell to 0.044 pp for the 45 Ah cell. Interestingly, the largest local capacity loss difference does not coincide with the largest thermal gradient, indicating that localised ageing is governed by the combined operating history of temperature, SoC, and current rather than the maximum thermal gradient alone. Although the differences between the cell-level distributed and lumped degradation predictions remain relatively small, the spatial variation in local capacity loss is more than an order of magnitude greater. Thus, cell-level capacity loss alone may conceal considerable localised ageing.

To illustrate the underlying electrothermal and degradation dynamics, Fig. 5 shows the evolution of temperature, current density, SoC, and capacity loss at four representative locations within both cells under SC at 3C. The selected locations are near the cell core, lateral surface, top, and bottom. As observed in Fig. 5(a), the larger cell exhibits substantially greater temperature non-uniformity, with peak inner temperatures of approximately 42 °C compared with 33 °C for the smaller cell. While the temperature difference between the inner and lateral surface regions remains approximately 2 °C for the 2.3 Ah cell, it approaches 9 °C for the 45 Ah cell during discharging. The resulting spatial variations in the temperature-dependent electrical parameters induce non-uniform current distributions, as shown in Fig. 5(b). Despite the relatively small differences in the local SoC trajectories in Fig. 5(c), the combined influence of local temperature, SoC, and current gives rise to different ageing rates.

Consequently, the local capacity loss trajectories in Fig. 5(d) progressively diverge, particularly for the 45 Ah cell, where the accumulated capacity loss at the inner and lateral surface regions reaches approximately 0.44% and 0.40%, respectively. The preferential cooling of the lateral surface therefore results in less degradation near the cooled interface than in the inner region under the considered operating conditions.

The preceding results demonstrate the importance of accounting for coupled spatial electrothermal dynamics in battery degradation modelling, par-

ticularly for large-format cells under high electrical loading. The observed variations in local degradation arise from the combined influence of temperature, SoC, and current distributions, which cannot be fully characterised by lumped cell-level quantities alone. These findings motivate the explicit consideration of distributed battery degradation in the design of battery control strategies

The relative sensitivities observed above are, however, specific to the selected T_v, SoC, and charge-current condition, since the individual degradation mechanisms respond differently to each of these operating variables. The following subsection systematically examines their sensitivities to temperature, SoC, and current to identify the operating regions in which each mechanism becomes dominant.

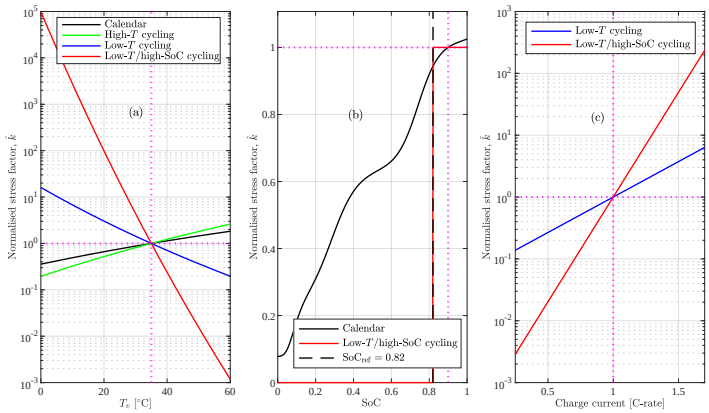

Figure 6: Sensitivity of the degradation stress factors to the principal operating variables, normalised with respect to $T_v = 35$ °C, $Z = 0.9$, and a charging rate of 1C. (a) Temperature sensitivity at $Z = 0.9$ and 1C. (b) SoC sensitivity at $T_v = 35$ °C and 1C, with the dashed black line indicating the high-SoC activation threshold $Z_{\mathrm{ref}} = 0.82$. (c) Charge-current sensitivity at $T_v = 35$ °C and $Z = 0.9$. The magenta dotted lines indicate the normalisation conditions at $\hat{k} = 1$.

3.3 Sensitivity of degradation mechanisms

The four degradation mechanisms respond differently to temperature, SoC, and charge current. To isolate the influence of these operating variables, the stress factors are evaluated using lumped operating quantities. Since the stress factors differ in magnitude and units, their sensitivities are compared using the normalised stress factor

$$\hat{k} = \frac{k(T_v, Z, I_{\mathrm{ch}})}{k_{\mathrm{norm}}}, \tag{C.53}$$

where k_{norm} denotes the corresponding stress factor evaluated at the normalisation condition $T_v = 35$ °C, $Z = 0.9$, and a charging rate of 1C. Accordingly, $\hat{k} = 1$ indicates the same ageing stress as at the normalisation condition, whereas values above or below unity indicate an increase or decrease, respectively. This normalisation enables comparison of the relative sensitivities of the individual mechanisms rather than their absolute contributions to capacity loss.

Fig. 6(a) shows the temperature sensitivity at fixed $Z = 0.9$ and a charging rate of 1C. The calendar and high-temperature cycling stress factors increase with temperature, whereas both low-temperature mechanisms exhibit the opposite trend. Relative to their values at 35 °C, the calendar and high-temperature cycling stress factors increase to approximately 1.83 and 2.61, respectively, at 60 °C. At 0 °C, the low-temperature cycling stress factor increases by a factor of approximately 16, while that of the low-temperature/high-SoC mechanism increases by nearly five orders of magnitude. These opposing trends demonstrate that reducing cell temperature does not uniformly reduce ageing stress, as it suppresses the high-temperature mechanisms while increasing the susceptibility to low-temperature degradation.

The SoC sensitivity is shown in Fig. 6(b) at $T_v = 35$ °C and a charging rate of 1C. The high-temperature and low-temperature cycling stress factors have no explicit SoC dependence and therefore remain unchanged. Calendar ageing, in contrast, increases with increasing SoC. The low-temperature/high-SoC mechanism exhibits a threshold-dependent behaviour, remaining inactive below $Z_{\mathrm{ref}} = 0.82$ and becoming active above this threshold, where its stress factor remains independent of further increases in SoC.

Fig. 6(c) illustrates the charge-current sensitivity at $T_v = 35$ °C and $Z = 0.9$. The calendar and high-temperature cycling stress factors have no explicit dependence on charge current, whereas both low-temperature mechanisms in-

crease exponentially (appears linear in the logarithmic ordinate) with charging rate. The low-temperature/high-SoC mechanism exhibits considerably greater sensitivity, with $\hat{k}$ reaching approximately 49.4 and 235 at charging rates of 1.5C and 1.7C, respectively. By comparison, the corresponding values for the low-temperature cycling mechanism are 3.74 and 6.35. Elevated charging rates can therefore strongly amplify low-temperature ageing stress, particularly when the high-SoC mechanism is active.

Taken together with the preceding analysis of spatial electrothermal nonuniformity, these results demonstrate the complex and competing effects of temperature, SoC, and current on battery degradation. The opposing temperature sensitivities of the ageing mechanisms indicate that a fixed thermal target may not be suitable across different operating conditions. The cooling action should therefore account jointly for the spatial temperature distribution, SoC, and charge current. These coupled and competing ageing effects motivate the formulation of the degradation-aware control problem presented in the following section.

4 MPC formulation

Building on the battery–coolant models developed in Section 2 and the degradation analysis in Section 3, this section formulates the degradation-aware control problem for the ITSC system. The control problem is first introduced, followed by the objective function, constraints, and controller implementation.

The distributed degradation model is employed throughout the control formulation. Accordingly, for notational simplicity, the superscript dist is omitted hereafter, such that X_{deg}, $\mathcal{Y}_{\mathrm{deg}}$, and Q_{loss} refer to the distributed degradation state vector, degradation output, and cell-level capacity loss, respectively.

4.1 Control problem statement

The control problem is formulated as the following optimal coolant allocation problem: *The degradation-aware optimal coolant allocation problem (DOCAP) is to determine, over a finite prediction horizon, the allocation of the available coolant flow among the lateral surface, top-tab, and bottom-tab cooling channels such that battery degradation is minimised while satisfying battery safety and actuator constraints.*

Given a battery current profile I_b, the DOCAP determines the optimal distribution and utilisation of the available coolant flow to regulate the spatial temperature distribution within the battery. The battery current is treated as a known exogenous input over the prediction horizon. The analysis in Section 3 demonstrates that battery degradation is governed by the coupled spatial distributions of temperature, SoC, and current, with the individual ageing mechanisms exhibiting different and competing sensitivities to these operating variables. Owing to the multivariable and history-dependent nature of battery degradation, the DOCAP explicitly accounts for the predicted distributed degradation in determining the optimal coolant allocation.

The DOCAP exhibits several properties that motivate a predictive control framework. First, the coupled electrical, thermal, degradation, and coolant dynamics are nonlinear, and the three complementary cooling channels provide multiple control degrees of freedom. Second, battery safety and actuator constraints must be satisfied throughout operation. Third, battery degradation depends not only on the instantaneous operating conditions but also on their evolution over time and the accumulated charge throughputs. Finally, prediction of the battery current enables the controller to anticipate changes in heat generation and their subsequent effects on the spatial temperature distribution and degradation. These features make MPC well suited for solving the DOCAP.

Previous studies have used MPC for battery thermal management to enhance temperature control, reduce energy consumption, and extend battery life [27], [39]. For real-time EV battery thermal control, RTI-MPC has also been employed as a computationally efficient alternative to conventional nonlinear MPC (NMPC) [31], [40], [41]. These studies, however, generally consider a limited set of cooling inputs such as coolant flow rate, coolant temperature, and/or auxiliary cooling power and do not address the dynamic allocation of coolant among multiple complementary cooling channels. Our previous work [16] addressed this limitation by using RTI-MPC to dynamically distribute the available coolant among the lateral surface and tab cooling channels of the ITSC system. The control objective, however, was formulated in terms of temperature regulation and thermal uniformity, without explicitly accounting for battery degradation in the optimisation. In the present work, this framework is extended to formulate and solve the DOCAP, with predicted distributed battery degradation explicitly incorporated into the control objec-

tive.

4.2 Cost function

At each sampling instant k, the MPC solves the DOCAP by determining the sequence of valve duty-cycles $\mathbf{u}_{\mathrm{mod},k}$ over a finite prediction horizon that minimises

$$\begin{aligned} J(k) := \lambda_1 \left[\hat{\mathcal{Y}}_{\mathrm{deg}}(k+N_p) - \hat{\mathcal{Y}}_{\mathrm{deg}}(k)\right] \\ + \sum_{i=0}^{N_p-1} \left[\lambda_2 \left(1^T u_{\mathrm{mod}}(k+i)\right)^2 + \lambda_3 \left\|\Delta u_{\mathrm{mod}}(k+i)\right\|_{R_\Delta}^2\right], \end{aligned} \tag{C.54}$$

where N_p is the prediction horizon, $\|x\|_D^2 := x^\top D x$, $1^T = [1,\ 1,\ 1]$, and $\lambda_1, \lambda_2 \geq 0$ and $\lambda_3 > 0$ are trade-off weights. The weighting matrix satisfies $R_\Delta \succ 0$. The three terms in (C.54) are described below.

Degradation penalty

The first term penalises the predicted increase in accumulated capacity loss over the prediction horizon. Specifically, $\hat{\mathcal{Y}}_{\mathrm{deg}}(k)$ denotes the accumulated capacity loss at the current sampling instant, while $\hat{\mathcal{Y}}_{\mathrm{deg}}(k+N_p)$ is the corresponding predicted value at the end of the horizon. Their difference represents the additional degradation expected to occur over the prediction horizon.

Through the dependence of the degradation rates on the spatial distributions of temperature, SoC, and current, as well as the accumulated local charge throughputs, this term directs the coolant allocation towards minimising predicted battery degradation.

Coolant utilisation penalty

The second term penalises the total fraction of the available coolant flow utilised by the ITSC system. From the valve actuation model in (C.37), the total coolant utilisation is given by

$$\frac{\dot{\mathcal{V}}_s(k+i) + \dot{\mathcal{V}}_t(k+i) + \dot{\mathcal{V}}_b(k+i)}{\dot{\mathcal{V}}_{\max}} = 1^T u_{\mathrm{mod}}(k+i), \tag{C.55}$$

such that the penalty discourages unnecessary coolant use without favouring a particular distribution of the given total coolant flow among the three

channels.

Control variation penalty

The third term penalises changes in the duty-cycle commands,

$$\Delta u_{\mathrm{mod}}(k+i) = u_{\mathrm{mod}}(k+i) - u_{\mathrm{mod}}(k+i-1). \tag{C.56}$$

This regularises the control action, discouraging abrupt changes in coolant routing and promoting smooth valve duty-cycle commands.

4.3 Constraints

The objective in (C.54) is minimised subject to the discrete-time battery–coolant dynamics and the battery safety and actuator constraints,

$$\hat{X}(k) = X(k), \tag{C.57a}$$
$$\hat{X}(k+i+1) = F_d\big(\hat{X}(k+i), U(k+i), W(k+i), k+i\big), \tag{C.57b}$$
$$\hat{Y}(k+i) = H\big(\hat{X}(k+i), U(k+i), W(k+i), k+i\big), \tag{C.57c}$$
$$\hat{\mathcal{Y}}_{\mathrm{deg}}(k+i) = \sum_{j=1}^{N_e} w_j \hat{Q}_{\mathrm{loss},j}(k+i), \tag{C.57d}$$
$$\hat{T}_{\max}(k+i+1) \le T_{\max}^{\lim}, \tag{C.57e}$$
$$0 \le u_{\mathrm{mod}}(k+i) \le 1, \tag{C.57f}$$
$$u_{\mathrm{mod},s}(k+i) + u_{\mathrm{mod},t}(k+i) + u_{\mathrm{mod},b}(k+i) \le 1, \tag{C.57g}$$

$\forall i \in \{0, 1, \ldots, N_p - 1\}$, where $F_d(\cdot)$ denotes the discrete-time nonlinear state-transition mapping obtained from the coupled continuous-time dynamics in (C.45). The notation $\hat{\cdot}$ denotes the predicted quantities, with the initial prediction state given by the state available at sampling instant k. The predicted degradation output is obtained by volume-weighted averaging of the accumulated local capacity losses, including at the terminal prediction instant $k+N_p$. The predicted maximum cell temperature is defined as

$$\hat{T}_{\max}(k+i) = \max_{j \in \{1,\ldots,N_e\}} \hat{T}_j(k+i), \tag{C.58}$$

where $\hat{T}_j$ denotes the reconstructed temperature at the jth quadrature node. Constraint (C.57e) limits the maximum predicted cell temperature to the

maximum allowable temperature $T_{\max}^{\lim}$ throughout the prediction horizon. Finally, (C.57f) bounds the individual valve duty-cycles, while (C.57g) ensures that the combined channel flow demand does not exceed the total coolant flow available from the pump.

4.4 Controller implementation

The DOCAP in (C.54)–(C.57) is implemented in a receding-horizon fashion. At each sampling instant k, the controller computes the optimal duty-cycle sequence $\{u_{\mathrm{mod}}^{*}(k+i)\}_{i=0}^{N_p-1}$. Only the first element of this sequence is then applied to the valve actuation system, after which the optimisation horizon is shifted and the problem is solved again using the updated state information.

The DOCAP is nonlinear owing to the coupled electrical, thermal, degradation, and coolant dynamics. The nonlinearities arise primarily from the operating-condition-dependent electrical parameters, spatial current redistribution, nonlinear degradation stress factors, and bilinear coolant dynamics. An NMPC formulation is thus implemented, where the nonlinear prediction dynamics are retained in their original forms. The resulting nonlinear programme is solved to convergence at each sampling instant and serves as a numerical reference for evaluating the RTI-MPC implementation. Since repeatedly solving the nonlinear optimisation problem to convergence can impose a considerable online computational burden, the proposed DOCAP controller is implemented using the RTI-MPC scheme.

RTI-MPC for the DOCAP

The RTI scheme comprises a warm-start step, a preparation phase, and a feedback phase at every sampling instant k.

Warm-start step Let $u_{\mathrm{mod}}^{*,\mathrm{prev}}(i)$, $i = 0, \ldots, N_p - 1$, represent the optimal duty-cycle sequence computed at sampling instant $k-1$. At k, this sequence is shifted by one step to initialise the nominal control trajectory,

$$\tilde{u}_{\mathrm{mod}}(k+i) = u_{\mathrm{mod}}^{*,\mathrm{prev}}(i+1), \qquad i = 0, \ldots, N_p - 2, \tag{C.59}$$

and the last duty-cycle command from the previous solution is repeated to complete the shifted horizon, i.e.,

$$\tilde{u}_{\mathrm{mod}}(k+N_p-1) = u_{\mathrm{mod}}^{*,\mathrm{prev}}(N_p-1). \tag{C.60}$$

The corresponding nominal volumetric flow trajectory $\tilde{U}(k+i)$ is obtained from (C.37). Using this trajectory together with the known battery current and coolant inlet temperature predictions, the nonlinear battery–coolant model in (C.57b) is propagated over the prediction horizon to obtain the nominal state trajectory

$$\tilde{X}(k+i+1) = F_d\big(\tilde{\xi}_i\big), \tag{C.61}$$

for $i = 0, \ldots, N_p - 1$, where

$$\tilde{\xi}_i = \big(\tilde{X}(k+i), \tilde{U}(k+i), W(k+i), k+i\big)\,.$$

The corresponding nominal output trajectory is obtained from (C.57c).

Preparation phase The nonlinear prediction dynamics in (C.57b) are locally linearised about the nominal state and control trajectories obtained from the warm-start step. Define the perturbation variables

$$\delta\hat{X}(k+i) = \hat{X}(k+i) - \tilde{X}(k+i), \tag{C.62a}$$
$$\delta U(k+i) = U(k+i) - \tilde{U}(k+i), \tag{C.62b}$$
$$\delta u_{\mathrm{mod}}(k+i) = u_{\mathrm{mod}}(k+i) - \tilde{u}_{\mathrm{mod}}(k+i). \tag{C.62c}$$

From (C.37), the perturbations in the volumetric flow rates and valve duty-cycles are related through

$$\delta U(k+i) = \dot{\mathcal{V}}_{\max}\delta u_{\mathrm{mod}}(k+i). \tag{C.63}$$

A first-order Taylor expansion of (C.57b) about the nominal trajectory (C.61) yields

$$\delta\hat{X}(k+i+1) = A_i\delta\hat{X}(k+i) + B_i\delta U(k+i), \tag{C.64}$$

where the Jacobians are

$$A_i = \left.\frac{\partial F_d}{\partial X}\right|_{\tilde{\xi}_i}, \qquad B_i = \left.\frac{\partial F_d}{\partial U}\right|_{\tilde{\xi}_i}. \tag{C.65}$$

Similarly, the output mapping in (C.57c) is locally linearised as

$$\delta\hat{Y}(k+i) = C_i\delta\hat{X}(k+i) + D_i\delta U(k+i), \tag{C.66}$$

where

$$C_i = \left.\frac{\partial H}{\partial X}\right|_{\tilde{\xi}_i}, \qquad D_i = \left.\frac{\partial H}{\partial U}\right|_{\tilde{\xi}_i}. \tag{C.67}$$

The known exogenous inputs are held fixed during the linearisation. The terminal degradation output appearing in the objective is linear in the degradation state according to

$$\hat{\mathcal{Y}}_{\mathrm{deg}}(k+N_p) = \mathcal{C}_{\mathrm{deg}}^{\mathrm{agg}} \hat{X}_{\mathrm{deg}}(k+N_p), \tag{C.68}$$

where $\mathcal{C}_{\mathrm{deg}}^{\mathrm{agg}}$ denotes the output matrix incorporating the normalised volume weights of the quadrature nodes and $\hat{X}_{\mathrm{deg}} = \tilde{X}_{\mathrm{deg}} + \delta\hat{X}_{\mathrm{deg}}$. Thus, the degradation penalty is affine in the terminal degradation state perturbation and requires no additional linearisation. The maximum temperature constraint in (C.57e) is equivalently imposed on the predicted temperatures. The corresponding local approximation is expressed as

$$\tilde{T}_j(k+i+1) + \delta\hat{T}_j(k+i+1) \le T_{\max}^{\lim}, \tag{C.69}$$

for $j = 1, \ldots, N_e$, where $\delta\hat{T}_j$ denotes the perturbation of the reconstructed temperature at the jth quadrature node. Substituting the linearised prediction dynamics and outputs into the original DOCAP yields a local quadratic programme (QP) at each sampling instant k,

$$\begin{aligned} \min_{\delta u_{\mathrm{mod},k}} \quad & J_{\mathrm{QP}}(k) \\ \text{s.t.} \quad & \text{(C.63), (C.64), (C.66),} \\ & \text{(C.69),} \\ & 0 \le \tilde{u}_{\mathrm{mod}}(k+i) + \delta u_{\mathrm{mod}}(k+i) \le 1, \\ & 1^\top \left(\tilde{u}_{\mathrm{mod}}(k+i) + \delta u_{\mathrm{mod}}(k+i)\right) \le 1, \\ & i = 0, \ldots, N_p - 1. \end{aligned} \tag{C.70}$$

Here, $J_{\mathrm{QP}}(k)$ is obtained from (C.54) using the affine degradation output in (C.68) and the corrected duty-cycle sequence

$$u_{\mathrm{mod}}(k+i) = \tilde{u}_{\mathrm{mod}}(k+i) + \delta u_{\mathrm{mod}}(k+i). \tag{C.71}$$

Feedback phase Once the updated battery–coolant state information at sampling instant k becomes available, the initial state perturbation is determined as $\delta\hat{X}(k) = X(k) - \tilde{X}(k)$. The prepared QP in (C.70) is then solved

for the optimal correction sequence $\delta u^*_{\mathrm{mod},k}$. The first corrected duty-cycle command,

$$u^*_{\mathrm{mod}}(k) = \tilde{u}_{\mathrm{mod}}(k) + \delta u^*_{\mathrm{mod}}(k), \tag{C.72}$$

is applied to the ITSC system through (C.37). The remaining elements of the corrected sequence are stored and shifted to initialise the warm-start trajectory at sampling instant $k + 1$.

The RTI-MPC performs one local linearisation of the nonlinear battery–coolant prediction dynamics and solves one QP per sampling instant, rather than repeatedly solving the nonlinear DOCAP to convergence. This reduces the online computational burden while retaining the receding-horizon and constraint-handling capabilities of the original DOCAP.

5 Results and discussion

5.1 Controller specification and benchmarks

The DOCAP is evaluated using the ITSC system and the 2.3 Ah and 45 Ah LFP cylindrical cells specified in Section 3.1. Unless otherwise stated, the battery–coolant parameters, actuator specifications, and total available coolant-flow capacity are retained from that section.

To generate a realistic and dynamically varying current load, the driving current is obtained from the Worldwide Harmonised Light Vehicle Test Procedure (WLTP), which represents urban, suburban, and highway driving conditions [42]. The original WLTP current trajectory is scaled and limited to a maximum discharging C-rate of 3C and regenerative charging C-rate of 1C, yielding a modified WLTP (mWLTP) profile with a mean absolute C-rate of approximately 2C. To represent repeated vehicle operation, each operating sequence comprises an mWLTP driving phase, a 10 min rest period, a constant-current charging phase, and a further 10 min rest period. The cell is initialised at a volume-averaged SoC of $Z_v = 0.90$ and driven until the end of the drive cycle, with $Z_v = 0.20$. Following the first rest period, the cell is charged at 1C, corresponding to $I_b = -2.3$ A and $I_b = -45$ A for the 2.3 Ah and 45 Ah cells, respectively, to restore the nominal SoC to 0.90. Each complete operating sequence lasts approximately 87.4 min and is repeated ten times to accumulate degradation arising from driving, rest, and charging. A representative SoC evolution and corresponding operating sequence

are illustrated in Fig. 3(b) and (c), respectively.

The initial cell temperature distribution $T(r, z, 0)$, coolant mean temperatures $X_{cl}(0)$, and coolant inlet temperatures $T_{cl,\text{in}}$ are all set to 35 °C. For the thermal prediction model, $M = N = 2$ (four states) basis functions are used for both the RTI-MPC and NMPC implementations, whereas the plant model employs $M = N = 10$ (100 states).

The MPC sampling time is 1 s, with a prediction horizon of $N_p = 100$ steps. The optimisation problems are implemented in YALMIP [43], with Gurobi used for the RTI-MPC and `fmincon` for the NMPC. The maximum permissible cell temperature is set to $T_{\max}^{\lim} = 50$ °C. The objective weights are selected as $\lambda_1 = 2 \times 10^6$, $\lambda_2 = 0.02$, and $\lambda_3 = 0.1$, with $R_\Delta = I_3$. The degradation weight λ_1 accounts for the comparatively small numerical magnitude of the predicted capacity loss increment over the prediction horizon, while λ_2 and λ_3 penalise coolant utilisation and intersample control variation, respectively. The same prediction horizon, constraints, and objective weights are employed for the RTI-MPC and NMPC implementations.

The degradation-aware ITSC strategy is evaluated against the five baseline cooling configurations defined in Section 3.1, namely SC, bTC, bTSC, btTC, and ES. With the exception of ES, the proposed ITSC system and all benchmark configurations are controlled using the RTI-MPC. The same prediction models, objective function, prediction horizon, thermal constraints, controller weights, and maximum available coolant-flow capacity are employed. They differ only in the cooling channels available to the controller. For SC, only $u_{\text{mod},s}$ is available, whereas bTC permits only $u_{\text{mod},b}$. The bTSC configuration permits $u_{\text{mod},b}$ and $u_{\text{mod},s}$, while btTC permits $u_{\text{mod},t}$ and $u_{\text{mod},b}$. All three duty-cycle inputs $u_{\text{mod},s}$, $u_{\text{mod},t}$, and $u_{\text{mod},b}$ are available in the proposed ITSC configuration. The ES benchmark also employs all three channels, but without online optimisation, with $u_{\text{mod},s} = u_{\text{mod},t} = u_{\text{mod},b} = \frac{1}{3}$, thereby continuously distributing the full available coolant flow equally among the three channels.

The primary ageing performance metric is the accumulated normalised capacity loss Q_{loss} obtained from the distributed degradation model. The contributions of the individual degradation mechanisms and the spatial distribution of capacity loss are also examined. Thermal behaviour is assessed using the volume-averaged temperature T_v, the root-mean-square spatial temperature gradient δT, and the core and surface temperatures, T_c and T_s, respectively.

The spatial variations in temperature, current density, SoC, and capacity loss are further examined to assess the local electrothermal and degradation responses. The coolant response is evaluated using the mean coolant temperatures and volumetric flow rates of the individual cooling channels. Coolant utilisation is assessed using the total duty-cycle $1^\top u_{\mathrm{mod}}$ and the cumulative total and channel-wise duty-cycles over the evaluation period. The latter represent the equivalent operating durations at the maximum available coolant-flow rate.

5.2 RTI-MPC vs. NMPC

Fig. 7 compares the closed-loop responses of the RTI-MPC and NMPC implementations for the 2.3 Ah cell over one complete drive–rest–charge–rest sequence. The accumulated capacity loss trajectories in Fig. 7(a) are nearly indistinguishable throughout the sequence, with a relative difference of only 0.00162% in the final accumulated capacity loss. Similarly, T_v, δT, T_c, and T_s in Fig. 7(b)–(d) show close agreement between the two implementations, even during the most thermally demanding portion of the driving cycle. The corresponding absolute errors in Q_{loss}, T_v, and δT are shown in Fig. 7(e). The maximum errors in T_v and δT are 0.0796 °C and 0.01324 °C/mm, respectively, while those in T_c and T_s are 0.0593 °C and 0.1456 °C, respectively. The close agreement in the degradation and thermal responses indicates that the local linearisation and single-update structure of the RTI-MPC closely reproduce the closed-loop behaviour of the NMPC under the considered operating conditions.

The computational benefit of the RTI implementation is evident from Fig. 7(f). The mean online execution time is 0.1 s for RTI-MPC, compared with 4.6 s for NMPC, corresponding to an average speed-up of 46 times. RTI-MPC satisfies the 1 s sampling interval for all of the control updates, whereas NMPC exceeds this interval in the entire sequence. Given its close agreement with the NMPC benchmark and substantially lower computational burden, RTI-MPC is employed for the evaluations in the remaining subsections.

5.3 ITSC vs. conventional cooling configurations

Fig. 8(a) compares the accumulated capacity loss under the different cooling configurations for the 2.3 Ah and 45 Ah cells over ten consecutive drive–rest–

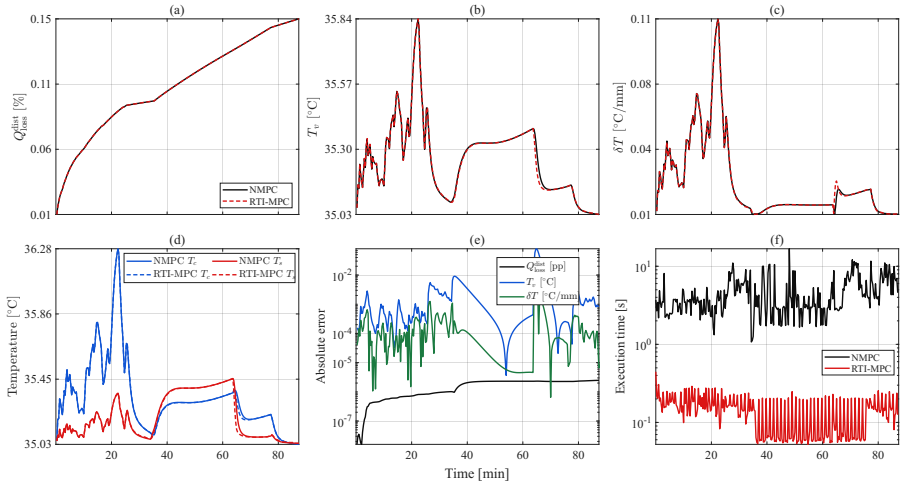

Figure 7: Closed-loop comparison of RTI-MPC and NMPC over one complete drive–rest–charge–rest sequence: (a) accumulated capacity loss, (b) volume-averaged cell temperature, (c) RMS spatial temperature gradient, (d) battery core and surface temperatures, (e) absolute differences between the RTI-MPC and NMPC responses, and (f) online execution time per controller update.

charge–rest sequences, corresponding to approximately 14.56 h of operation ($\approx$ 7.73 equivalent full cycles (EFCs)).

Among the feedback-controlled configurations, ITSC achieves the lowest total degradation for both cells, followed by bTSC, SC, and btTC, with bTC experiencing the largest capacity loss. Although the differences are relatively small for the 2.3 Ah cell, their systematic separation is visible in the enlarged final portion of the trajectories in Fig. 8(a-i). The influence of cooling configuration becomes considerably more pronounced for the 45 Ah cell in Fig. 8(a-ii). Compared with bTC, ITSC reduces the accumulated capacity loss by 1.51% and 7.98% for the 2.3 Ah and 45 Ah cells, respectively. ES achieves comparable degradation performance to ITSC for both cells but requires continuous full-flow cooling, corresponding to a cumulative coolant duty of 14.56 h. In contrast, ITSC requires only 2.19 h and 5.46 h for the 2.3 Ah and 45 Ah cells, respectively, corresponding to reductions in coolant utilisation of 85% and 62.5%.

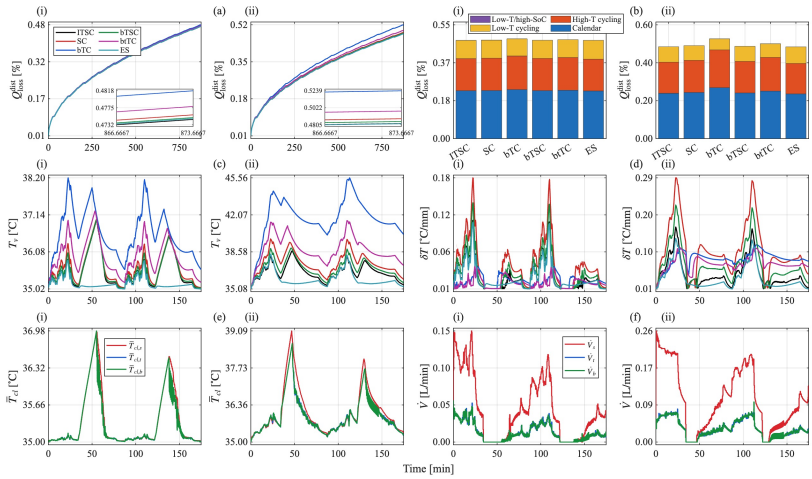

Figure 8: Comparison of the degradation-aware ITSC strategy with conventional cooling configurations for (i) the 2.3 Ah cell and (ii) the 45 Ah cell. (a) Accumulated capacity loss over ten operating sequences, with enlarged views of the final separation between the configurations. (b) Corresponding contributions of the four degradation mechanisms. (c) Volume-averaged cell temperature T_v. (d) RMS spatial temperature gradient δT. ITSC coolant response: (e) mean coolant temperatures and (f) volumetric flow rates of the individual cooling channels. The trajectories in (c)–(f) are shown for the first two operating sequences.

The ageing-mechanism decomposition in Fig. 8(b) explains the observed degradation behaviour. Calendar ageing Q_C provides the largest contribution for both cells, followed by high-temperature cycling ageing Q_H. The former remains active throughout the complete operating sequence, while its stress factor increases with both temperature and SoC. Moreover, its square-root dependence on elapsed time results in a larger incremental degradation rate during the early stages of the ageing trajectory. Q_H on the other hand, accumulates with the total charge throughput and is therefore active during both charging and discharging. The repeated high-current driving phases simultaneously generate substantial charge throughput and the largest temperature excursions, making this contribution significant. The low-temperature cycling

mechanism Q_L remains non-negligible despite the cells operating at conventionally high temperatures. This follows from the structure of the adopted degradation model, in which the mechanism is not activated by a hard temperature threshold. Rather, its stress factor decreases continuously with increasing temperature and increases with charge current, while the associated capacity loss accumulates with charge-only throughput. Accordingly, this mechanism remains active during both the 1C recharge and the regenerative-current portions of the driving profile. By contrast, the low-temperature/high-SoC mechanism Q_{LS} contributes negligibly to the accumulated capacity loss, even though the local SoC repeatedly exceeds its activation threshold of 0.82. This is because only charge throughput occurring above this threshold contributes to the mechanism. Furthermore, as shown in Section 3.3, this mechanism exhibits the strongest inverse-temperature dependence among the four ageing terms. At the prevailing 35–46 °C temperatures and limited charge current, its stress factor is strongly suppressed. A meaningful contribution from Q_{LS} requires high SoC, charging, and sufficiently low temperature to occur simultaneously for a sustained period.

The competing temperature dependence of the ageing mechanisms is evident when comparing ITSC and ES in Fig. 8(b). The lower cell temperatures maintained by ES reduce both calendar and high-temperature cycling ageing but increase the contribution of low-temperature cycling ageing. The opposite behaviour is particularly evident for bTC in the 45 Ah cell, where T_v reaches approximately 45.5 °C in Fig. 8(c-ii). Although the higher operating temperature suppresses low-temperature cycling ageing, this benefit is outweighed by the increased calendar and high-temperature cycling contributions. This behaviour is consistent with the mechanism-sensitivity analysis in Section 3.3 and illustrates that the degradation-optimal thermal condition depends on the combined contributions of the individual ageing mechanisms at the prevailing operating conditions.

The spatial temperature gradients in Fig. 8(d) differ substantially among the cooling configurations. SC produces the largest RMS temperature gradients, reaching approximately 0.176 °C/mm and 0.287 °C/mm for the 2.3 Ah and 45 Ah cells, respectively. The tab-only configurations maintain comparatively small temperature gradients by exploiting the favourable axial heat-conduction path. Nevertheless, both accumulate more degradation than SC due to their higher operating temperatures. Similarly, although ES main-

tains lower temperature gradients than ITSC during several portions of the operating sequence, their accumulated capacity losses are comparable. These observations indicate that minimising thermal non-uniformity does not necessarily minimise the spatially averaged capacity loss under the considered operating conditions.

The coolant temperature and flow responses of ITSC are shown in Fig. 8(e) and (f), respectively. The controller does not attempt to return the battery to its initial temperature of 35 °C whenever it heats up. Rather, coolant allocation is determined by the predicted degradation reduction relative to the associated coolant-utilisation and control-variation costs, while satisfying the imposed thermal constraint. During the thermally demanding portions of each sequence, the surface channel receives the largest coolant allocation. This follows from its effectiveness in extracting heat via the large surface area and reducing the volume-averaged cell temperature. The top and bottom channels are used more selectively, providing additional control authority to shape the internal thermal field and limit the non-uniformity introduced by SC. Over the complete evaluation period, the surface channel accounts for approximately 63% and 60% of the total coolant utilisation for the 2.3 Ah and 45 Ah cells, respectively. The remaining coolant resource is distributed approximately equally between the two tab channels. The coolant demand is substantially higher for the 45 Ah cell, particularly during driving. This reflects the greater thermal regulation requirements of the larger cell, as evidenced by the higher temperatures and spatial temperature gradients in Fig. 8(c-ii) and (d-ii). Furthermore, the close degradation performance of bTSC to ITSC suggests that combining a surface cooling path with a tab cooling path captures much of the useful thermal control authority, whereas the tab-only configurations improve thermal uniformity at the expense of a higher mean temperature.

To further examine the influence of spatial non-uniformity under closed-loop operation, Fig. 9 shows the local electrothermal and degradation responses at four representative quadrature nodes for both cells under ITSC. The corresponding locations are selected near the core, surface, top, and bottom regions of the cell. The local temperature trajectories in Fig. 9(a) show that both the magnitude and spatial distribution of temperature vary throughout the operating sequence. During the high-current driving phases, the inner region generally experiences the largest temperature excursions. However, the relative temperatures of the different regions change during charging and rest

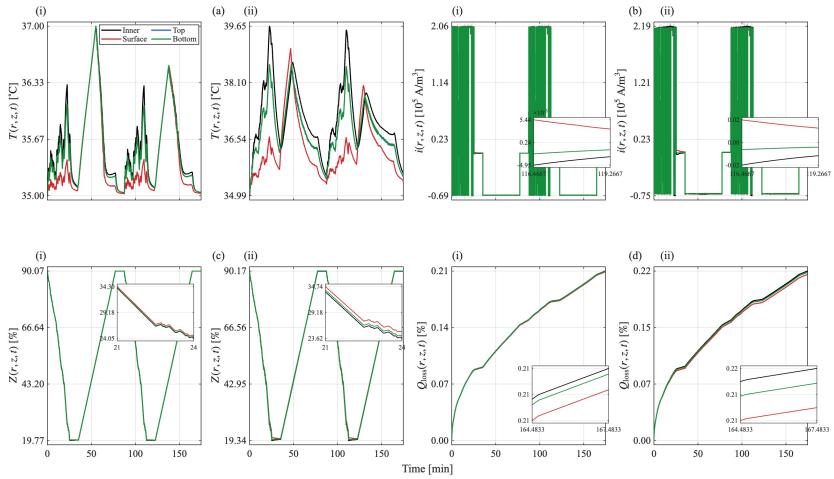

Figure 9: Local electrothermal and degradation responses under the degradation-aware ITSC strategy for (i) the 2.3 Ah cell and (ii) the 45 Ah cell over the first two operating sequences: (a) local temperatures, (b) local volumetric current densities, (c) local SoCs, and (d) local accumulated capacity losses at four representative quadrature nodes located near the core, surface, top, and bottom regions of the cell. The enlarged views highlight the spatial differences in the respective quantities.

as the heat generation and coolant allocation vary. The spatial temperature separation is considerably more pronounced for the 45 Ah cell, where the local temperature approaches 39.7 °C during driving.

The non-uniform electrothermal conditions give rise to spatial current redistribution, as illustrated in Fig. 9(b). Although the local current densities largely follow the imposed battery current profile, their magnitudes differ among the quadrature nodes. During driving, the relatively cooler surface region tends to exhibit higher electrical resistance, resulting in a smaller share of the discharge current and consequently a higher local SoC, as evident from the enlarged views in Fig. 9(c). During rest, the differences in local SoCs and polarisation voltages drive internal redistribution currents despite the net battery current being zero. In particular, the surface region locally discharges, while the inner region undergoes charging. These spatial differences are more

pronounced in the 45 Ah cell, consistent with its greater electrothermal non-uniformity. An important implication of this behaviour is that certain regions of the cell may reach their local SoC limits before the volume-averaged SoC reaches its prescribed operating bounds. Consequently, monitoring only the cell-level SoC may not adequately capture the internal electrochemical operating conditions, particularly in larger cells exhibiting substantial thermal non-uniformity.

The combined influence of the local temperature, current density, and SoC histories is reflected in the accumulated local capacity losses in Fig. 9(d). For both cells, the inner region accumulates more degradation than the surface region at the end of the displayed interval, with a considerably larger separation for the 45 Ah cell. Over the complete evaluation period, the maximum spatial differences in local capacity loss under ITSC are approximately 0.00254 and 0.01245 percentage points for the 2.3 Ah and 45 Ah cells, respectively. Thus, despite their comparable spatially averaged capacity losses, the larger cell exhibits substantially greater non-uniformity in its degradation.

Overall, the results demonstrate the ability of the degradation-aware ITSC strategy to exploit the complementary characteristics of surface and tab cooling according to their combined influence on battery degradation. ITSC achieves the lowest accumulated capacity loss among the feedback-controlled configurations for both cells, while attaining comparable degradation performance to continuous full-flow ES cooling with substantially lower coolant utilisation. Furthermore, the greater spatial non-uniformity observed in the larger cell highlights the importance of accounting for the distributed electrothermal and degradation dynamics when evaluating battery thermal management strategies.

5.4 Degradation minimisation vs. thermal-proxy minimisation

For the OCAP, battery health is represented indirectly through temperature tracking and thermal gradient minimisation. The original formulation in [16] is retained, with the addition of the coolant utilisation term employed in

DOCAP, giving

$$\min_{u_{\mathrm{mod},k}} J(k) := \sum_{i=0}^{N_p-1} \Bigg[\lambda_T \left\| T^r(k+i) - \hat{\mathcal{Y}}_t(k+i) \right\|_{Q_e}^2 + \lambda_\delta \left\| \delta\hat{T}(k+i) \right\|_{Q_\delta}^2 + \lambda_u \left(1^T u_{\mathrm{mod}}(k+i) \right)^2 + \lambda_{\Delta u} \left\| \Delta u_{\mathrm{mod}}(k+i) \right\|_R^2 \Bigg]. \tag{C.73}$$

The first term in (C.73) regulates the predicted temperature outputs $\hat{\mathcal{Y}}_t$ towards a desired reference T^r, while the second penalises internal temperature non-uniformity through the gradient-based measure $\delta\hat{T}$ defined in (C.16b). The third and fourth terms penalise coolant utilisation and control input variation, respectively.

The comparison between DOCAP in (C.54) and OCAP in (C.73) is conducted for both the 2.3 Ah and 45 Ah cells. Since both cells exhibit qualitatively similar responses, only the results for the 45 Ah cell are presented for brevity. Both controllers employ the same ITSC configuration, prediction model, RTI-MPC settings, ten-sequence input current profile, constraints, and available coolant-flow capacity. The initial conditions $T(r, z, 0)$ and $X_{cl}(0)$, coolant inlet temperature $T_{cl,\mathrm{in}}$, and OCAP temperature reference T^r are all set to 35 °C. In addition, the coolant utilisation and control variation weights are kept identical between the two formulations, such that $\lambda_2 = \lambda_u$ and $\lambda_3 = \lambda_{\Delta u}$, respectively. The original OCAP thermal weights are retained as $\lambda_T = \lambda_\delta = 1$. The same distributed degradation model is propagated under both controllers to evaluate the capacity loss resulting from their respective electrothermal and coolant allocation trajectories.

Fig. 10(a) shows that DOCAP and OCAP achieve similar accumulated capacity loss over the ten operating sequences, with OCAP achieving a marginally lower capacity loss of 0.48147%, compared with 0.48209% under DOCAP. However, this small reduction in degradation is obtained with substantially greater coolant utilisation, as shown in Fig. 10(b). The cumulative coolant duty increases from 5.46 h under DOCAP to 14.56 h under OCAP, corresponding to approximately 2.67 times greater coolant utilisation. Thus, DOCAP achieves comparable degradation performance while reducing coolant utilisation by approximately 62.5%.

This difference arises from the quantities represented in the two objectives. OCAP continuously penalises deviations from the reference temperature and

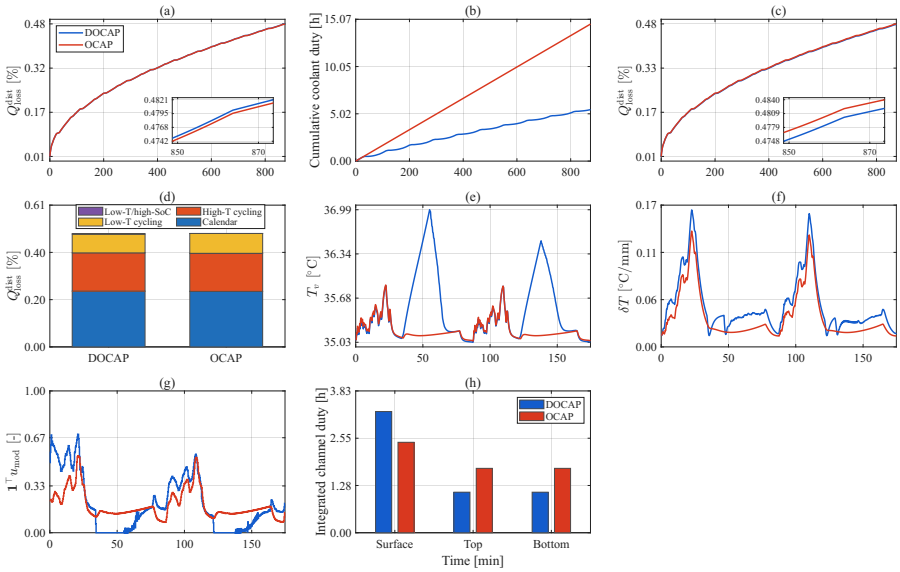

Figure 10: Comparison of DOCAP and OCAP for the 45 Ah cell over ten operating sequences. (a) Accumulated capacity loss and (b) cumulative coolant utilisation using the original OCAP thermal weights. (c)–(h) Comparison under comparable coolant utilisation: (c) accumulated capacity loss, (d) individual degradation mechanism contributions, (e) volume-averaged temperature, (f) RMS spatial temperature gradient, (g) total duty-cycle, and (h) cumulative channel-wise coolant duty. The trajectories in (e)–(g) are shown for the first two operating sequences.

internal temperature gradients, resulting in continuous full-flow cooling under the considered thermal weights. DOCAP, however, determines the cooling action according to its predicted influence on degradation relative to the associated coolant utilisation cost. As demonstrated in Section 3.3, although lowering the cell temperature reduces calendar and high-temperature cycling ageing, it simultaneously increases low-temperature cycling ageing. Consequently, the substantial increase in coolant utilisation under OCAP produces only a marginal reduction in total degradation.

To examine the influence of the control objective under comparable coolant utilisation, the OCAP thermal weights are uniformly scaled while preserving

their relative weighting. The common scaling gives $\lambda_T = \lambda_\delta = 3.906 \times 10^{-3}$, yielding a cumulative coolant duty of 5.92 h for OCAP, compared with 5.46 h for DOCAP. The corresponding responses are shown in Fig. 10(c)–(h). Despite utilising approximately 8.37% more coolant, OCAP accumulates a higher capacity loss of 0.48401%, corresponding to approximately 0.40% greater degradation than DOCAP.

Interestingly, the lower degradation obtained with DOCAP does not correspond to improved thermal metrics. As shown in Fig. 10(e) and (f), OCAP generally maintains a lower volume-averaged temperature and smaller spatial temperature gradients throughout the operating sequence. Nevertheless, these improvements in thermal regulation and uniformity do not translate into lower accumulated capacity loss. The degradation mechanism decomposition in Fig. 10(d) provides insight into this behaviour. The differences between the individual contributions arise because the degradation mechanisms respond to the local temperature, current density, SoC, and throughput histories rather than to the mean temperature or spatial temperature gradients alone. In particular, high-temperature cycling ageing accumulates with the total charge throughput and is therefore sensitive to the temperatures prevailing during high-current operating periods, whereas low-temperature cycling ageing accumulates during charging and increases with decreasing temperature. Consequently, reducing the mean temperature or thermal gradients during operating periods in which the associated degradation benefit is small may result in inefficient utilisation of the available coolant resource. DOCAP instead evaluates the predicted thermal response through the individual degradation mechanisms and thereby explicitly accounts for the prevailing local electrothermal operating conditions when determining the cooling action.

The different objectives also result in markedly different spatial and temporal coolant allocations. Fig. 10(h) shows that DOCAP allocates a greater proportion of its coolant resource to the surface channel, whereas OCAP distributes a greater proportion through the two tab channels. This behaviour is consistent with the gradient minimisation term in OCAP, which favours the temperature homogenising capability of tab cooling. DOCAP instead allocates coolant according to its combined influence on the local degradation mechanisms, resulting in greater utilisation of the surface cooling path. The temporal allocation in Fig. 10(g) further reveals that DOCAP employs stronger cooling actions during the driving phases, while substantially reduc-

ing coolant utilisation during charging. OCAP, in contrast, distributes its coolant resource more continuously throughout the operating sequence. The resulting temperature trajectories in Fig. 10(e) reflect these different allocations, with DOCAP permitting larger temperature excursions during charging. These warmer charging conditions reduce the low-temperature cycling stress, illustrating how a degradation-aware controller can exploit the competing temperature dependencies of the individual degradation mechanisms.

Overall, the results demonstrate that temperature regulation and thermal uniformity are not necessarily equivalent to degradation minimisation. With the original thermal weights, OCAP achieves comparable degradation performance at substantially greater coolant utilisation. When the thermal weights are adjusted to obtain comparable coolant utilisation, DOCAP achieves lower accumulated degradation despite allowing larger temperature excursions and spatial temperature gradients. These findings highlight the benefits of explicitly accounting for the operating-condition-dependent degradation mechanisms when determining both the temporal and spatial allocation of the available cooling resource.

6 Conclusion

This article has developed a degradation-aware predictive battery–coolant control framework for dynamically allocating a finite coolant resource among the surface and tab cooling channels of battery cells. Starting from an experimentally parametrised semi-empirical degradation model, the ageing dynamics were reformulated as state equations and extended to account for the spatial distributions of temperature, current density, SoC, and charge throughput. The resulting distributed electrothermal–degradation model captures the coupled effects of thermal non-uniformity, current redistribution, and spatially varying degradation within the cell. The results showed that these spatial differences become more pronounced in larger cells, even when their spatially averaged degradation is comparable to that of smaller cells.

The degradation-aware optimal coolant allocation problem (DOCAP) was implemented using a real-time iteration model predictive control (RTI-MPC) scheme to reduce the online computational burden. Through local linearisation of the nonlinear prediction dynamics and quadratic program-based optimisation, RTI-MPC closely reproduced the closed-loop degradation and

thermal responses of nonlinear MPC, while achieving an approximately 46-fold computational speed-up and execution times well below the prescribed sampling interval.

Evaluation of the 2.3 Ah and 45 Ah cells over ten consecutive drive–rest–charge–rest sequences revealed the competing effects of the different cooling configurations. While tab-only cooling maintained comparatively uniform internal temperature fields, its higher operating temperatures increased the dominant calendar and high-temperature cycling ageing contributions. The integrated tab–surface cooling (ITSC) system instead exploited the complementary characteristics of the cooling interfaces according to their combined influence on degradation. ITSC achieved the lowest accumulated capacity loss among the feedback-controlled configurations for both cells, while attaining comparable degradation performance to continuous equal-split full-flow cooling with 85% and 62.5% lower coolant utilisation for the 2.3 Ah and 45 Ah cells, respectively. The distributed modelling framework further revealed that spatially non-uniform temperatures induce current redistribution and local SoC differences, including internal currents during rest despite a zero net battery current. These effects were more pronounced in the 45 Ah cell, which also exhibited approximately five times greater spatial variation in accumulated capacity loss under ITSC than the 2.3 Ah cell. Relying solely on the volume-averaged SoC may therefore lead to undetected local violations of the prescribed SoC limits. Comparison with conventional thermal-proxy control showed that improved temperature regulation and thermal uniformity do not necessarily translate into lower degradation. Under comparable coolant utilisation, DOCAP achieved lower accumulated capacity loss despite the conventional controller maintaining a lower mean temperature and smaller spatial temperature gradients. Effective ageing mitigation depends not only on reducing temperature or thermal non-uniformity, but also on allocating the available coolant spatially and temporally according to its influence on the active degradation mechanisms. Direct degradation minimisation enables these competing effects to be accounted for concurrently.

Future work will focus on experimentally validating ageing-induced changes in cell electrothermal dynamics and extending the proposed framework to degradation-aware fast-charging control with local electrothermal operating constraints. Further investigations will consider joint optimisation of pump energy consumption and coolant allocation, as well as extending the modelling

and control framework to battery modules and packs.

Appendix A Electrical model matrices

This appendix provides the explicit expressions of the system matrices in (C.8).

$$\mathcal{A}_e = \begin{bmatrix} 0 & 0 & 0 \\ 0 & -\frac{1}{R_1 C_1} & 0 \\ 0 & 0 & -\frac{1}{R_2 C_2} \end{bmatrix}, \qquad \mathcal{B}_e = \begin{bmatrix} -\frac{1}{3600 C_{\text{avail}}} \\ \frac{1}{C_1} \\ \frac{1}{C_2} \end{bmatrix}. \tag{C.74}$$

For notational simplicity, the spatial and temporal arguments of the ECM parameters are omitted here.

Appendix B Degradation modelling

This appendix presents the detailed degradation modelling, including the lumped model and its extension to account for spatial variations in temperature, SoC, and current.

B.a Lumped degradation model

The stress factors corresponding to the four degradation mechanisms in (C.18) are expressed as follows [26].

$$k_C(\bar{T}, Z) = k_{C,\text{ref}} \exp\left[-\frac{E_{a,C}}{R_g}\left(\frac{1}{\bar{T}} - \frac{1}{T_{\text{ref}}}\right)\right] \cdot g_C(Z), \tag{C.75a}$$

$$k_H(\bar{T}) = k_{H,\text{ref}} \exp\left[-\frac{E_{a,H}}{R_g}\left(\frac{1}{\bar{T}} - \frac{1}{T_{\text{ref}}}\right)\right], \tag{C.75b}$$

$$k_L(\bar{T}, I_{\text{ch}}) = k_{L,\text{ref}} \exp\left[\frac{E_{a,L}}{R_g}\left(\frac{1}{\bar{T}} - \frac{1}{T_{\text{ref}}}\right)\right] \cdot \exp\left[\beta_L \frac{I_{\text{ch}} - I_{\text{ch,ref}}}{C_0}\right], \tag{C.75c}$$

$$k_{LS}(\bar{T}, I_{\text{ch}}, Z) = k_{LS,\text{ref}} \cdot \exp\left[\frac{E_{a,LS}}{R_g}\left(\frac{1}{\bar{T}} - \frac{1}{T_{\text{ref}}}\right)\right] \cdot \exp\left[\beta_{LS} \frac{I_{\text{ch}} - I_{\text{ch,ref}}}{C_0}\right] \cdot g_{H\text{SoC}}(Z). \tag{C.75d}$$

The SoC-dependent terms are

$$g_C(Z) = \exp\left[\frac{\alpha F}{R_g}\left(\frac{U_{a,\text{ref}} - U_a(Z)}{T_{\text{ref}}}\right)\right] + k_0, \tag{C.76a}$$

$$g_{H\text{SoC}}(\text{SoC}) = \frac{\text{sgn}\left(Z - Z_{\text{ref}}\right) + 1}{2}. \tag{C.76b}$$

E_a denotes the activation energy, R_g, the universal gas constant, F, the Faraday constant, α, the transfer coefficient, U_a, the anode open-circuit potential, k_0, a constant offset, C_0, the nominal cell capacity, β, the charge-current sensitivity coefficient, and the subscript ref, the reference values.

B.b Distributed degradation model

The distributed degradation model (C.28)–(C.29) is obtained by evaluating the stress factors in (C.75) using the local temperature, SoC, reference-capacity-scaled local current, and accumulated throughputs. The individual

local capacity loss rates are given by

$$\dot{Q}_C(r,z,t) = \frac{k_C(T,Z)}{7200\sqrt{\tau+\epsilon_\tau}}, \tag{C.77a}$$

$$\dot{Q}_H(r,z,t) = k_H(T)\frac{|\tilde{I}_\ell|}{7200\sqrt{\varphi_{\text{tot}}+\epsilon_{\text{tot}}}}, \tag{C.77b}$$

$$\dot{Q}_L(r,z,t) = k_L(T,I_{\text{ch}})\frac{I_{\text{ch}}}{7200\sqrt{\varphi_{\text{ch}}+\epsilon_{\text{ch}}}}, \tag{C.77c}$$

$$\dot{Q}_{LS}(r,z,t) = k_{LS}(T,I_{\text{ch}},Z)\frac{I_{\text{ch}}}{3600}. \tag{C.77d}$$

For notational simplicity, the spatial and temporal arguments of the operating variables and degradation states on the right-hand side of (C.77) are omitted.

References

[1] M. Safari, "Battery electric vehicles: Looking behind to move forward," *Energy Policy*, vol. 115, pp. 54–65, 2018.

[2] A. Škegro, T. Wik, B. Bijlenga, A. Bessman, and C. Zou, "System-level assessment of dynamic reconfiguration for lifetime and cost outcomes in electric vehicle battery packs," *Nat. Commun.*, vol. 17, no. 1, 2026, Art. no. 5980.

[3] S. Ma et al., "Temperature effect and thermal impact in lithium-ion batteries: A review," *Prog. Nat. Sci.: Mater. Int.*, vol. 28, no. 6, pp. 653–666, 2018.

[4] J. Lin, X. Liu, S. Li, C. Zhang, and S. Yang, "A review on recent progress, challenges and perspective of battery thermal management system," *Int. J. Heat Mass Transf.*, vol. 167, p. 120 834, 2021.

[5] C. Bibin, M. Vijayaram, V. Suriya, R. S. Ganesh, and S. Soundarraj, "A review on thermal issues in Li-ion battery and recent advancements in battery thermal management system," *Mater. Today Proc.*, vol. 33, pp. 116–128, 2020.

[6] R. Carter et al., "Directionality of thermal gradients in lithium-ion batteries dictates diverging degradation modes," *Cell Rep. Phys. Sci.*, vol. 2, no. 3, 2021, Art. no. 100351.

[7] X. Du et al., "Experimental studies of effects of temperature gradients on performance of pouch type large format NMC/C lithium-ion battery," *J. Power Sources*, vol. 587, 2023, Art. no. 233688.

[8] S. Li, C. Zhang, Y. Zhao, G. J. Offer, and M. Marinescu, "Effect of thermal gradients on inhomogeneous degradation in lithium-ion batteries," *Commun. Eng.*, vol. 2, no. 1, 2023, Art. no. 74.

[9] Y. Zhao, L. B. Diaz, Y. Patel, T. Zhang, and G. J. Offer, "How to cool lithium ion batteries: Optimising cell design using a thermally coupled model," *J. Electrochem. Soc.*, vol. 166, no. 13, A2849–A2859, 2019.

[10] Y. Zhao, Y. Patel, T. Zhang, and G. J. Offer, "Modeling the effects of thermal gradients induced by tab and surface cooling on lithium ion cell performance," *J. Electrochem. Soc.*, vol. 165, no. 13, A3169–A3178, 2018.

[11] O. Dondelewski et al., "The role of cell geometry when selecting tab or surface cooling to minimise cell degradation," *eTransportation*, vol. 5, p. 100 073, 2020.

[12] S. Li et al., "Optimal cell tab design and cooling strategy for cylindrical lithium-ion batteries," *J. Power Sources*, vol. 492, p. 229 594, 2021.

[13] C. Bolsinger and K. P. Birke, "Effect of different cooling configurations on thermal gradients inside cylindrical battery cells," *J. Energy Storage*, vol. 21, pp. 222–230, 2019.

[14] T. Ahmad, A. Mishra, S. Ghosh, and C. S. Casari, "Identifying efficient cooling approach of cylindrical lithium-ion batteries," *Energy Technol.*, vol. 10, no. 2, p. 2 100 888, 2022.

[15] G. K. Peprah, T. Wik, Y. Huang, F. Altaf, and C. Zou, "Model predictive cooling control of cylindrical battery cells through tab and surface channels," in *Proc. IEEE Amer. Control Conf. (ACC)*, 2025, pp. 4071–4076.

[16] G. K. Peprah, T. Wik, M. Tamadondar, and C. Zou, *Model predictive coolant allocation for integrated tab-surface cooling of battery cells*, arXiv:2607.10872, 2026.

[17] U. S. Kim, J. Yi, C. B. Shin, T. Han, and S. Park, "Modelling the thermal behaviour of a lithium-ion battery during charge," *J. Power Sources*, vol. 196, no. 11, pp. 5115–5121, 2011.

[18] X. Lin et al., "Online parameterization of lumped thermal dynamics in cylindrical lithium ion batteries for core temperature estimation and health monitoring," *IEEE Trans. Control Syst. Technol.*, vol. 21, no. 5, pp. 1745–1755, 2013.

[19] X. Lin et al., "A lumped-parameter electro-thermal model for cylindrical batteries," *J. Power Sources*, vol. 257, pp. 1–11, 2014.

[20] G. K. Peprah, F. Liberati, F. Altaf, G. Osei-Dadzie, A. Di Giorgio, and A. Pietrabissa, "Optimal load sharing in reconfigurable battery systems using an improved model predictive control method," in *Proc. 29th Mediterranean Conf. Control Autom. (MED)*, IEEE, 2021, pp. 979–984.

[21] G. K. Peprah, Y. Huang, T. Wik, F. Altaf, and C. Zou, "Thermal modeling of battery cells for optimal tab and surface cooling control," *IEEE Trans. Control Syst. Technol.*, vol. 34, no. 4, pp. 1831–1844, 2026.

[22] G. K. Peprah, T. Wik, Y. Huang, F. Altaf, and C. Zou, "Control-oriented 2D thermal modelling of cylindrical battery cells for optimal tab and surface cooling," in *2024 IEEE American Control Conf. (ACC)*, pp. 4651–4656.

[23] X. Jin et al., "Applicability of available Li-ion battery degradation models for system and control algorithm design," *Control Eng. Pract.*, vol. 71, pp. 1–9, 2018.

[24] J. S. Edge et al., "Lithium ion battery degradation: What you need to know," *Phys. Chem. Chem. Phys.*, vol. 23, no. 14, pp. 8200–8221, 2021.

[25] W. Vermeer, G. R. C. Mouli, and P. Bauer, "A comprehensive review on the characteristics and modeling of lithium-ion battery aging," *IEEE Trans. Transp. Electrific.*, vol. 8, no. 2, pp. 2205–2232, 2022.

[26] M. Schimpe, M. E. von Kuepach, M. Naumann, H. C. Hesse, K. Smith, and A. Jossen, "Comprehensive modeling of temperature-dependent degradation mechanisms in lithium iron phosphate batteries," *J. Electrochem. Soc.*, vol. 165, no. 2, A181–A193, 2018.

[27] Y. Xie, C. Wang, X. Hu, X. Lin, Y. Zhang, and W. Li, "An MPC-based control strategy for electric vehicle battery cooling considering energy saving and battery lifespan," *IEEE Trans. Veh. Technol.*, vol. 69, no. 12, pp. 14 657–14 673, 2020.

[28] C. Wei, T. Hofman, E. I. Caarls, and R. van Iperen, "Zone model predictive control for battery thermal management including battery aging and brake energy recovery in electrified powertrains," *IFAC-PapersOnLine*, vol. 52, no. 5, pp. 303–308, 2019.

[29] Y. Wu et al., "Optimal battery thermal management for electric vehicles with battery degradation minimization," *Appl. Energy*, vol. 353, 2024, Art. no. 122090.

[30] D. Li, Q. Hu, W. Jiang, H. Dong, and Z. Song, "Integrated power and thermal management for enhancing energy efficiency and battery life in connected and automated electric vehicles," *Appl. Energy*, vol. 396, 2025, Art. no. 126213.

[31] M. Diehl, H. G. Bock, and J. P. Schlöder, "A real-time iteration scheme for nonlinear optimization in optimal feedback control," *SIAM J. Control Optim.*, vol. 43, no. 5, pp. 1714–1736, 2005.

[32] S. Gros, M. Zanon, R. Quirynen, A. Bemporad, and M. Diehl, "From linear to nonlinear MPC: Bridging the gap via the real-time iteration," *Int. J. Control*, vol. 93, no. 1, pp. 62–80, 2020.

[33] H. E. Perez, J. B. Siegel, X. Lin, A. G. Stefanopoulou, Y. Ding, and M. P. Castanier, "Parameterization and validation of an integrated electro-thermal cylindrical LFP battery model," in *Proc. ASME 5th Annu. Dyn. Syst. Control Conf. Joint JSME 11th Motion Vib. Conf. (DSCC/MOVIC)*, doi: 10.1115/DSCC2012-MOVIC2012-8782, Oct. 2012, pp. 41–50.

[34] Y. Xie, S. Shi, J. Tang, H. Wu, and J. Yu, "Experimental and analytical study on heat generation characteristics of a lithium-ion power battery," *Int. J. Heat Mass Transf.*, vol. 122, pp. 884–894, 2018.

[35] M. Naumann, M. Schimpe, P. Keil, H. C. Hesse, and A. Jossen, "Analysis and modeling of calendar aging of a commercial $LiFePO_4$/graphite cell," *J. Energy Storage*, vol. 17, pp. 153–169, 2018.

[36] M. Naumann, F. B. Spingler, and A. Jossen, "Analysis and modeling of cycle aging of a commercial $LiFePO_4$/graphite cell," *J. Power Sources*, vol. 451, 2020, Art. no. 227666.

[37] Y.-S. Duh, J.-H. Theng, C.-C. Chen, and C.-S. Kao, "Comparative study on thermal runaway of commercial 14500, 18650 and 26650 LiFePO4 batteries used in electric vehicles," *J. Energy Storage*, vol. 31, p. 101 580, 2020.

[38] G.-H. Kim and A. Pesaran, "Battery thermal management system design modeling," National Renewable Energy Lab. (NREL), Golden, CO, Tech. Rep., 2006.

[39] L. Acker, P. Hofmann, and J. Konrad, "Predictive battery thermal management for fast charging of electric vehicles using nonlinear model predictive control and dynamic programming," *Automot. Engine Technol.*, vol. 11, no. 1, 2026, Art. no. 1.

[40] J. Lopez-Sanz et al., "Thermal management in plug-in hybrid electric vehicles: A real-time nonlinear model predictive control implementation," *IEEE Trans. Veh. Technol.*, vol. 66, no. 9, pp. 7751–7760, 2017.

[41] T. Fischer, T. Kraus, C. Kirches, and F. Gauterin, "Demonstration of a nonlinear model predictive control of a thermal management system for electric vehicles in real-time," in *Proc. IEEE Conf. Control Technol. Appl. (CCTA)*, 2018, pp. 676–682.

[42] P. Mock et al., "The WLTP: How a new test procedure for cars will affect fuel consumption values in the EU," *Intern. council on clean transp.*, vol. 9, no. 3547, 2014.

[43] J. Lofberg, "Yalmip: A toolbox for modeling and optimization in matlab," in *2004 IEEE Int. Conf. on Robotics and Automation (IEEE Cat. No. 04CH37508)*, pp. 284–289.