HIERARCHICAL CONTROL FOR MULTI-DOMAIN COORDINATION OF VEHICLE ENERGY SYSTEMS WITH SWITCHED DYNAMICS

BY

HERSCHEL CYRUS PANGBORN

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Mechanical Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2019

Urbana, Illinois

Doctoral Committee:

Professor Andrew Alleyne, Chair
Professor Geir Dullerud
Professor Daniel Liberzon
Associate Professor Carolyn Beck
Professor John Wen, Rensselaer Polytechnic Institute

Abstract

This dissertation presents a hierarchical control framework for vehicle energy management. As a result of increasing electrification, legacy integration and control approaches for vehicle energy systems have become limiting factors of performance and cannot accommodate the requirements of next-generation systems. Addressing this requires control frameworks that coordinate dynamics across multiple physical domains and timescales, enabling transformative improvements in capability, efficiency, and safety.

To capture multi-domain storage and exchange of energy, a graph-based dynamic modeling approach is proposed and experimentally validated. This modeling approach is then leveraged for model-based control, in which the complex task of energy management is decomposed into a hierarchical network of model predictive controllers that coordinate decision-making across subsystems, physical domains, and timescales. The controllers govern both continuous and switched dynamic behaviors, addressing the hybrid nature of modern vehicle energy systems.

The proposed hierarchical control framework is evaluated in application to a hardware-in-the-loop electro-thermal testbed representative of a scaled aircraft energy system, where it achieves significantly improved capability, efficiency, and safety as compared to legacy control approaches. Next, the structural information embedded in the graph-based modeling approach is shown to facilitate analysis. Closed-loop stability of decentralized MPC frameworks is guaranteed by analyzing the passivity of switched nonlinear graph-based systems and augmenting their controllers with a local passivity-based constraint. Lastly, a hierarchical control formulation guaranteeing satisfaction of state and input constraints for a class of switched graph-based systems is presented. This formulation is demonstrated in application to thermal management using both simulation and experimental implementation.

To my friends, family, and teachers.

Acknowledgments

I would first like to thank my advisor, Professor Andrew Alleyne, for his invaluable guidance throughout my graduate studies. I am fortunate to have a mentor who teaches his mentees how to think, rather than telling them what to think. If my future career can serve to pay forward the experiences and opportunities that he has facilitated for me, I will consider it a success.

I would like to acknowledge the members of my doctoral examination committee, Professors Geir Dullerud, Daniel Liberzon, Carolyn Beck, and John Wen, for their support and feedback throughout this process.

My graduate studies have benefited immensely from participation in several research centers. I would like to thank the leadership of the Center of Excellence for Integrated Thermal Management of Aerospace Vehicles (CITMAV), especially Professor Tim Fisher, for the opportunity to engage in its important research. I also wish to thank the faculty and staff of the NSF Engineering Research Center for Power Optimization of Electro-Thermal Systems (POETS) for their commitment to setting students up for success. These individuals include Professors Andrew Alleyne and Alan Mantooth, Dr. Karen Bender, Sonya Chambers, Joe Muskin, and Dr. Jessica Perez. The experimental results in the penultimate chapter of this dissertation would not have been possible without Cyril Rybicki's expertise in data acquisition software development. Special thanks to Jodi Gritten, whose hard work has made unparalleled contributions to the success of the center, all while brightening every day with a smile at the door. I have also been fortunate to work with the outstanding team at CU Aerospace. Thanks in particular to Dr. David Carroll and Dr. Joe Zimmerman for their engagement and support.

Thank you to the the past and present members of the Alleyne Research group (ARG) with whom my tenure intersected. My research has benefitted greatly from many discussions and collaborations with these individuals, and their friendship has truly enriched my time at UIUC. Thanks to Yangmin Xie, Nanjun Liu, Erick Sutanto, and Tim Deppen for welcoming me to ARG and continuing a precedent of success for me to follow. Thanks to Katie Amber, Lindsey Gonzales, Bryan Keating, Spencer Kieffer, Sunny Sharma, Sarah Garrow, Nate Weir, and Malia Kawamura for their companionship and steadying presence through the ups and downs of this journey. Thanks to Ashley Armstrong, Pamela Tannous, Donald Docimo, Oyuna Angatkina, Spencer Igram, Chris Aksland, and Cary Laird for continuing the ARG culture of friendship, support, and success. Thanks especially to Justin Koeln and Matt Williams, with whom I have enjoyed close collaboration and friendship, and look forward to more to come.

I must also thank my family. My parents enduring love and support have given me the confidence to pursue my passions. My father, Robert, has equipped me with a lifetime of lessons in leadership, exemplifying what it means to support a family and community. My mother, Ann, has taught me to have compassion, seek balance, and pursue peace. My sisters have always been kind enough to include their younger brother in their lives, and in doing so have served as shining examples to follow in my own search for self.

I would not have made it to the finish line without Lori Morgan at my side. On countless occasions she has helped me to stay the course under stress and smoothed a furrowed brow, not to mention filling an empty stomach.

Lastly, I would like to acknowledge the support of the National Science Foundation (NSF) Graduate Research Fellowship Program (GRFP) under Grant Number DGE-1144245, and the NSF Engineering Research Center (ERC) for Power Optimization of Electro-Thermal Systems (POETS) under cooperative agreement EEC-1449548. This support has given me the flexibility to pursue research that I believe will not only contribute to my field, but also serve society.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation and Background

1.1.1 Vehicle Electrification

Modern vehicles are complex machines consisting of interconnected subsystems that interact in multiple physical domains across a variety of dynamic timescales. A ubiquitous and longstanding megatrend in the design of these "systems of systems" is their electrification. For example, Fig. 1.1 depicts the exponential increase in on-board electrical power of military and commercial aircraft [1, 2]. Beyond electrification within conventional platforms and powertrains, many new electrified vehicle classes are undergoing rapid development. These include automotive electric vehicles [3], all-electric passenger aircraft [4], electric-drive off-road vehicles [5], and electric-drive naval ships [6].

The replacement of traditional mechanical, hydraulic, and pneumatic components of vehicles by electrical systems provides many benefits, including advanced capability [7], enhanced safety [8], improved efficiency [2], reduced mass and volume [9], decreased lifetime costs [10], smaller carbon footprint [9], and greater ability to leverage renewable energy sources. However, electrification is accompanied by increased thermal loading due to inefficiencies of electrical components. This places additional burden on thermal management systems responsible for transferring, storing, and rejecting thermal energy to maintain temperature constraints throughout the vehicle [2, 7, 11]. This challenge is compounded by the continual desire to reduce the size and weight of thermal management systems. It is further exacerbated in aviation by a decreased ability to reject thermal energy due to the use of composite skin materials with high thermal resistance and reduction in ram air heat exchanger cross

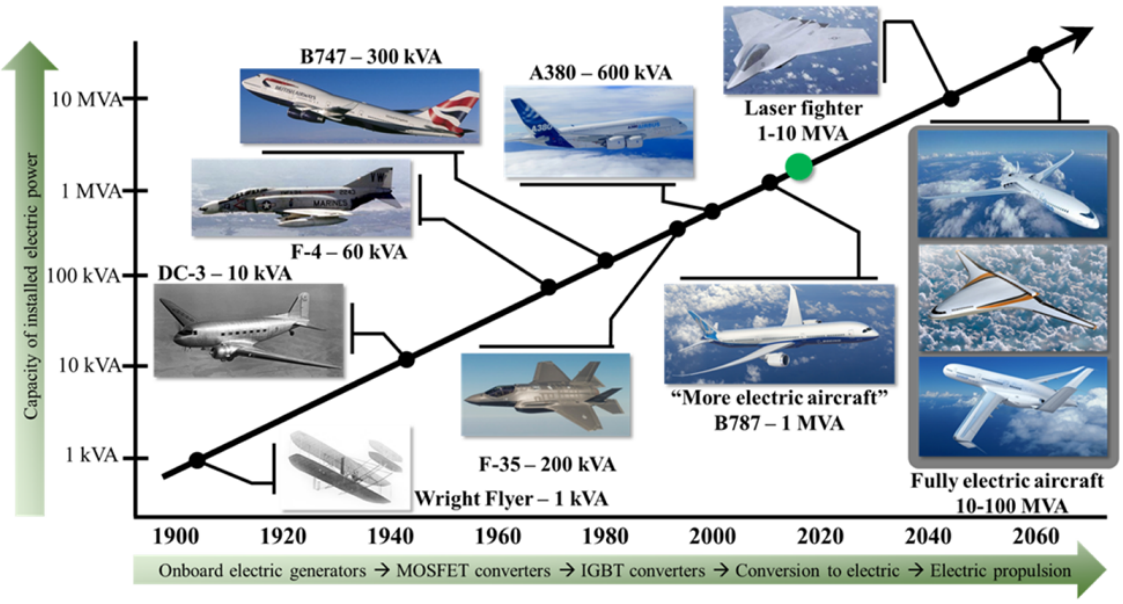

Figure 1.1: Electrification in aviation: Historical and projected trends [1].

sections [12]. The difficulties of sufficiently extracting thermal energy from electronic systems is exemplified by the attribution of greater than 50% of failures in military electronics to insufficient thermal management, with similar issues occurring in commercial systems [13, 14].

1.1.2 Multi-Domain Coupling and Switching

Vehicle energy systems are composed of subsystems and components that interact with each other in multiple physical domains, including the mechanical, electrical, thermal, hydraulic, and pneumatic domains. The associated dynamic timescales rage across orders of magnitude, from sub-milliseconds for electrical states [15] to minutes for thermal states [12]. This coupling across domains can play a strong role in key figures of merit associated with capability, safety, and efficiency.

Of particular relevance to this dissertation are the energy interactions within and between electrical and thermal systems. In many vehicle powertrains, electrical energy is generated by extracting mechanical power from one or more engines. This extraction affects the efficiency and performance of the propulsion system, and must be constrained to ensure that high

mechanical stress, torque ripple, and speed transients do not cause events such as compressor stall [16]. Electrical energy is stored in batteries and capacitors, passed through transformers and converters, and used to power electrical and mechanical devices. Electrical energy is also converted to thermal energy as a result of inefficiencies in solid state electronics, motors, and other equipment.

The efficiency of electrical systems can depend strongly on their operating temperature [17]. Reliability and degradation can also be strongly temperature-dependent [14, 18]. Therefore, thermal management systems are tasked with transporting thermal energy away from electrical components to maintain constraints on their operating temperature. As an intermediate step before its rejection to the environment, thermal energy can be stored in the thermal capacitances of metal cold plates, air volumes, liquid tanks, and tanks of phase-change material. Air, single-phase liquid, or phase-change refrigerant are often used as a thermal transport medium. The resulting fluid-based thermal management systems, or "fluid-thermal systems," are therefore governed by both conservation of energy and conservation of mass of the working fluid. The fluid is circulated by actuators such as fans, pumps, and compressors, which are often driven electrically. As a result, increasing power to actuators of the thermal management system can result in additional inefficiencies in the electrical system, generating more thermal energy. Actuators of the thermal management system may also produce thermal energy more directly, for example by friction heating in pumps [19].

The above discussion highlights the bidirectional coupling between electrical and thermal systems in vehicles. The efficiency and degradation of electrical systems depends on their thermal state, while thermal systems operate subject to the thermal energy generated by electrical systems and themselves consume electrical energy. While traditional analyses have often treated this coupling as quasi-steady-state, it is becoming increasingly transient in nature [11], in large part due to the integration of electrical devices characterized by high ratios of peak to average power that may only be active during particular mission segments (e.g., takeoff or landing) or are operated by pulsed power [20].

Both stationary and mobile energy systems are governed by both continuous behavior and discrete switching events. Examples of switching include the turning on and off of

solenoid valves in fluid-thermal systems, reconfiguration of multi-evaporator vapor compression systems [21], addition and removal of suppliers or consumers in an electric grid [22], switching of electrical buses and cooling modes in advanced vehicles [23–25], and switching between charging, discharging, and idle modes of a battery. This switching necessitates special attention in modeling and control design. Modeling approaches must accurately capture switching dynamics, while control designs must perform decision-making for both continuous and switched actuators. The presence of switching has significant implications to theoretical analysis as well. For example, even when it can be proved that each mode of a switched system is stable in independence, this alone does not guarantee that the system will remain stable under switching [26].

1.1.3 Traditional and Emerging Practices

1.1.3.1 Dynamic Modeling and Analysis

In response to the shifting paradigms of vehicle energy systems, a variety of complimentary efforts have focused on employing transient analyses to improve upon traditional steady-state methods. For advanced aircraft, these include integrated "tip-to-tail" dynamic modeling [12, 27, 28], characterization of component performance under transient operation [29, 30], design optimization [31–33], improved architectures for thermal management [19, 34], and enhanced mission planning [35]. A particularly strong trend has been the development of simulation-based toolboxes for the dynamic modeling of energy systems. In addition to facilitating system design, these can serve as a simulated plant on which to evaluate candidate control strategies. Examples of such toolboxes for the dynamic modeling of air-conditioning and refrigeration systems in buildings and aircraft include Thermosys [36] and the AFRL Transient Thermal Modeling and Optimization (ATTMO) [37] toolbox. The Thermal Management System (TMS) [38] toolbox includes dynamic models for a variety of aircraft thermal systems and components, while the PowerFlow [39] and Aerospace Power System (APS) [1] toolboxes include models spanning the electrical, thermal, mechanical, hydraulic, and pneumatic domains of aircraft energy systems.

Each of the above toolboxes consists of individual component models (e.g., a library of MATLAB/Simulink blocks) that can be interconnected to form complete systems. This modularity facilitates the sizing and validation of individual components within candidate architectures and permits a wide range of configurations and system scales to be implemented in simulation. However, extracting the underlying model equations and structure of coupling for implementation in model-based control can be an arduous task. Black box system identification techniques are often applied for these purposes, however these may not preserve the physical interpretation of the system states and must be repeated each time a change is made to the system architecture or any component parameter. This motivates the need for modeling approaches that retain the modularity of the above toolboxes, but also directly produce a system of equations (e.g., a state space representation) that can be readily employed in model-based control frameworks.

1.1.3.2 Conventional Control Approaches

Accounting for dynamic coupling within vehicle energy systems is a particularly exigent challenge in the field of controls. The complexity, scale, and multi-timescale nature of vehicle energy systems often render centralized control intractable due to computational limitations and communication bandwidth requirements [40]. Historically, decentralized, or "siloed," system design and control approaches have been employed to decompose the overall energy management problem into more manageable sub-problems, with limited consideration of the dynamic interactions between systems [1, 11, 41, 42]. Single-input, single-output (SISO) control approaches provide regulation about predetermined steady-state operating conditions, while logic-based decision trees govern switching and address potential fault scenarios. Coupling between subsystems, components, or domains is treated as a disturbance to which controllers and setpoints are made robust through extensive tuning to ensure constraint satisfaction. This practice necessitates conservative designs and can be time consuming and expensive to verify and validate, as well as significantly sub-optimal due to a lack of overall coordination in decision-making across the vehicle. Consequently, legacy system integration and control design methods have increasingly become the limiting factors of energy sys-

tem capability as electrification accelerates, and cannot accommodate the requirements of next-generation systems [11, 41]. Therefore, approaches must be developed that more closely integrate the operation of propulsion, power, and thermal management systems.

1.1.3.3 Model Predictive Control of Vehicle Energy Systems

Model Predictive Control (MPC) can explicitly consider constraints on both actuator inputs and states, allowing operational limits to be directly considered in optimizing system performance. In contrast to regulation-based control approaches that would seek to track a reference value, MPC can freely operate within constraints in minimizing an objective function. Furthermore, by continually evaluating over a time horizon into the future, MPC can leverage preview of upcoming disturbances and operational requirements to improve performance. For example, this allows thermal storage elements to be pre-cooled in advance of large heat loads, peak-shaving thermal transients and reducing violations of upper bounds on temperatures [43]. In the predictive control of aircraft energy systems, this preview information can be extracted from weather data and knowledge of the mission or flight plan. The emergence of connected and autonomous technologies for other vehicle classes will allow traffic conditions and path planning to also be leveraged as preview information in predictive energy management strategies [44].

Centralized MPC has been implemented throughout the literature for the control of specific subsystems and/or physical domains of vehicle energy systems. Examples include hybrid hydraulic and hybrid electric automotive powertrains [45, 46], vapor compression systems [47], aircraft fuel thermal management systems [48, 49], aircraft propulsion and electrical power systems [50–52], and all-electric ship propulsion and electrical power systems [53].

1.1.3.4 Hierarchical MPC of Vehicle Energy Systems

Centralized MPC is generally not tractable for control over all aspects of a vehicle energy system due to the problem scale and limited on-board computational power. Furthermore, leveraging the predictive capabilities of MPC to optimize slow dynamics becomes more computationally expensive as the range of dynamic timescales present in the system widens.

This is because control of the fast dynamic states requires a fast update rate, necessitating that the prediction horizon consist of many steps to extend long enough into the future to capture slow transients.

The shortcomings of conventional control approaches and intractability of centralized MPC has motivated the development of decentralized, distributed, and hierarchical MPC frameworks for a variety of large scale applications. These include building thermal systems [54], water distribution networks [55, 56], chemical process networks [57], smart grids and microgrids [58–60], and aircraft electrical and propulsion systems [61]. While this dissertation focuses specifically on hierarchical MPC of vehicle energy systems, broader reviews are available in the literature, including [40, 62].

Hierarchical MPC frameworks are particularly suited to systems characterized by dynamics spanning a wide range of timescales (e.g., time constants spanning multiple orders of magnitude), as is the case for multi-domain vehicle energy systems. Figure 1.2 provides a notional example of a hierarchical control framework. A single control formulation in the "vehicle level" at the top of the hierarchy is responsible for coordinating overall behavior at a relatively slow update rate, with a time horizon sufficiently long to optimize slow transients. However, the slow update rate prohibits this level from governing faster dynamics, compensating for faster timescale model error and disturbances, or leveraging inter-sample updates to preview information. Attempting to correct this by increasing the update rate of the vehicle level controller while maintaining the same time horizon may not be computationally tractable, as this would increase the number of prediction steps to be solved by the optimization program while decreasing the duration between consecutive updates in which the program must be solved for real-time implementation. Therefore, the vehicle level instead sends objectives and preview information to the "system level" below, where the overall control problem is decomposed into multiple sub-problems and solved with a faster update rate, allowing faster dynamic behavior to be governed. This communication and decomposition continues in the formation of additional levels below, with controllers at each level governing smaller portions of the system at faster update rates than in the level above. Controllers at the lower levels of the hierarchy interface directly with the plant to issue input commands. In addition to this downward flow of information, output feedback from sensors

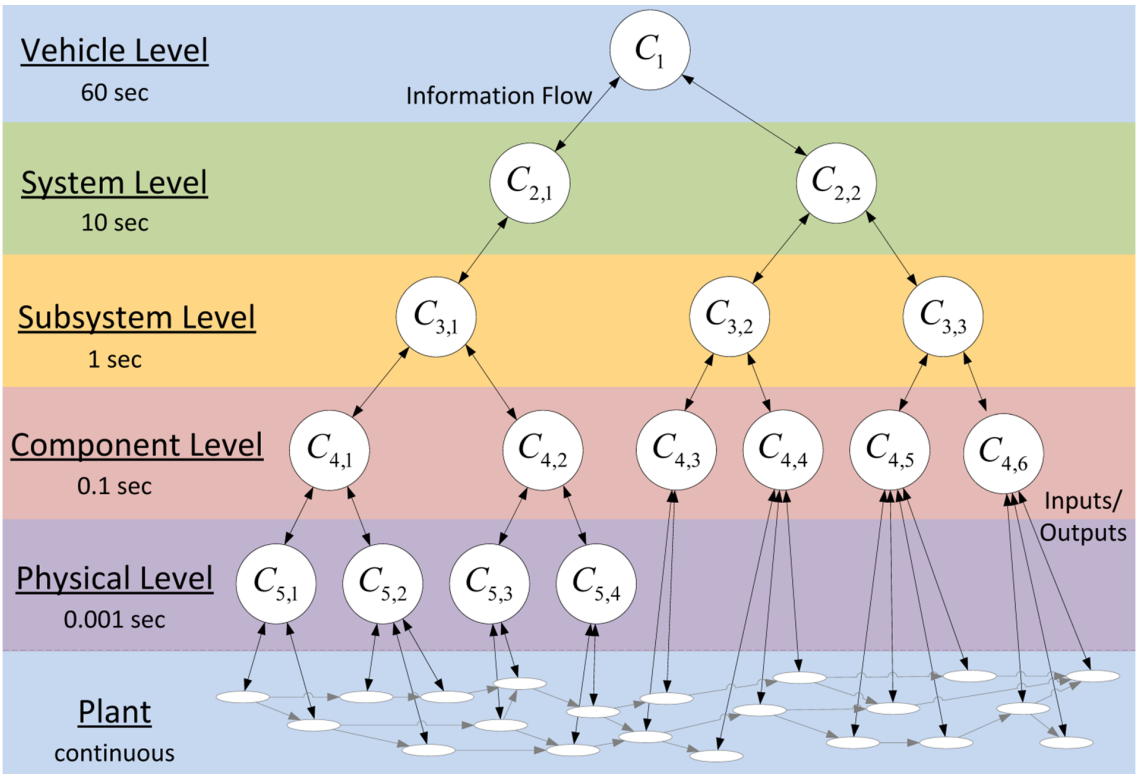

Figure 1.2: Notional hierarchical control framework with sample controller update intervals for each level. Modified from [63].

throughout the system is communicated up the hierarchy. While the notional example in Fig. 1.2 has five levels, in general it is possible for hierarchies to have more or fewer levels, as well as different numbers of controllers within each level from the second level down.

A defining feature of the hierarchical structure is that controllers at the same level do not directly exchange information with each other. This greatly reduces the communication requirements of the hierarchy and precludes the use of the iterative or cascaded methods from distributed control approaches that can be computationally expensive or inhibit parallelization [55, 62]. However, accounting for coupling within each level is necessary to achieve effective coordination of the system and avoid the conservatism inherent in many decentralized control approaches [40]. Therefore, in lieu of direct communication within each level, coordination among each level's controllers is achieved using communication from the levels above, as discussed further in later chapters of this dissertation.

The ability of hierarchical control to coordinate multi-timescale coupling throughout a system has inspired its application to vehicle energy systems, especially for advanced aircraft.

The majority of these efforts have focused on the control of a single physical domain, such as the thermal domain [25, 63, 64] or electrical domain [23]. Multi-domain efforts such as [1, 65, 66] have been implemented in simulation only. This motivates the experimental application in Chapter 3 of this dissertation, which evaluates the effectiveness of hierarchical control subject to the model error, communication delays, computational limitations, and other phenomena that can occur in experimental implementation.

A significant absence in the efforts referenced above is the development of a generalizable hierarchical control framework for vehicle energy management that can govern both continuous and switched behavior. Switching decisions are often restricted to the top level of the hierarchy, as in [23], or integrated ad hoc into a framework constructed primarily to govern continuous behavior. Therefore, a focus of this dissertation is the formulation and demonstration of a switched hierarchical MPC framework for vehicle energy management. This includes both experimental application and the development of approaches for ensuring stability and robustness under switching. For further literature review associated with the specific contributions of this dissertation, readers are referred to the relevant discussion within each chapter.

1.1.4 Experimental Testbeds

This dissertation balances analysis with implementation. To do so successfully, it is key that experimental application be used to evaluate the validity of assumptions in analytical constructions. Where assumptions are found to be limiting, the analysis can be augmented to better support the target applications.

To validate both modeling and control approaches for energy systems, laboratory-scale experimental testbeds have been developed across a range of application areas. Examples include the vapor compression refrigeration testbeds of [21, 36, 67], the hydraulic hybrid vehicle testbed of [68], the aircraft fuel thermal management system (FTMS) testbeds of [43, 69], the shipboard chilled water and electrical power distribution testbed of [70], and the shipboard electric propulsion and thermal management testbed of [71]. For convenience, cost, or safety, these testbeds are often not constructed at full scale, and may take further

departures from their real-world equivalents. For example, the FTMS testbed in [43, 49] was constructed at approximately $1/12^{\text{th}}$ scale and uses deionized water as a working fluid rather than jet fuel. However, by including the same physical domains and dynamic interactions as their target applications, these testbeds play a crucial role in demonstrating the efficacy of candidate modeling approaches and control frameworks for real-world implementation. This is especially true when the modeling and control methods are rooted in first-principles analysis that can be applied regardless of the system scale or specific architecture, as is the case in this dissertation.

1.2 Research Objectives

1.2.1 Desired Capabilities

To facilitate continued electrification, control frameworks for vehicle energy systems must adopt a more holistic approach, communicating within a network to coordinate energy management decisions across subsystems, components, timescales, and physical domains. This decision-making must extended to systems governed by both continuous and switched behavior. Theoretical tools must be developed to ensure the stability and robustness of these frameworks, while experimental validation bridges the gap between theory and practice, certifying their applicability as enabling technology for the continued electrification of vehicles.

With these needs established, the primary objective of this dissertation is the development and demonstration of model-based control strategies for vehicle energy systems that exhibit the following capabilities:

- *Applicable* to the energy systems of a wide range of vehicle classes and architectures by employing generalizable modeling and analysis approaches.
- *Scalable* to systems with many components, actuators, and dynamic states by using modular modeling approaches and control frameworks, as well as computationally efficient and parallelizable algorithms.

- *Hybrid* in the sense that actuators and system representations characterized by both continuous and switched behavior are supported.
- *High performance* by explicitly leveraging knowledge of system mission, constraints on inputs and states, and coupling within and between physical domains.

1.2.2 Dissertation Scope

To achieve the capabilities described above, this dissertation provides contributions in the following areas:

1. Formulation and validation of a modular, first-principles, and control-oriented modeling approach that captures continuous and switched dynamics of energy systems across multiple physical domains.
2. Development of a switched hierarchical MPC framework that coordinates among subsystems, timescales, and physical domains of energy systems, improving their overall capability, safety, and efficiency as compared to traditional approaches.
3. Demonstration of the hierarchical control framework on an experimental testbed representative of vehicle energy systems, highlighting improvements in key figures of merit as compared to traditional control approaches.
4. Derivation and application of approaches for ensuring the stability and robustness of the switched hierarchical control framework.

Figure 1.3 provides a visual outline of the relationships between these contributions and indicates the chapters of this dissertation in which each can be found.

1.3 Organization of Dissertation

The remainder of this dissertation is organized as follows. Chapter 2 presents a dynamic graph-based modeling approach that inherently captures the structure of coupling within

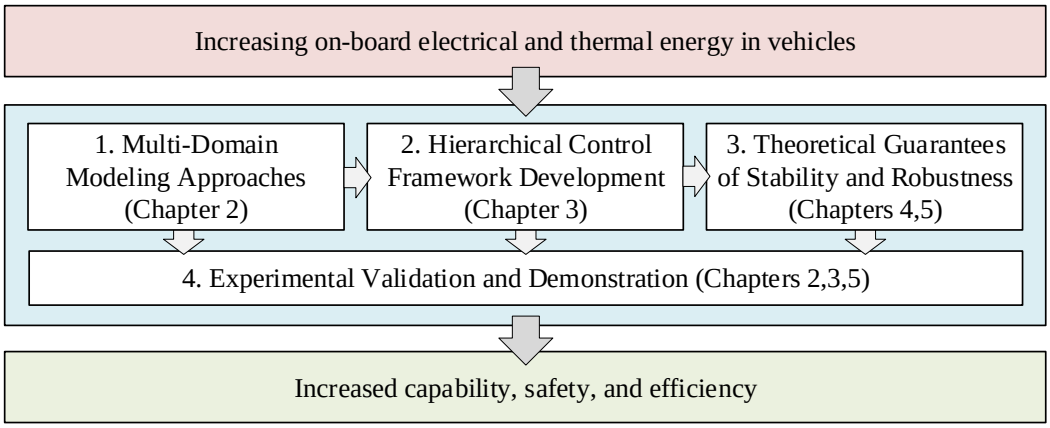

Figure 1.3: Outline of dissertation contributions and relevant chapters.

an energy system. Derived from first, principles, this approach can be applied to capture interactions within and between multiple physical domains. To demonstrate the application of the approach, hydraulic and thermal graph-based models for components of a single phase fluid-thermal system are presented and validated using an experimental testbed.

Chapter 3 demonstrates the efficacy of switched hierarchical MPC frameworks for vehicle energy systems. A hierarchical control framework is constructed and applied to a hardware-in-the-loop experimental testbed representative of a scaled aircraft electro-thermal system. In comparison to a baseline controller consisting of traditional SISO loops and logic-based decision trees, the hierarchical control framework is shown to achieve significantly greater thermal and electrical performance in key figures of merit associated with capability, efficiency, and safety.

Chapter 4 proposes a passivity-based approach to guaranteeing stability of switched graph-based systems governed by decentralized, distributed, or hierarchical MPC frameworks. This is achieved by augmenting controllers of a framework with a local, passivity-based constraint. A numerical example demonstrates the efficacy of the proposed approach on a fluid tank system controlled by a hierarchical framework.

Chapter 5 presents a two-level switched hierarchical control framework that leverages properties of cooperative systems to guarantee satisfaction of state and input constraints for a class of graph-based dynamic models. The overall objective of the control design is to track a reference trajectory for the energy transferred to a system, subject to bounded uncertainty

in the exogenous disturbances associated with its ability to dissipate energy to sinks. The proposed approach is demonstrated in both simulation and experimental application using a fluid-thermal system with dynamic behavior representative of a scaled aircraft fuel thermal management system.

Lastly, Chapter 6 concludes this dissertation by summarizing key contributions and suggesting future research directions.

1.4 Notation

The symbol $\mathbb{R}$ denotes the set of real numbers, and $\mathbb{R}_+$ denotes the set of non-negative real numbers. The set of all integers in the range of N to P is written as $[N{:}P]$. The notation $[x_i]$ indicates a vector $x \in \mathbb{R}^N$ of elements x_i, $i \in [1{:}N]$. Similarly, the notation $[m_{ij}]$ indicates a matrix $M \in \mathbb{R}^{N \times P}$ of elements m_{ij}, $i \in [1{:}N]$, $j \in [1{:}P]$. Lower case superscripts are used throughout this dissertation in the naming of variables, while upper case text or numbers in superscripts indicate mathematical functions, such as a transpose or exponent. For example, x^t is the vector of sink states of a graph-based model, while M^T is the transpose of a matrix M. The eigenvalues of a square matrix $A \in \mathbb{R}^{N \times N}$ are written as $\lambda_i(A)$, $i \in [1{:}N]$. The real components of these eigenvalues are denoted as $\mathrm{Re}(\lambda_i(A))$. For a function $f(x)\colon \mathbb{R}^N \to \mathbb{R}$, the zero set of $f(x)$ is denoted as $\mathcal{N}(f(x)) \equiv \{x | f(x) = 0\}$. Lastly, $\|x\|_Q \equiv x^T Q x$ for a vector $x \in \mathbb{R}^N$ and positive semi-definite matrix $Q \in \mathbb{R}^{N \times N}$.

Chapter 2

Graph-Based Modeling

2.1 Motivation and Background

2.1.1 Desired Model Features

Regardless of the physical domains governing their dynamics, many complex systems are fundamentally characterized by the storage, transport, and conversion of conserved quantities such as mass and energy. In this dissertation, such systems are termed "power flow systems." Examples in the literature include building thermal systems [54, 72, 73], water distribution networks [55, 56, 74], chemical process networks [57, 75], electrical power grids [22, 76], and energy systems in land, sea, and air vehicles [20, 45, 66, 71, 77]. As discussed in Section 1.1.2, the capability, safety, and efficiency of these systems can be strongly dependent on dynamic interactions both within and between physical domains. Furthermore, these interactions can occur across a wide range of timescales and be governed by both continuous and switched behaviors. As a result, control-oriented modeling approaches for these systems must include the following features to successfully facilitate model-based control design:

- *Modular:* Complete system models should be formed by interconnecting models of individual components and subsystems. This greatly reduces the time required to generate and validate system models, and permits models for a wide range of configurations and system scales to be readily formulated.

- *Physical domain and timescale agnostic:* Dynamic interactions within and between all relevant physical domains and timescales should be captured within a unified modeling framework.

- *Structure-preserving:* In addition to capturing coupling within the system, the modeling approach should make explicit the underlying structure of that coupling. The utility of this feature for system analysis and control design is a recurring theme of this dissertation, particularly in Section 2.5 and Chapters 4-5.

- *Hybrid:* The modeling approach must capture both continuous and switched behavior.

- *Flexible in representation:* The required model fidelity may vary with application and computational resources. As such, the modeling approach must be flexible to accommodate various level of complexity in representation. This includes allowing the number of dynamic states used to model a given component or subsystem to be varied, and allowing the assumed form of model equations (linear, bilinear, nonlinear, etc.) to be adjusted.

2.1.2 Modeling Approaches

As discussed in Section 1.1.3.1, recent approaches for the multi-domain dynamic modeling of vehicle energy systems have focused largely on the development of simulation-based toolboxes, and do not directly provide a system of equations that can be readily employed for model-based control. An alternative and powerful tool for deriving the governing equations of multi-domain systems is bond graph modeling [78], in which bonds represent power as a function of generalized effort and flow variables. Within each physical domain, these variables are associated with specific properties. For example, in the mechanical domain, the flow variable is velocity and the effort variable is force. In the electrical domain, the flow variable is current and the effort variable is voltage. While bond graph modeling is a valuable tool for simulation and representation, a shortcoming of this approach is that the assembly of component and subsystem models to form system models can result in excessively complex structures [79]. In addition, bond graph models do not explicitly convey the underlying structure of coupling within the system, and as a result are not often used to support analysis of the model for tasks such as model reduction and clustering. As shown in

Section 2.5 and Chapters 4-5, capturing this structure explicitly can be valuable to system analysis and control design for energy systems.

In this dissertation, a dynamic graph-based modeling approach derived from conservation equations is employed to achieve the desired features of Section 2.1.1. Graph-based approaches have been shown to facilitate the control-oriented dynamic modeling and analysis of power flow systems in a number of applications, including chemical plants [80–82], building thermal systems [83, 84], fluid flow and data transmission networks [85], and information flow in power electronics [86]. The specific approach in this dissertation has been applied to model vehicle energy systems across the hydraulic, electrical, thermal, and mechanical domains [64, 66, 69, 77].

The remainder of this chapter is organized as follows. Section 2.2 describes the generic graph-based modeling approach employed throughout this dissertation. Section 2.3 demonstrates the application of this approach by presenting both nonlinear and linearized models of the hydraulic and thermal domains of fluid-thermal components. Section 2.4 describes a fluid-thermal testbed and demonstrates with experimental validation that graph-based models of individual components can be combined to accurately represent complete systems. While the primary use of graph-based modeling in this dissertation is for model-based hierarchical control and supporting analysis, Section 2.5 briefly discusses further applications to model decomposition and system design optimization. Section 2.6 revisits the desired model features of Section 2.1.1 to explain how the examples in this chapter demonstrate achievement of these goals. Section 2.7 provides a concluding summary of the chapter.

2.2 Generic Approach

In this dissertation, the structure of interconnections of a system under study is described by the oriented graph $\mathbf{G} = (v, e)$ of order N_v with vertices $v = [v_i], i \in [1{:}N_v]$, and size N_e with edges $e = [e_j], j \in [1{:}N_e]$. As shown in the notional graph example of Fig. 2.1, each edge e_j is incident to two vertices and indicates directionality from its *tail* vertex v_j^{tail} to its *head* vertex v_j^{head}. The set of edges directed into vertex v_i is given by $e_i^{head} = \{e_j | v_j^{head} = v_i\}$, while the set of edges directed out of vertex v_i is given by $e_i^{tail} = \{e_j | v_j^{tail} = v_i\}$ [87].

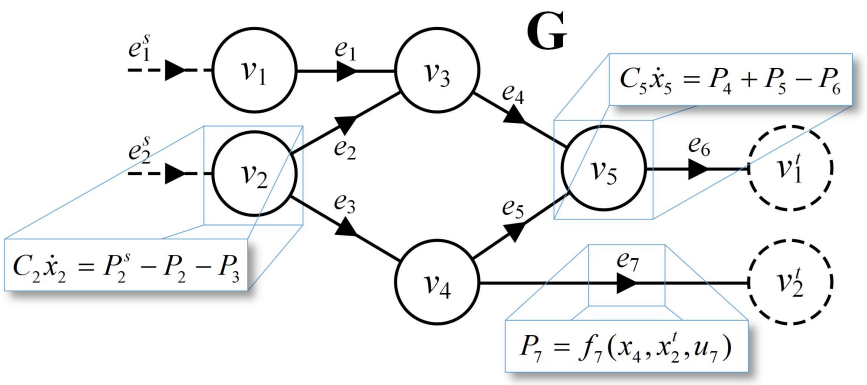

Figure 2.1: Notional example to demonstrate key features of the graph-based modeling approach. Dashes indicate states and power flows that are exogenous to system **S**. Here, $N_v = 7, N_e = 7, N_s = 2$, and $N_t = 2$. Modified from [64, 69].

The dynamic model **S** resulting from graph **G** is derived from application of conservation equations. Each vertex v_i is assigned a assigned dynamic state x_i associated with storage of a conserved quantity of interest. Similarly, each edge e_j is assigned a value P_j describing the rate of transfer of the conserved quantity between adjacent vertices. Adopting the terminology used when conservation of energy is applied, these edge transfer rates will often be referred to as "power" or "power flow" in this dissertation. The orientation of each edge indicates the convention assigned to positive power flow, from v_j^{tail} to v_j^{head}. Therefore, the dynamics of each state of **S** satisfy the conservation equation:

$$C_i \dot{x}_i(t) = \sum_{\{j | e_j \in e_i^{head}\}} P_j(t) - \sum_{\{j | e_j \in e_i^{tail}\}} P_j(t), \tag{2.1}$$

where $C_i > 0$ is the storage capacitance of the vertex. Physically, (2.1) states that the rate of storage in the vertex is equal to the total power flow into the vertex minus the total power flow out of the vertex.

The power flow associated with each edge of graph **G** is a function of the states of the two vertices to which it is incident and may also be a function of an input signal u_j. In this dissertation, this is represented using a generic function f_j of three arguments:

$$P_j = f_j\left(x_j^{tail}, x_j^{head}, u_j\right). \tag{2.2}$$

Fig. 2.1 includes examples of (2.1)-(2.2) as applied to the example graph.

In addition to capturing the exchange of a conserved quantity within the system under study, the modeling approach must account for exchange with entities external to the system. Sources to graph $\mathbf{G}$ are modeled by source edges $e^s = [e^s_j]$, $j \in [1{:}N_s]$ with associated power flows $P^s = [P^s_j]$, which are treated as disturbances to the system that may come from neighboring systems or the environment. Therefore, edges belonging to e^s are not counted among the edges e of graph $\mathbf{G}$, and transfer rates in P^s are not counted among the internal transfer rates P of system $\mathbf{S}$. The vector of all vertices of $\mathbf{G}$ that are incident to a source edge is denoted as v^s, with x^s as the corresponding states of $\mathbf{S}$.

Sinks of graph $\mathbf{G}$ are modeled by sink vertices $v^t = [v^t_j]$, $j \in [1, N_t]$ with associated states $x^t = [x^t_j]$. The sink vertices are counted among the vertices v of graph $\mathbf{G}$, but the sink states x^t are not included in the state vector x of system $\mathbf{S}$. Instead, the sink states x^t are treated as disturbances to the system associated with neighboring systems or the environment. Each edge may be incident to at most one sink vertex.

To describe the structure of edge and vertex interconnections of a graph, the incidence matrix $M = [m_{ij}] \in \mathbb{R}^{N_v \times N_e}$ is defined as:

$$m_{ij} = \begin{cases} 1 & \text{if } v_i \text{ is the tail of } e_j, \\ -1 & \text{if } v_i \text{ is the head of } e_j, \\ 0 & \text{else.} \end{cases} \tag{2.3}$$

M can be partitioned as:

$$M = \begin{bmatrix} \bar{M} \\ \underline{M} \end{bmatrix} \text{ with } \bar{M} \in \mathbb{R}^{(N_v - N_t) \times N_e}, \tag{2.4}$$

where the indexing of edges is ordered such that $\bar{M}$ is a structural mapping from power flows P to states x, and $\underline{M}$ is a structural mapping from P to sink states x^t.

The structural mapping from source power flows P^s to x is similarly given by $D = [d_{ij}] \in \mathbb{R}^{(N_v - N_t) \times N_s}$, where:

$$d_{ij} = \begin{cases} 1 & \text{if } v_i \text{ is the head of } e_j^s, \\ 0 & \text{else.} \end{cases} \tag{2.5}$$

For example, M and D for the graph in Fig. 2.1 are given by:

$$M = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ -1 & -1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & -1 & -1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix}, \quad D = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}, \tag{2.6}$$

where $\bar{M}$ is given by the top five rows of M, and $\underline{M}$ is given by the bottom two rows of M.

Following from the conservation equation for each vertex (2.1) and the above definitions of $\bar{M}$ and D, the dynamics of system $\mathbf{S}$ are given by:

$$\mathbf{S}: \quad C\dot{x}(t) = -\bar{M}P(t) + DP^s(t), \tag{2.7}$$

where $C = diag([C_i])$ is a diagonal matrix of the capacitances of the states in $\mathbf{S}$.

Following from (2.2), the vector of all power flows P in $\mathbf{S}$ is given by

$$P = F(x, x^t, u) = [f_j(x_j^{tail}, x_j^{head}, u_j)]. \tag{2.8}$$

2.3 Domain-Specific Modeling

The generic graph-based modeling approach of Section 2.2 has been successfully applied to represent dynamics of a wide range of physical domains. Hydraulic and thermal modeling of single phase fluid-thermal components is found in [69], which includes detailed derivations

and experimental validation of graph-based models for a fluid reservoir (also referred to as a fluid tank in this dissertation), flow split/junction, pump, pipe, cold plate, and liquid-to-liquid brazed plate heat exchanger. Further validation of these components is found in [64]. A graph-based model capturing temperatures throughout a multi-level DC/AC inverter is presented and experimentally validated in [88]. Graph-based models for turbomachinery and electrical devices are found in [77], which includes an air cycle machine (ACM), electrical generator, bus, constant power load, constant current load, and constant impedance load. Electrical and mechanical graph-based models for components of a hybrid unmanned aerial vehicle are found in [89], including a battery, DC and AC electric machines, DC/DC converter, electrical bus, and vehicle dynamics. Graph-based models for components of an electric automotive powertrain are found in [66], which includes a battery, DC electric machine, DC/DC converter, vapor compression system (VCS), and cabin thermal model.

While a complete exposition of all the components listed above falls outside the scope of this dissertation, the hydraulic and thermal models for the fluid-thermal components of [69] are summarized below to demonstrate the practical application of the modeling approach and provide context for the experimental validation that follows in Section 2.4. Beyond the focus on heat and mass transfer in this chapter, Section 3.3.2 provides further details on graph-based modeling in the electrical domain.

Fig. 2.2 shows the hydraulic (left) and thermal (right) graphs for key components of single-phase fluid-thermal systems. Assemblies of these component graphs can be used to represent thermal management systems for applications including aircraft Fuel Thermal Management Systems (FTMSs), liquid-cooled automotive powertrains, and server farms [19, 33, 48]. The hydraulic graphs are derived from application of conservation of fluid mass, while the thermal graphs are derived from application of conservation of thermal energy. Dashes indicate sources or sinks of each component, representing signals originating from neighboring components or disturbances. For hydraulic graphs, vertices represent dynamic states of pressure, while edges represent the rate of mass transfer between vertices. For thermal graphs, vertices represent dynamic states of temperature, while edges represent thermal power flow between vertices due to convection in heat exchangers or fluid transport. Table 2.1 summarizes the quantities associated with each element of the graph-based approach, including

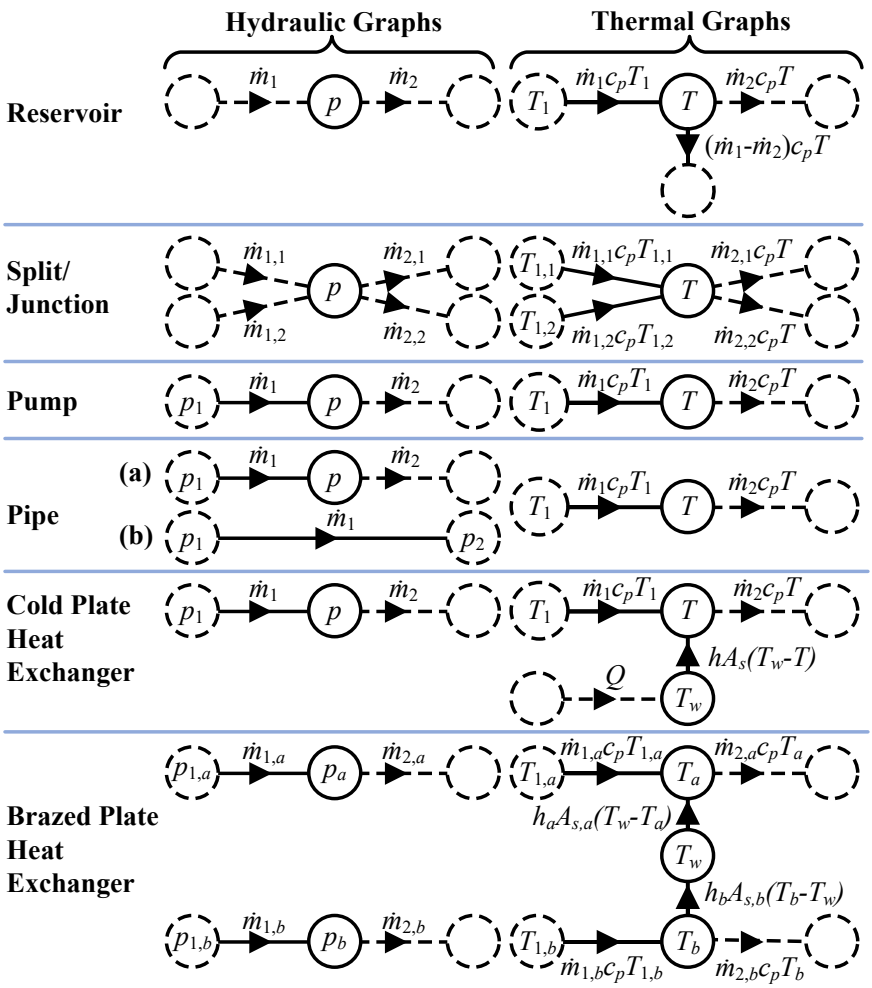

Figure 2.2: Hydraulic and thermal graphs for fluid-thermal component models. Modified from [64, 69].

in the generic sense of Section 2.2 and in the specific hydraulic and thermal domains of Sections 2.3.1 and 2.3.2, respectively.

2.3.1 Hydraulic Modeling

For notational clarity, a superscript m denotes some capacitances, functions, matrices, and inputs associated with hydraulic graphs. The reader is referred to [69] for a detailed derivation of the model equations that follow.

For all hydraulic vertices except those of a reservoir, the hydraulic capacitance is given by $C_i^m = V_i\rho/E$, where V_i is the fluid volume in the component and both the fluid density ρ

Table 2.1: Summary of quantities in graph-based models of fluid-thermal components.

Generic graph, $\mathbf{G}$	Hydraulic graph, $\mathbf{G}^m$	Thermal graph, $\mathbf{G}^r$
Conserved quantity	Fluid mass	Thermal energy
Vertex storage state, x	Pressure, p	Temperature, T
Edge transfer rate, P	Fluid mass flow rate, $\dot{m}$	Thermal power flow, P^r
Edge input, u	Actuator effort, u^m	Fluid mass flow rate, $\dot{m}^r$

and the bulk modulus E are assumed to be constant across all components. For reservoirs, $C_i^m = A_{c,i}/g$, where $A_{c,i}$ is the reservoir cross sectional area and g is the gravitational constant.

Following from (2.2), the fluid mass flow rate of each edge of the graph is given generically by $\dot{m}_j = f_j^m\left(p_j^{tail}, p_j^{head}, u_j^m\right)$. For all hydraulic edges except those of a pump, the fluid mass flow rate $\dot{m}_j$ is given specifically by:

$$\dot{m}_j = \rho A_{c,j}\sqrt{\frac{2\left(p_j^{tail} - p_j^{head} + \rho g \Delta h_j\right)}{\rho\left(k_j \frac{L_j}{D_j} + K_{L,j}\right)}}, \tag{2.9}$$

where L_j, D_j, and $A_{c,j}$ are the fluid flow length, diameter, and cross sectional area of the component, respectively, Δh_j is the height difference between the inlet and outlet flow, k_j is the friction factor, and $K_{L,j}$ is the minor loss coefficient. For pumps, the fluid mass flow rate is given by:

$$\dot{m}_j = \rho A_{c,j}\sqrt{2g\left(H_j - \frac{p_j^{head} - p_j^{tail}}{\rho g}\right)}. \tag{2.10}$$

Here, the pump head H_j is determined using an empirical map as a linear function of pump effort u_j^m and the pressure differential across the pump:

$$H_j = \alpha_{1,j} + \alpha_{2,j}\left(p_j^{head} - p_j^{tail}\right) + \alpha_{3,j} u_j^m, \tag{2.11}$$

where $\alpha_{1,j}$, $\alpha_{2,j}$, and $\alpha_{3,j}$ are constants. The pump effort may correspond to, for example, the percent duty cycle of pulse width modulation (PWM) of the power supply to an electric pump.

When hydraulic graphs of multiple components are interconnected to represent a system, the hydrodynamics can be represented in the form of (2.7)-(2.8). The fluid system configuration used for demonstration in Section 2.4 consists of closed loops such that fluid mass does not enter or exit the system. Therefore there are no source edges, and so in the notation of the hydraulic graph variables, (2.7) reduces to:

$$C^m \dot{p} = -\bar{M}^m \dot{m}, \tag{2.12}$$

while (2.8) is given by:

$$\dot{m} = F^m\left(p, u^m\right) = [f_j^m\left(p_j^{tail}, p_j^{head}, u_j^m\right)]. \tag{2.13}$$

Among the component models in Fig. 2.2, only edges corresponding to pumps require edge inputs. However, note that the pipe model can be modified to represent an on/off solenoid valve by multiplying its mass flow rate in (2.9) by an edge input u_j^m, where $u_j^m = 0$ when the valve is closed, and $u_j^m = 1$ when the valve is open.

2.3.2 Thermal Modeling

For notational clarity, a superscript r denotes some capacitances, functions, matrices, and inputs associated with thermal graphs. The reader is referred to [69] for a detailed derivation of the model equations that follow. Note that all the thermal component models in Fig. 2.2 except the reservoir assume that the mass of fluid stored in the component remains constant. However, this assumption can easily be removed by the addition of an edge that captures the thermal energy lost or gained due to a change in stored fluid mass, as shown for the reservoir.

For all vertices associated with a fluid temperature, the thermal capacitance is given by $C_i^r = \rho V_i c_p$, where the specific heat capacitance of the fluid c_p is assumed to be constant across all components. If the mass of fluid stored in the component is time-varying, as may be the case for a reservoir, then the volume used in calculating this thermal capacitance is also time-varying. For all vertices associated with heat exchanger wall temperatures,

$C_i^r = M_{w,i} c_{p,w,i}$, where $M_{w,i}$ is the mass of the wall and $c_{p,w,i}$ is the specific heat capacitance of the wall material.

Following from (2.2), the thermal power flow of each edge of the graph is given generically by $P_j^r = f_j^r \left(T_j^{tail}, T_j^{head}, \dot{m}_j^r\right)$. For advective thermal power flow due to fluid transport, this is given by:

$$P_j^r = \dot{m}_j^r c_p T_j^{tail}. \tag{2.14}$$

When fluid flows in a loop between two thermal elements, the advective power flow in each direction can be combined into a single "bidirectional advection" power flow, given by:

$$P_j^r = \dot{m}_j^r c_p \left(T_j^{tail} - T_j^{head}\right). \tag{2.15}$$

Convective power flow between fluid and the wall of heat exchangers is given by:

$$P_j^r = h_j A_{s,j} \left(T_j^{tail} - T_j^{head}\right). \tag{2.16}$$

Here, $A_{s,j}$ is the convective surface area and h_j is the heat transfer coefficient, which is typically an empirically-derived function [90, 91]. As in [64, 69], in this dissertation this function is assumed to be of the form:

$$h_j = \beta_{1,j} + \beta_{2,j} \dot{m}_j^r T_j^{head}, \tag{2.17}$$

where $\beta_{1,j}$ and $\beta_{2,j}$ are constants.

When thermal graphs of multiple components are interconnected to represent a system, the thermodynamics can be represented in the form of (2.7)-(2.8). In the notation of the thermal graph variables, (2.7) is given by:

$$C^r \dot{T} = -\bar{M}^r P^r + D^r P^{s,r}, \tag{2.18}$$

while (2.8) is given by:

$$P^r = F^r \left(T, T^t, \dot{m}^r\right) = [f_j^r \left(T_j^{tail}, T_j^{head}, \dot{m}_j^r\right)]. \tag{2.19}$$

2.3.3 Multi-Graph System Representation

The combined hydraulic and thermal dynamics of a system can be simultaneously modeled using two coupled graphs. As shown in Fig. 2.3, the mass flow rates calculated for edges of the hydraulic graph serve as inputs to the equations governing thermal power flow along edges of the thermal graph. The edges in $\mathbf{G}^m$ may not map one-to-one to the edges in $\mathbf{G}^r$. For example, a single mass flow rate may affect multiple edges of the thermal graph. It is also possible that some mass flow rate inputs to the thermal system are exogenous, and not modeled within the hydraulic graph. For example, this could include mass flow rates on the secondary side of heat exchangers by which thermal energy is transferred to and from neighboring systems. These external mass flow rates are denoted by $\dot{m}^{ext} = [\dot{m}_i^{ext}], i \in [1{:}N_{ext}]$ and treated as disturbances to the thermal model. The mapping from $\dot{m}^{ext}$ and the mass flow rates calculated within the hydraulic graph $\dot{m}$ to the mass flow rate inputs to the thermal graph $\dot{m}^r$ can be represented as:

$$\dot{m}^r = Z \begin{bmatrix} \dot{m} \\ \dot{m}^{ext} \end{bmatrix}, \tag{2.20}$$

where $Z \in \{0,1\}^{N_e^r \times (N_e^m + N_{ext})}$.

To capture the dynamics of pumps, including rate limits and time delays between each pump command $u_i^p, i \in [1{:}N_p]$ and the actual pump effort u_i^m input to the hydraulic graph, each u_i^m is paired with a single-input-single-output (SISO) system $\mathbf{S}_i^p$ as shown in Fig. 2.3. Each $\mathbf{S}_i^p$ models the dynamic effort state of the i^{th} pump u_i^m as a function of its commanded value u_i^p.

In this dissertation, pump states and inputs are expressed in units of percent duty cycle of PWM. The dynamic of each pump is modeled as a first-order response with time constant τ_i^p and delay ξ_i^p, given as a transfer function by:

$$u_i^m(s) = \frac{e^{-\xi_i^p s}}{\tau_i^p s + 1} u_i^p(s). \tag{2.21}$$

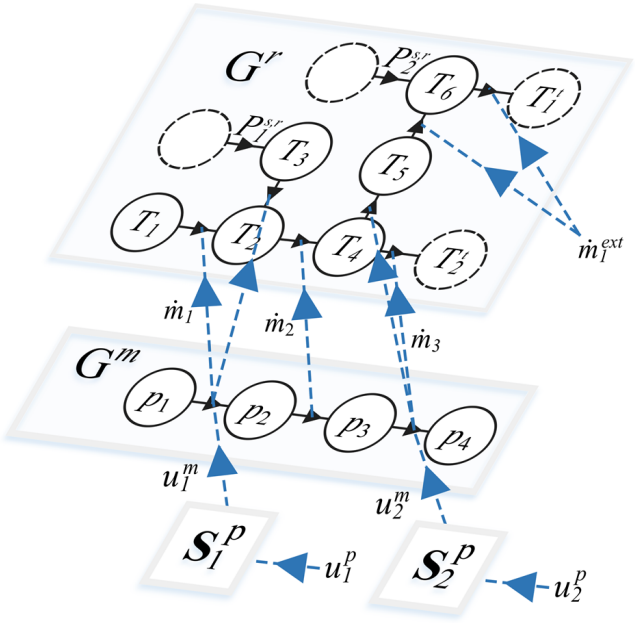

Figure 2.3: Interconnection between a hydraulic graph (middle) and a thermal graph (top), with pump dynamics (bottom) input to the hydraulic graph [64].

2.3.4 Hydraulic Linearization

In general, the graph-based models have a nonlinear form but satisfy the generic relationships of (2.1) and (2.2) for each vertex and edge. For control design, it is often useful to use a linear representation of the system dynamics. A benefit of the graph-based approach is that a linear model of the full system can be generated by individual linearization of each edge relationship, as shown for hydraulic graphs in this section and for thermal graphs in Section 2.3.5.

From (2.9)-(2.10), the nonlinear hydraulic mass flow rate equations for all components follow the general form:

$$\dot{m}_j = c_{1,j}\sqrt{c_{2,j} + c_{3,j}\left(p_j^{tail} - p_j^{head}\right) + c_{4,j}u_j^m}, \tag{2.22}$$

where the coefficients $c_{i,j}$ are constant for each i, j. Linearizing this expression about an equilibrium operating condition using a first-order Taylor Series gives linear mass flow rate

equations of the form:

$$\Delta\dot{m}_j = a_j^m \left(\Delta p_j^{tail} - \Delta p_j^{head}\right) + b_j^m \Delta u_j^m, \tag{2.23}$$

where a_j, b_j are constant coefficients and for a generic signal $y(t)$, $\Delta y(t) \equiv y(t) - y_0$ with y_0 as the equilibrium value of y about which the linearization is performed.

Following substitution into (2.12), the linear equations for a hydraulic system model are given by:

$$\dot{p} = A^m \Delta p + B^m \Delta u^m, \tag{2.24}$$

where

$$\begin{aligned} A^m &= -\left(C^m\right)^{-1} \bar{M}^m diag\left(\left[a_j^m\right]\right) \left(\bar{M}^m\right)^T, \\ B^m &= -\left(C^m\right)^{-1} \widetilde{\bar{M}}^m diag\left(\left[\tilde{b}_k^m\right]\right). \end{aligned} \tag{2.25}$$

Here, $\widetilde{\bar{M}}^m$ represents the sub-matrix of columns of $\bar{M}^m$ corresponding to edges associated with pumps, and $[\tilde{b}_k^m]$ is the vector of input coefficients for edges associated with pumps (i.e., edges k for which $c_{4,k} \neq 0$ in (2.22)).

The mass flow rates of the linearized hydraulic model are given by:

$$\Delta\dot{m} = V_1^m \Delta p + V_2^m \Delta u^m, \tag{2.26}$$

where

$$\begin{aligned} V_1^m &= diag\left(\left[a_j^m\right]\right) \left(\bar{M}^m\right)^T, \\ V_2^m &= [v_{jk}^m] \in \mathbb{R}^{N_e^m \times N_p}, \end{aligned} \tag{2.27}$$

and

$$v_{jk}^m = \begin{cases} b_j^m & \text{if } e_j \text{ is associated with pump } k, \\ 0 & \text{else.} \end{cases} \tag{2.28}$$

2.3.5 Thermal Linearization

From the power flow equations of (2.14)-(2.16) and the assumed expression for the heat transfer coefficient of (2.17), the nonlinear power flow flow equations for all components follow the general form:

$$P_j = c_{1,j}T_j^{tail} + c_{2,j}T_j^{head} + c_{3,j}T_j^{tail}\dot{m}_j^r + c_{4,j}T_j^{head}\dot{m}_j^r + c_{5,j}\left(T_j^{tail} - T_j^{head}\right)T_j^{head}\dot{m}_j^r, \quad (2.29)$$

where the coefficients $c_{i,j}$ are constant for each i, j. Linearizing this expression about an equilibrium operating condition using a first-order Taylor Series gives linear power flow equations of the form:

$$\Delta P_j^r = a_{1,j}^r \Delta T_j^{tail} + a_{2,j}^r \Delta T_j^{head} + b_j^r \Delta\dot{m}_j^r, \quad (2.30)$$

where $a_{1,j}^r$, $a_{2,j}^r$, and b_j^r are constant coefficients. Following substitution into (2.18), the linear equations for a thermal system model are given by:

$$\dot{T} = A^r \Delta T + B_1^r \Delta T^{out} + B_2^r \Delta\dot{m}^r + B_3^r \Delta P^{s,r}, \quad (2.31)$$

where

$$\begin{aligned} A^r &= -\left(C^r\right)^{-1}\bar{M}^r\left(\bar{M}_a^r\right)^T, \\ B_1^r &= -\left(C^r\right)^{-1}\bar{M}^r\left(\underline{M}_a^r\right)^T, \\ B_2^r &= -\left(C^r\right)^{-1} diag\left([b_j^r]\right), \\ B_3^r &= -\left(C^r\right)^{-1} D^r, \end{aligned} \quad (2.32)$$

and $M_a^r = [m_{ij}] \in \mathbb{R}^{N_v^r \times N_e^r}$ is a weighted incidence matrix for the thermal graph, with

$$m_{ij} = \begin{cases} a_{1,j}^r & \text{if } v_i \text{ is the tail of } e_j, \\ a_{2,j}^r & \text{if } v_i \text{ is the head of } e_j, \\ 0 & \text{else.} \end{cases} \quad (2.33)$$

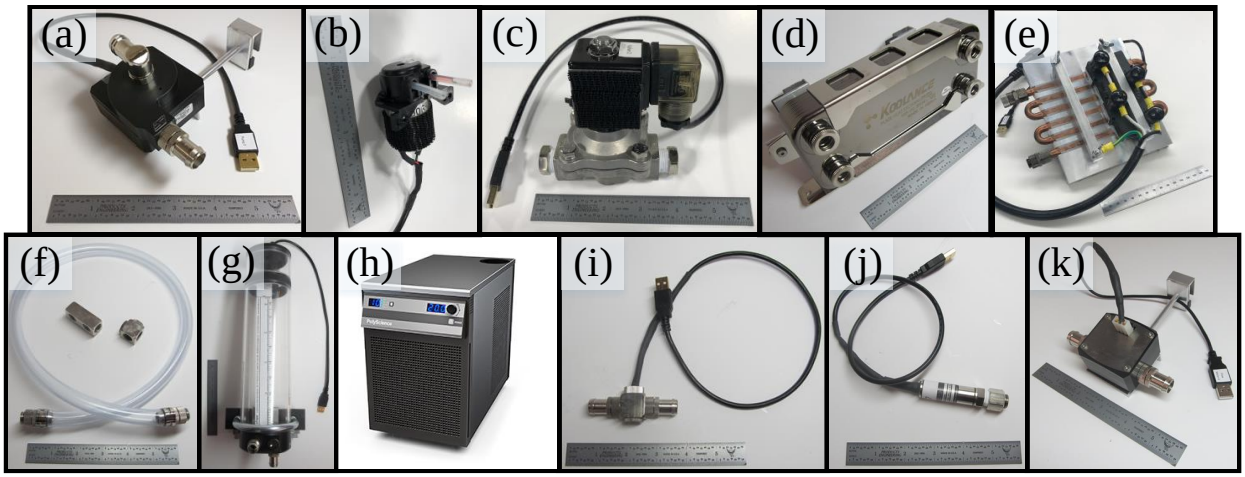

Figure 2.4: Individual components as labeled in Table 2.2, pictured with a 6" ruler for scale. Of the two types of temperature sensors, a fluid temperature sensor is pictured in (i). Modified from [64, 69].

2.4 Modeling Example and Validation

This section demonstrates how the graph-based models of fluid-thermal components in Section 2.3 can be assembled to represent a complete fluid-thermal system. This is supported by model validation with the experimental testbed described below.

2.4.1 Experimental Testbed Overview

This experimental testbed was developed to capture the key features of fluid-based thermal management systems on a laboratory scale. The use of flexible tubing and a slatted mounting stand allows the system architecture to be easily reconfigured to match a given application. Table 2.2 and Fig. 2.4 contain specifications and images of the components and sensors currently included in the testbed. The working fluid is an equal parts mixture of propylene glycol and water.

Centrifugal pumps are the primary fluid movers in the system. Speed is controlled via a PWM percent duty cycle, with less than 20% corresponding to a constant 1300 RPM, 65% and above corresponding to 4500 RPM, and a linear trend between. Peak power consumption of the pumps is 20 W with a peak efficiency of 35%. In addition to these centrifugal pumps, positive displacement pumps are used to drain fluid from the system. This emu-

lates behaviors like decreasing stored fuel mass in aircraft tanks due to fuel burn in engines. Solenoid valves are used to exert switched (on/off) flow control in fluid lines.

Table 2.2: Testbed component descriptions.

Component	Specifications	No. supported
(a) Centrifugal pump	· Swiftech MCP35X · 12 VDC, 1.5 A max, PWM controlled · 4.4 m max head · 17.5 LPM max flow · SparkFun ACS712 low current sensor	8
(b) Positive displacement pump	· Gikfun peristaltic pump · 12V DC, PWM controlled · 0.1 LPM max flow	2
(c) Solenoid valve	· WIC Valve 2SCW Series · Vacuum to 150 PSI operating pressure	8
(d) Heat exchanger	· Koolance HXP-193 · 12 plates · 4.0 kW at 5 LPM and 20°C inlet temperature difference	-
(e) Cold plate	· Wakefield-Vette 6-pass, 6" cold plate · Vishay LPS1100H47R0JB resistors, 47 Ω, 1100 W max power each · Crydom 10PCV2415 solid state relay · Echun Electronic Co. ECS1030-L72 non-invasive current sensor	4
(f) Pipe	· Koolance HOS-13CL · Clear PVC	-
(g) Reservoir	· Koolance 80x240mm · Acrylic · 8" eTape liquid level sensor	4
(h) Chiller	· Polyscience 6000 Series · Up to 2900 W at 20°C · -10°C to +70°C	2
(i) Temp. sensor	· Koolance SEN-AP008B (fluid) · Koolance SEN-AP007P (surface) · 10 KΩ thermistor	16
(j) Pressure sensor	· Measurement Specialties US300 · Up to 310 kPa gauge	7
(k) Flow rate sensor	· Aqua Computer High Flow · 0.5 to 25 LPM	8

Liquid-to-liquid brazed plate heat exchangers (HXs) transfer thermal energy between fluid loops in either a parallel-flow or counter-flow configuration.

Each cold plate (CP) consists of an aluminum plate with copper tubing passing through. Several 47 Ω resistive heaters are mounted to each CP and wired to a solid state relay actuating the heater power output. Up to four heaters can be mounted to each CP, however in this dissertation just two are used, allowing a maximum heat load of 1.7 kW to be applied to each CP.

The reservoirs act as thermal storage elements. A liquid level sensor inside each reservoir is used to calculate its liquid volume, which in turn can be used to calculate its thermal capacitance.

Two 1.5 HP (1.12 kW) industrial chillers act as thermal energy sinks. With variable temperature control from $-10°C$ to $70°C$, the chillers can emulate a wide range of sink conditions.

Infrared cameras were used to identify locations on the the HX and CP walls that closely represent the average wall temperature, at which surface temperature sensors are affixed. The infrared image in Fig. 2.5 shows CP1 and reservoir 1 of the example testbed configuration in Fig. 2.6. The cable for the CP1 wall temperature sensor leads from the center of the plate across its left side.

Sensors and actuators are connected to a National Instruments CompactDAQ, exchanging sensor measurements and actuator commands with National Instruments LabVIEW software on a desktop computer at a rate of 10 Hz. signals can be exchanged between LabVIEW and MATLAB/Simulink either by running the two programs simultaneously and communicating via the User Datagram Protocol (UDP), or by embedding MATLAB code in LabVIEW using a MATLAB script node.

2.4.2 Example Configuration Description

The testbed is pictured in Fig. 2.6 in an example configuration used for demonstration in this chapter. The corresponding schematic is shown in Fig 2.7. This configuration is notionally representative of a simplified aircraft fuel thermal management system (FTMS) into which

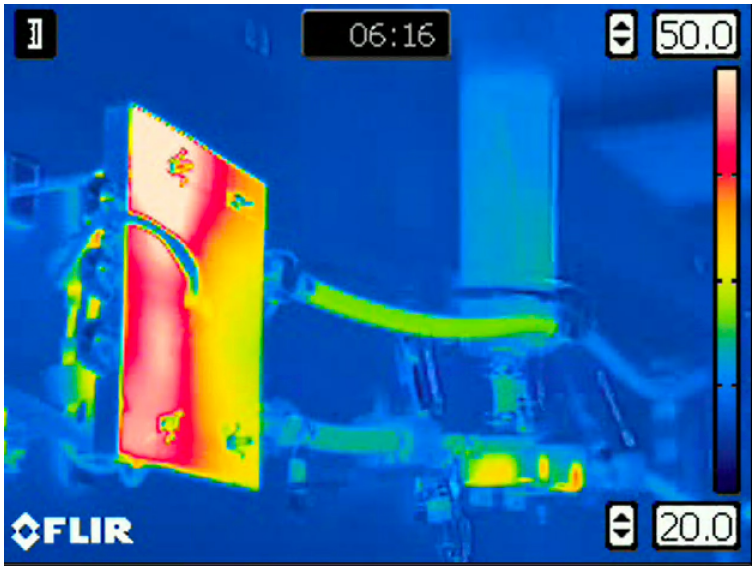

Figure 2.5: Infrared image of CP 1 and reservoir 1 from the example testbed configuration in Fig. 2.6 [64].

thermal energy from actuators, generators, engine oil, and other transient loads is absorbed, stored in liquid fuel, and rejected through transfer to neighboring systems or the environment [28].

The example configuration has eight pumps arranged in four sets of two. The 'a' and 'b' pumps of each set receive the same commands. Therefore, for notational convenience the two pumps in each set are referred to collectively. For example, pumps 1a and 1b are collectively termed "pump 1."

The secondary loop (identified as the left half of the system in Fig. 2.7) absorbs thermal energy from the heaters mounted to CP1, through which fluid is driven by pump 1. This loop has dedicated thermal storage available in reservoir 1, and the ability to exchange thermal energy across HX1 with fluid driven by pump 2.

The primary loop (identified as the right half of the system in Fig. 2.7), includes two parallel fluid flow paths out of reservoir 2. The path driven by pump 3 passes through HX1, exchanging thermal energy with the secondary loop. The path driven by pump 4 passes through CP2 and CP3, absorbing thermal energy produced by their their heaters. The two flow paths then junction and pass through HX2, by which thermal energy is transferred out of the system to the thermal sink.

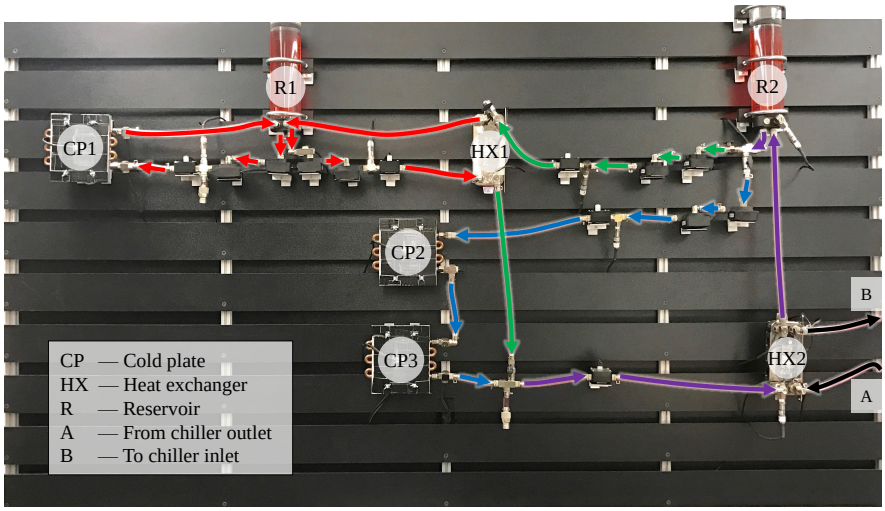

Figure 2.6: Example thermal-fluid testbed configuration for experimental validation [64].

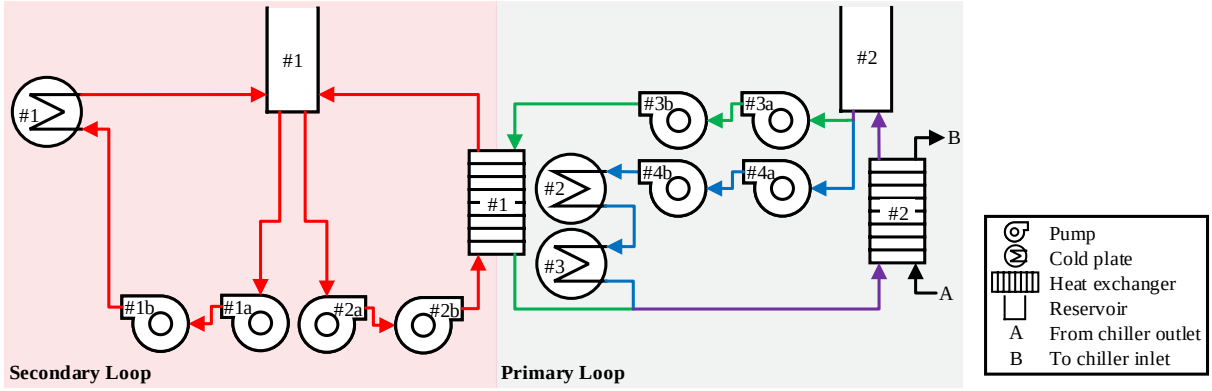

Figure 2.7: Schematic of example testbed configuration [64].

2.4.3 Graph-Based Representation of Example Configuration

The hydrodynamics of the example testbed configuration in Fig. 2.7 are represented by the system graph shown in Fig. 2.8, formed by interconnection of the individual hydraulic component graphs from Fig. 2.2. This hydraulic graph consists of 32 vertices and 34 edges, which in turn set the number of pressure states and mass flow rates in the corresponding graph-based hydraulic model.

Figure 2.9 shows the thermal graph for the example testbed configuration, formed by interconnection of the individual thermal component graphs from Fig. 2.2. The edges exiting the three leftmost dashed vertices indicate heat transfer from the resistive heaters to the CPs, treated as disturbances to the system. The right side of the graph includes a source power

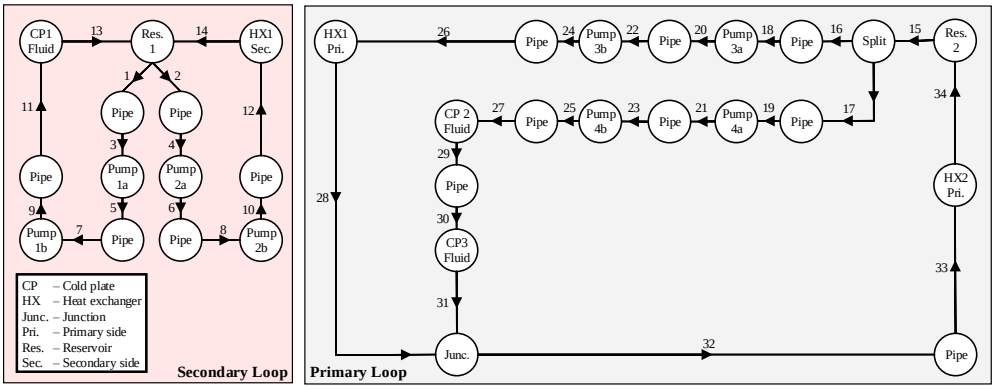

Figure 2.8: Hydraulic graph for example testbed configuration [64].

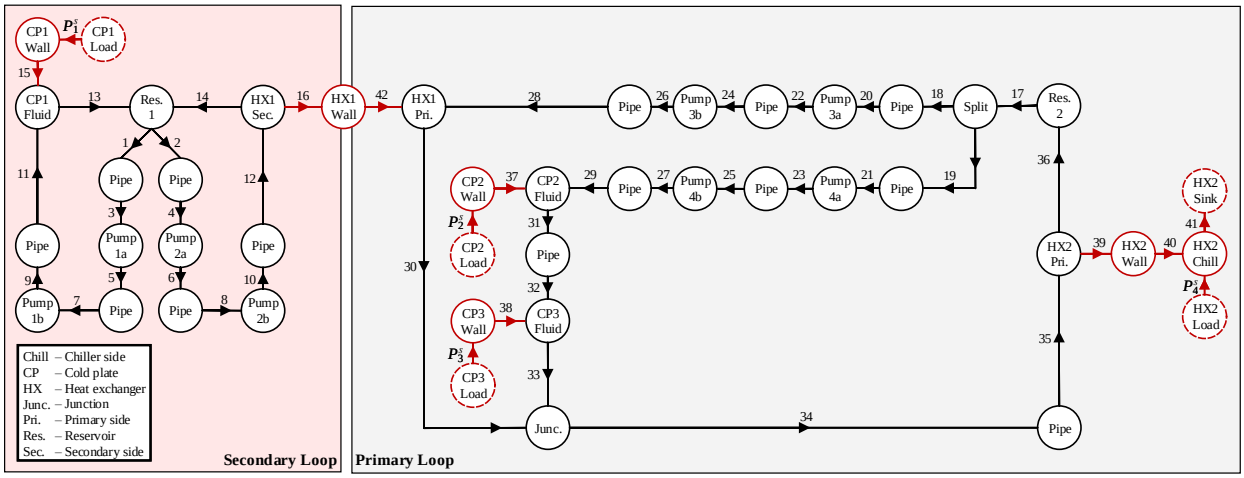

Figure 2.9: Thermal graph for example testbed configuration [64].

flow and sink state to capture the flow from and to the chiller, respectively. Thus in (2.18),

$$P^{s,r} = \begin{bmatrix} Q_1 & Q_2 & Q_3 & \dot{m}_{ext} c_p T_c \end{bmatrix}^T \tag{2.34}$$

where each Q is the heat load to the corresponding CP, $\dot{m}_{ext}$ is the mass flow rate of chilled fluid through the right side of HX2, and T_c is the temperature of the fluid exiting the chiller and entering the right side of HX2. The thermal graph consists of 39 vertices (one of which is a sink vertex), 41 edges, and 4 source edges. This results in a corresponding graph-based thermal model with 38 temperature states and 41 thermal power flows.

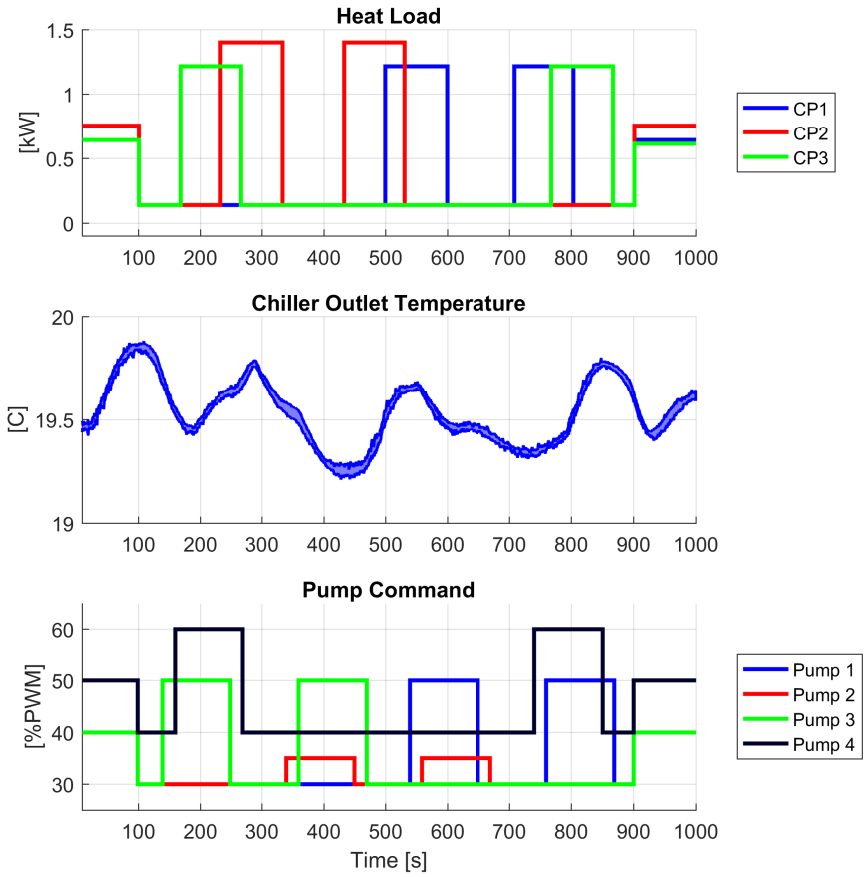

Figure 2.10: Inputs and disturbances used for model validation [64].

2.4.4 Validation of Fluid-Thermal Models

Figure 2.10 shows the commands and disturbances applied to the experimental system and models for validation. The linearization point used for the linear models is the steady-state operating condition of the nonlinear models subject to commands and disturbances that fall approximately in the middle of the operating range. To demonstrate the repeatability of the system across multiple runs, five experimental trials were conducted with the same commanded sequence. The traces for the chiller outlet temperature of Fig. 2.10 show the envelope between the maximum and minimum value measured at each time among the five trials.

The heat loads plotted in Fig. 2.10 are translated into a reference current for each CP's resistive heaters using an empirical map between the applied electrical current and the

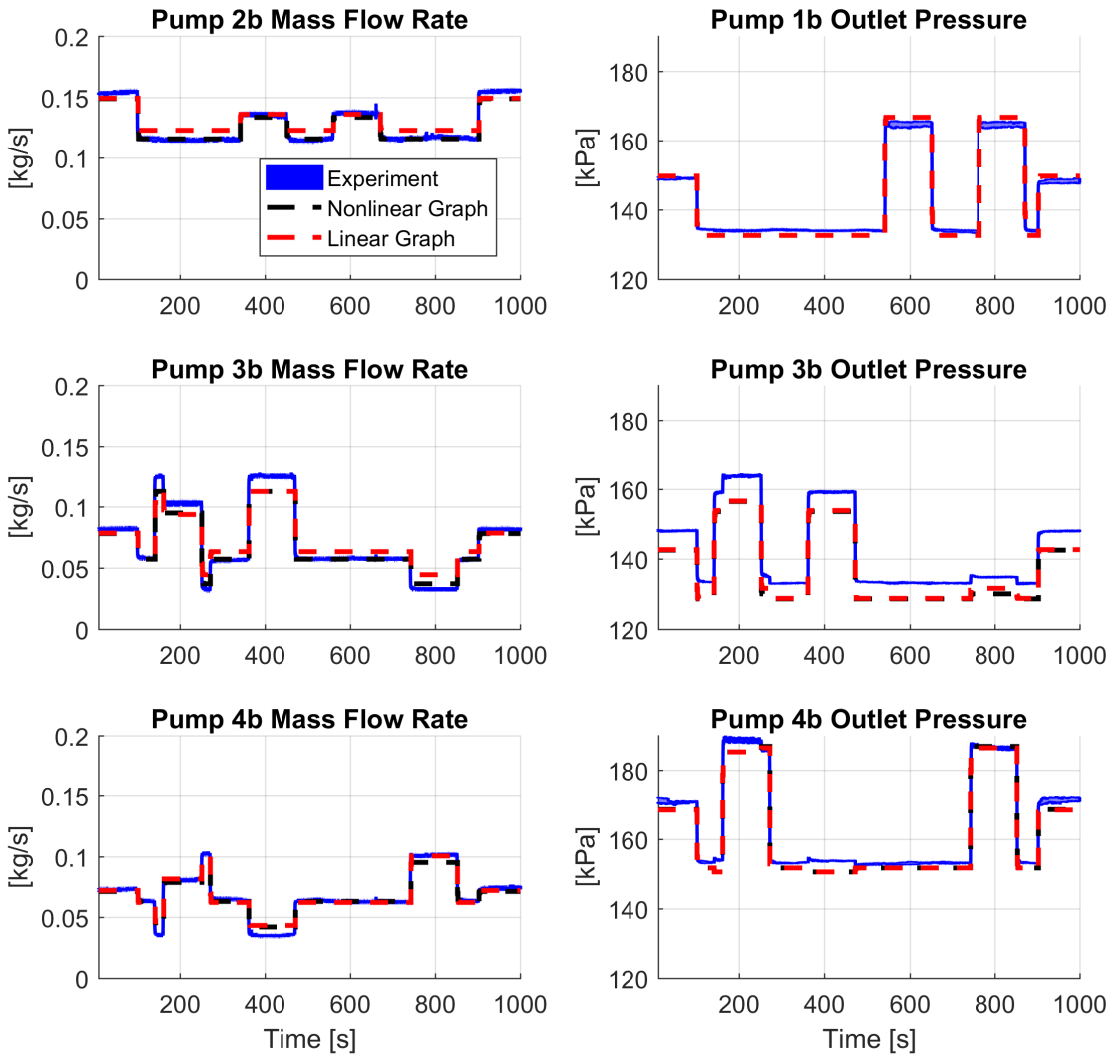

Figure 2.11: Selected signals for hydraulic validation of experimental data with nonlinear and linear graph-based models [64].

achieved heat load. Each reference current is tracked by proportional-integral (PI) control of the corresponding solid state relay.

The chiller is set to track a temperature set point of 20°C. Fig. 2.10 shows that deviation from this set point of about 0.5°C on average is present due to measurement and tracking error within the chiller's internal controller.

Figures 2.11-2.12 show a selection of hydraulic and thermal signals, respectively, that result from applying the inputs and disturbances of Fig. 2.10. All experimental traces plotted show the envelope between the maximum and minimum value measured at each time among five

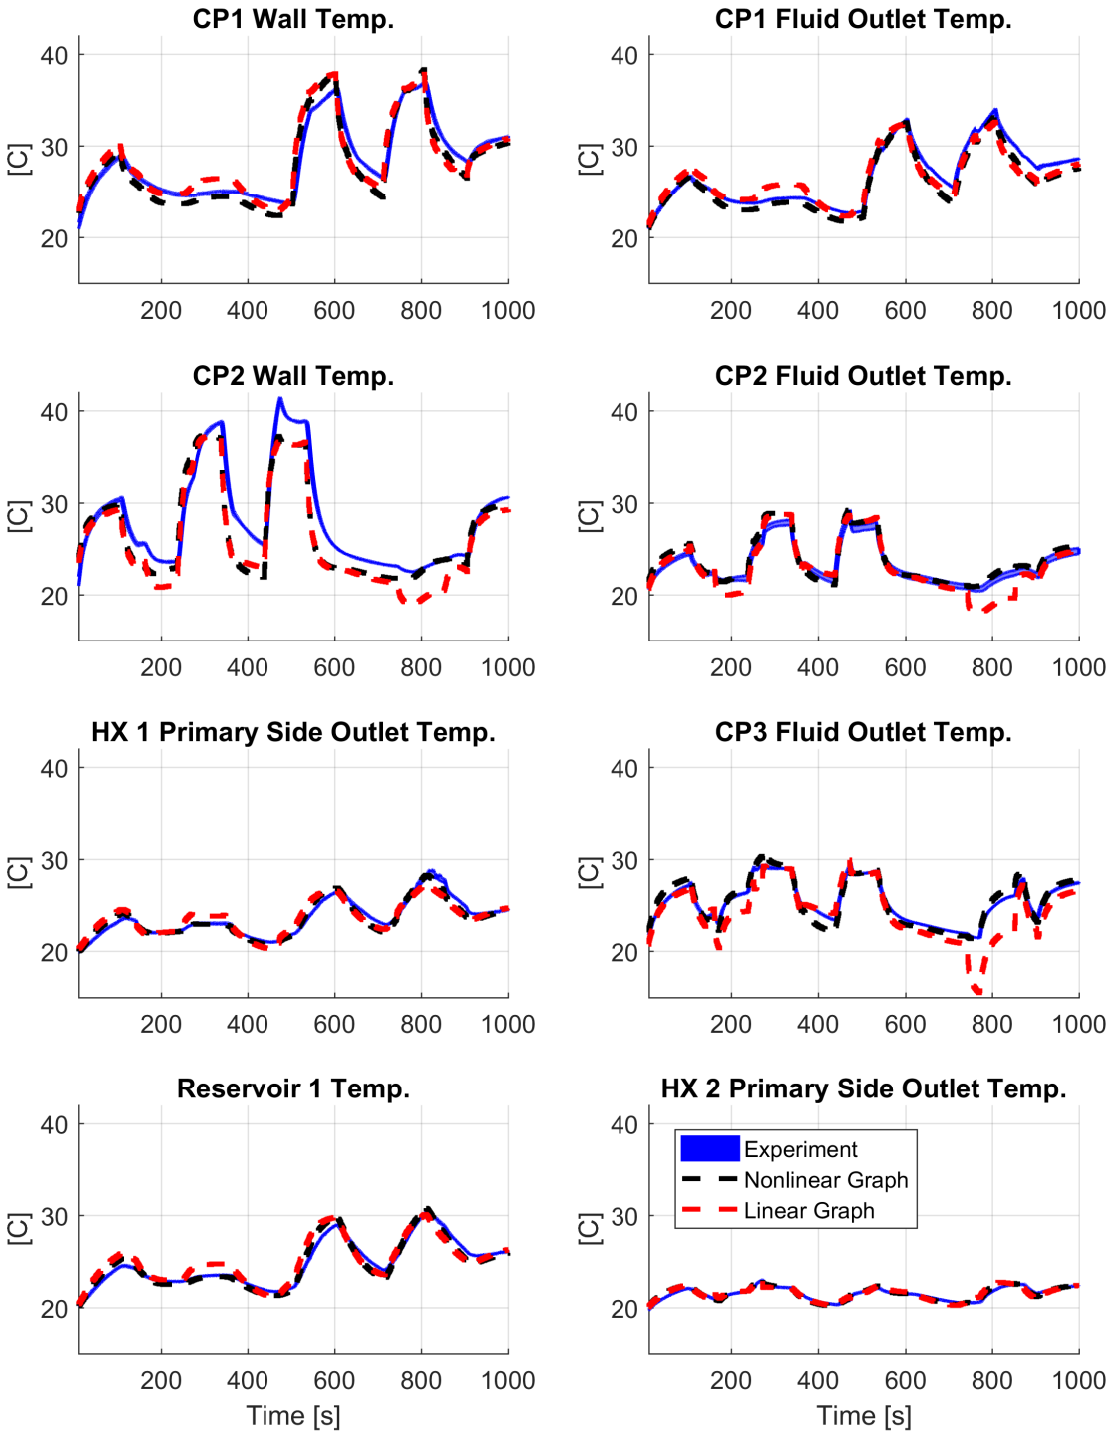

Figure 2.12: Selected signals for thermal validation of experimental data with nonlinear and linear graph-based models [64].

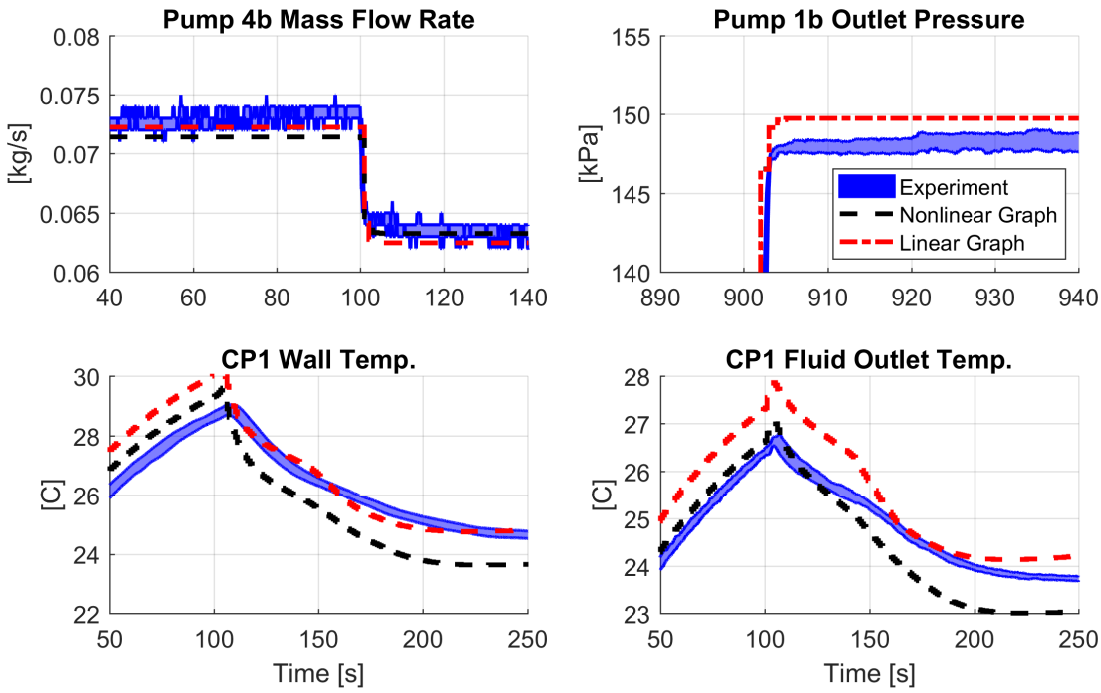

Figure 2.13: Closer view of several signals from Figs. 2.11-2.12. All experimental traces show the envelope between the maximum and minimum value measured at each time among five experimental trials [64].

experimental trials. To make the width of these envelopes more clear, a closer view of several signals is provided in Fig. 2.13, which demonstrates that the testbed exhibits a high degree of repeatability.

Figure 2.11 demonstrates a close matching between the experimental data and the hydraulic graph-based models. While offset occurs at times between the models and data, this is generally small relative to the magnitude of the gains when commands change. Where differences exist between the two models, especially in the traces for the pump 2b and 3b mass flow rates, the nonlinear model is more accurate than the linear model. This is due to the error incurred by linearization of the terms under the square root in (2.9) and (2.10).

Figure 2.12 similarly demonstrates a high degree of accuracy in the nonlinear thermal graph-based model. The discrepancies that occur can be attributed to unmodeled friction and thermal energy losses to ambient air, as well as errors in the heat transfer coefficient correlations. In the interval from 750-850 s, significant error occurs in signals of the linear model pertaining to CPs of the primary loop. This is largely due to the combination of a

low mass flow rate in pump 3 and a high mass flow rate in pump 4, which falls far from the linearization conditions. Linearization of the bilinear fluid thermal power flow equation (2.14) decouples the relationship between mass flow rates and temperatures, and this can result in large error under some operating conditions. However, the linear thermal model still preserves the correct signs in the gains during this time interval, as is critical to the design of stabilizing model-based controllers. The accuracy of the linear model at most other times across the 1000 s mission is close to that of the nonlinear model. While the hydraulic and thermal models could be made more accurate at the cost of increased complexity, the accuracy demonstrated here is sufficient for their intended use in model-based closed-loop control.

2.4.5 Hydraulic Coupling Constraints

As discussed in Chapter 3, successful control of the fluid-thermal system requires knowledge of the bounds on the achievable and/or allowable mass flow rates. From the perspective of the thermal graph-based models, this can be viewed as constraints on the edge inputs $\dot{m}^r$. For system architectures involving splits and/or junctions, the mass flow rate constraints associated with the fluid lines may be strongly coupled to each other. For example, in Fig. 2.7 the split in the primary loop results in coupling between the mass flow rates through pump 3 and pump 4. To determine these coupling constraints, the nonlinear hydraulic graph-based model is simulated to steady-state at all combinations of pump speeds in the range of 20-65% PWM in increments of 0.25%.

Fluid flow reversal (i.e., flow in the opposite direction from the arrows in Fig. 2.7) can be captured in the hydraulic graph-based modeling framework by including an absolute value under the square root in (2.9)-(2.10) and multiplying the full expression by a signum function of the pressure differential. When this is done, mass flow rates in the reverse direction from the edge orientation of the hydraulic graph are assigned a negative sign. Flow reversal is typically an undesirable behavior in fluid-based thermal management systems. Therefore, as a safety margin against flow reversal, any input combinations resulting in a mass flow rate less than 0.03 kg/s are excluded from the allowable hydraulic operating conditions. The

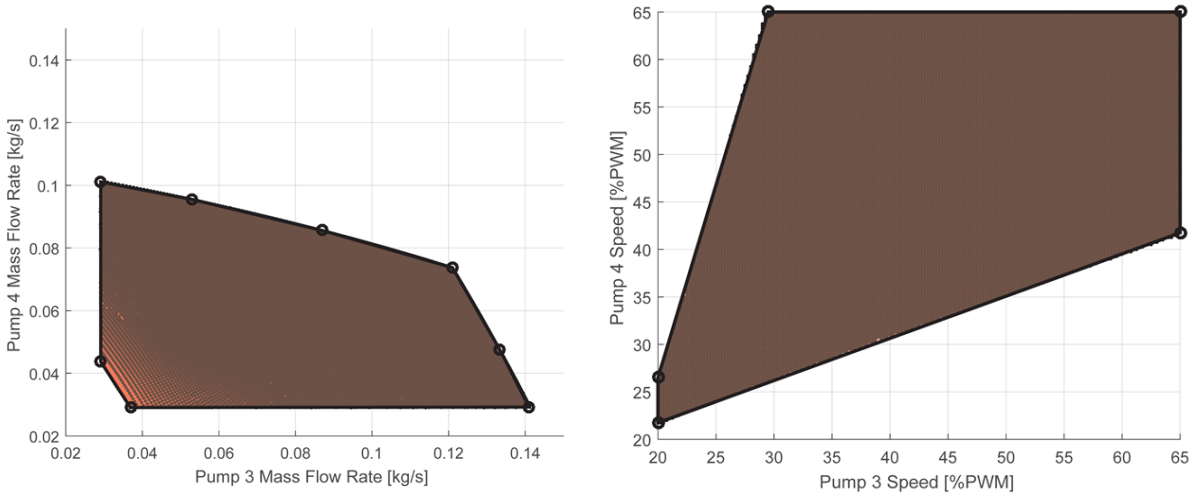

(a) Envelope of mass flow rates above 0.03 kg/s.

(b) Pump commands generating mass flow rates in Fig. 2.14(a).

Figure 2.14: Hydraulic coupling in Pumps 3 and 4 [64].

resulting envelope of allowable and achievable mass flow rates through pumps 3 and 4 is shown in Fig. 2.14(a). Figure 2.14(b) shows the corresponding envelope of pump commands generating mass flow rates in the envelope of Fig. 2.14(a). Combinations of commands involving a relatively high speed in one pump and a relatively low speed in the other pump are seen to be excluded from the envelope because they can result in flow reversal.

The envelopes in Figs. 2.14(a)-2.14(b) can be accurately represented as polytopes with a relatively small number of vertices. For example, eight vertices are sufficient to define the envelope in Fig. 2.14(a), as indicated by the circled points. The half-space representation, or "H-representation," defining this polytope is given by the linear inequality:

$$E_{pri} = \{\dot{m}_{pri} \mid H_{pri}\dot{m}_{pri} \leq z_{pri}\}, \tag{2.35}$$

where H_{pri} is a matrix and z_{pri} is a vector of appropriate dimensions, and

$$\dot{m}_{pri} = \begin{bmatrix} \dot{m}_{pump\ 3} & \dot{m}_{pump\ 4} \end{bmatrix}^T. \tag{2.36}$$

2.5 Utility for Energy Systems

The graph structure underlying the graph-based modeling approach facilitates the modular assembly of multi-domain system models from component and subsystem models. This structural information is found, for example, in the incidence matrix $\bar{M}$ that serves as a map from power flows P to states x and appears directly in the system dynamics of (2.7). In addition to supporting modularity in model formulation, the graph-theoretic nature of the approach has been shown to be valuable for model-based hierarchical control, model decomposition, and system design optimization. The first of these is an area of focus in this dissertation. As shown in Chapters 4-5, graph-theoretic notions such as paths and connectivity can be employed for analysis of the stability and robustness of hierarchical control frameworks. While model decomposition and design optimization are not areas of focus in this dissertation, a brief discussion of recent results for each is included below to illustrate the broad utility of the graph-based modeling approach.

2.5.1 Model Decomposition

A feature of graph-based models that makes them attractive for model-based hierarchical control is that a system model can easily be decomposed into sub-models by partitioning its graph into sub-graphs based on an analysis of the edges and vertices. Edges that are cut as a result of this partitioning represent coupling terms between the sub-models, for which controllers of the hierarchy can account by exchanging information. While the specific partitioning applied can be determined by engineering intuition, separation of physical domains, spatial configuration within the system, or other means, a key consideration must be to preserve the most significant state coupling terms within each sub-model.

A system decomposition that quantitatively takes this coupling into consideration can be achieved using tools from graph theory. In [66, 92], an impulse disturbance is injected into the input of a graph-based model. The resulting power flow along the edges is integrated in time to produce an energy-based measure of the strength of coupling between all adjacent vertices, which is in turn used as a distance metric within an agglomerative hierarchical

clustering algorithm. The resulting dendrogram informs on how to cluster vertices, and their associated dynamic states, into sub-models within each level of a hierarchical control framework such that the strength of coupling along cut edges is minimized. In [66], this method is demonstrated in application to electro-thermal hierarchical control of a simulated automotive electric vehicle powertrain.

2.5.2 Design Optimization

The graph-based modeling approach has also been demonstrated in application to system design optimization. From a library of component graphs and knowledge of their possible interconnections, graph-theoretic techniques can be leveraged to generate all possible graphs within a class of system architectures. The corresponding dynamic models can then be programmatically generated, optimized, and compared to reveal the highest performing architectures.

In [33, 93], this process is performed for a class of single-phase fluid-thermal systems consisting of cold plates (CPs) in series and/or parallel flow with a single junction and split, as shown in Fig. 2.15. For a given number of CPs, all candidate graph-based models within the class of architectures can be programatically generated using enumeration methods for rooted tree graphs [94, 95]. An optimization program is solved for each candidate to determine the time-varying valve commands that maximize its thermal endurance, defined as the duration until any temperature constraint in the system is violated. All candidates are ranked by their thermal endurance, revealing the best designs under this figure of merit. In [33, 93], this process is performed for architectures with up to six CPs, for which there are 4051 unique candidates within the considered class.

2.6 Review of Features

Section 2.1.1 establishes five desired features that control-oriented modeling approaches should include to facilitate model-based control design for complex energy systems. These

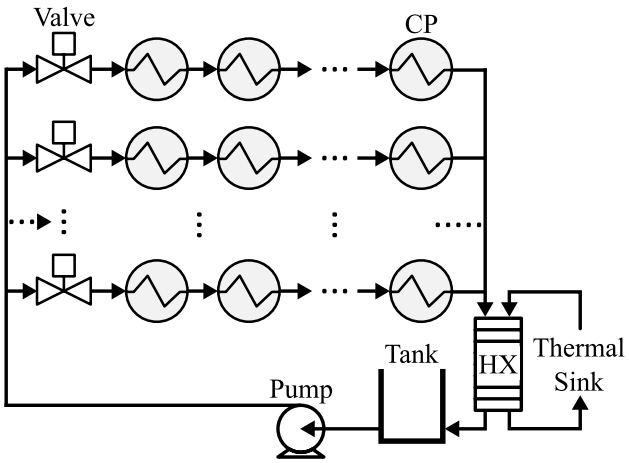

Figure 2.15: Class of fluid-thermal architectures considered for design study [33].

features are next revisited, using examples from this chapter to explain how they are achieved by graph-based modeling and how they support the objectives of this dissertation.

2.6.1 Modular

As shown in Sections 2.3-2.4, graph-based models of systems can be formed by interconnecting graphs of their constituent components and subsystems. Unlike simulation-based toolboxes as discussed in Section 1.1.3.1, this modularity is achieved while preserving access to the dynamic equations governing the system behavior. Furthermore, unlike bond graph models as discussed in Section 2.1.2, the assembly of component or subsystem models does not result in excessively complex structures necessitating further simplification.

Modularity is useful not only in constructing system models, but also in decomposing models for hierarchical model-based control. Any desired decomposition can be performed, regardless of which components or subsystems the vertices were associated with when the model was composed. A quantitative decomposition approach is discussed in Section 2.5.1, where the graph structure of the models facilitates the application of a clustering algorithm together with an energy-based distance metric that measures the strength of coupling between adjacent vertices of the graph.

2.6.2 Physical Domain and Timescale Agnostic

In the generic graph-based modeling approach of Section 2.2, vertices are associated with the storage of a conserved quantity, regardless of its physical domain. Edges represent the paths for exchange of that quantity between vertices. Section 2.3 demonstrates how the generic approach can be applied to capture the hydraulic and thermal domains of fluid-thermal components. The system used to demonstrate hierarchical control in Chapter 3 includes graph-based models of the electrical domain. Further graph-based models and experimental validation for electrical, thermal, and mechanical components of aircraft, hybrid unmanned aerial vehicles, and electric automotive vehicles are found in [66, 77, 88, 89, 92, 96]. In each case, the modeling approach is shown to capture dynamic behavior with sufficient accuracy for the intended use in model-based closed-loop control. Coupling between domains can be represented using the multi-graph approach in Section 2.3.3 or within a single graph, as in Chapter 3 and [66, 77, 89, 92]. This illustrates how the modeling approach serves as a unified framework by which to capture dynamic interactions within and between physical domains at their relevant timescales.

2.6.3 Structure-Preserving

In the graph-based system dynamics of (2.7), the incidence matrix $\bar{M}$ serves as a map of the underlying structure of coupling among the states. Chapters 4-5 illustrate how this knowledge of the coupling structure can support analysis of stability and robustness. This is also leveraged in the model decomposition approach discussed in Section 2.5.1, which provides a metric by which to evaluate the strength of each coupling term and informs on how to decompose a graph-based system into sub-models within a hierarchical control framework. As discussed in Section 2.5.2, the dynamic model of a candidate system architecture can be programatically generated from the structural information of its graph representation. This allows architecture exploration to be performed without requiring the model for each candidate architecture to be manually generated, which would be prohibitively time consuming in cases where there are hundreds or thousands of unique candidates to evaluate.

2.6.4 Hybrid

As discussed in Section 1.2.1, many energy systems are governed by both continuous and switched behavior. While this chapter focuses on continuous dynamics, later chapters include extensions to capture switched behavior. Two methods for this are employed, both of which involve modifying edge relationships.

The first method is to restrict the edge inputs u of (2.8) to fall within one of a collection of constraint sets $\mathbb{U}_i$, $i \in [1{:}N_m]$. The active constraint set at each time is selected by a switching signal $\sigma(t)\colon [0,\infty) \to [1{:}N_m]$, with $u(t) \in \mathbb{U}_{\sigma(t)}$. When each $\mathbb{U}_i$ is singleton, this can be equivalently represented using a unique switching signal for each edge $\sigma_j(t)\colon [0, inf) \to [1{:}N_{m,j}]$, selecting from among a finite set of values $\mathbb{V}_j = [u_{j,1}{:}u_{j,N_{m,j}}]$ such that $u_j(t) = u_{j,\sigma_j(t)}$. Another equivalent representation when each $\mathbb{U}_i$ is singleton is to eliminate the edge inputs entirely and instead consider switching of the functions governing the edge power flows directly, such that (2.2) becomes $P_{j,\sigma_j} = f_{j,\sigma_j}\left(x_j^{tail}, x_j^{head}\right)$.

The second method allows edges internal and external (i.e., source edges) to the graph to be switched independently from the action of the edge inputs. This is done by first defining indicator matrices $\Phi_\sigma = diag\left([\phi_j]\right)_\sigma$, $\phi_j \in \{0,1\}$, $j \in [1{:}N_e]$ and $\Gamma_\sigma = diag\left([\gamma_j]\right)_\sigma$, $\gamma_j \in \{0,1\}$, $j \in [1{:}N_s]$. These indicator matrices are then incorporated into (2.7) to obtain the switched system:

$$\mathbf{S}_\sigma\colon \quad C\dot{x} = -\bar{M}\Phi_\sigma P + D\Gamma_\sigma P^s. \tag{2.37}$$

More details on switched graph-based models and illustrative examples are provided in later chapters where applicable.

2.6.5 Flexible in Representation

While the dynamics of (2.7) are linear with respect to the power flows, the functions governing the power flows in (2.2) are nonlinear in general. However, as shown in Sections 2.3-2.4, linearization of these functions can be performed to generate linear graph-based models, at the cost of increased model error at some operating conditions. In later chapters, assump-

tions imposed on the form of the power flow functions are shown to facilitate analysis of stability and robustness.

In addition to the flexibility of the edge equations, the graph-based approach can have flexibility in the number of vertices used to capture the dynamics of specific components, systems, or subsystems. For example, while the wall of the thermal cold plate model in Fig. 2.2 is represented by a single lumped temperature, it could instead be represented using multiple coupled states to capture the spatial variation in temperature across the wall. In this way, more detailed and potentially more accurate models can be derived, at the cost of increased computational complexity due to the increased number of power flows and state equations. Conversely, the model complexity can be reduced by removing or combining vertices or edges of the graph, at the cost of loss of detail and potentially decreased accuracy.

2.7 Chapter Summary

This chapter motivates, presents, and experimentally validates a dynamic graph-based modeling approach for power flow systems. Derived from first principles, this approach can be applied to capture interactions within and between multiple physical domains and dynamic timescales. Modularity allows component and subsystem models to be generated independently and then interconnected to form system models. The governing equations of the model make explicit its underlying structure of state coupling. This can be leveraged in system analysis, model decomposition, design optimization, and hierarchical control. The approach can be extended to capture switched behavior and is highly flexible in representation, for example by admitting nonlinear or linearized governing equations. To demonstrate the graph-based modeling approach, hydraulic and thermal models of fluid-thermal components are presented and validated using an experimental testbed.

Chapter 3

Hierarchical Control of Aircraft Electro-Thermal Systems

3.1 Motivation and Background

As discussed in Chapter 1, hierarchical control frameworks represent a promising technology for addressing the exigent challenges of vehicle energy management. These frameworks supplant traditional decentralized control approaches with a paradigm that coordinates both continuous and switched dynamics across subsystems, components, timescales, and physical domains. To highlight the value of this coordination, this chapter serves as a case study in the construction and experimental application of a hierarchical MPC framework. This is performed for a hardware-in-the-loop (HIL) experimental testbed representative of a scaled aircraft electro-thermal system.

Many advanced controllers for vehicle energy systems have been proposed in the literature and demonstrated in simulation, including [1, 19, 23, 24, 44, 50–53, 61, 65, 66, 97–100]. These often take a centralized approach [19, 24, 50–53, 98], while decentralized approaches are typically limited to a two-level inner/outer loop structure with a single control formulation in each level [23, 44], limiting extensibility to more general distributed frameworks that can accommodate additional timescales and dynamic complexity. Furthermore, relatively few of these strategies have been implemented experimentally, where model error, communication delays, computational limitations, and other phenomena can result in significantly reduced performance from that achieved in simulation. The experimental demonstration that has been performed in the literature is typically limited to a single physical domain, such as the control approaches for thermal management in [48, 64, 71]. Therefore, a key contribution of this chapter is to demonstrate that the high performance of hierarchical MPC for multi-

domain coordination of energy systems, which has been previously been exhibited only in simulation, is also realizable in practice.

The goals of any controller for aircraft energy management are to 1) maximize the capability of the aircraft by achieving the desired operation of mission- and flight-critical hardware, 2) satisfy various system constraints for safe and reliable operation, and 3) minimize fuel consumption. In general, the sheer complexity of aircraft energy systems, high degree of coupling among their constituent elements, and wide range in dynamic timescales can present barriers to the simultaneous consideration of these goals during real time operation.

In this chapter, load shedding is adopted as a metric for comparing the electrical capability achieved under different control strategies, quantifying the first goal above. This mirrors the optimal electric power distribution and load management strategies of [23, 24], in which non flight-critical loading is reduced to prevent overloading of the electrical system. In this chapter, controllers that shed less of a desired load profile are said to provide greater capability to achieve a desired mission or flight plan. Performance in this figure of merit is considered together with additional figures of merit quantifying the second and third goals identified above.

Figure 3.1 shows a candidate hierarchical framework for aircraft energy management. To achieve coordination across fast and slow dynamics, relatively long timescale behavior is managed with long update intervals (and therefore long time horizons when MPC is used) by a controller at the top level of the hierarchy, while shorter timescales are managed by faster-updating controllers at lower levels. Within each level, decision-making can be distributed among multiple controllers governing partitions of the full system dynamics. As enclosed by the green triangle in Fig. 3.1, the case study in this chapter focuses on the vehicle-, system-, and subsystem-level coordination of electrical and thermal systems. Therefore, faster timescale component-level aspects such as voltage regulation of power electronic devices, cell balancing within batteries, and motor control of pumps lie beyond the present scope. However, the methodology presented here will permit these aspects to be incorporated in future work.

In an electro-thermal hierarchical control framework, coordination between controllers can account for coupling between the electrical and thermal domains to ensure that electrical

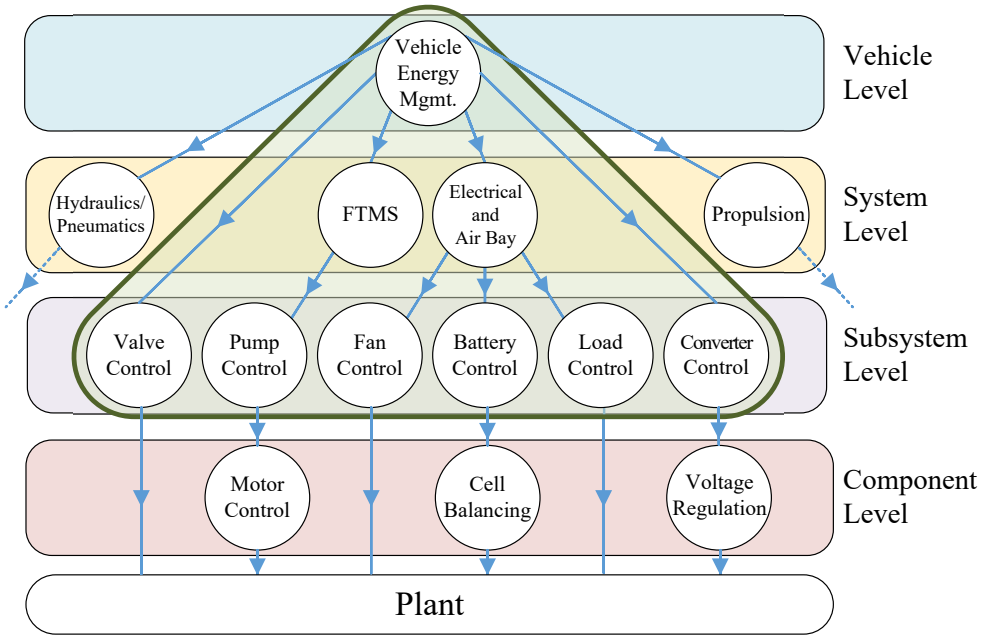

Figure 3.1: Candidate hierarchical framework for aircraft energy management. The green triangle designates the focus of this case study [101].

operation falls within the ability of thermal management to provide adequate cooling. At the top of the hierarchy, objectives and constraints across the entire vehicle including all relevant physical domains can be considered together. Using preview of the mission or flight profile and knowledge of anticipated environmental conditions, proactive action can be taken to prepare the thermal system in advance of large loads (known as pre-cooling) and strategically throttle or reschedule the operation of non-critical electrical systems if necessary to remain within electrical and thermal constraints. These coordinated decisions from high-level controllers are communicated down the hierarchy as references for lower-level controllers managing specific partitions of the system at faster timescales.

This hierarchical decomposition of the complete energy management task into multiple communicating controllers provides significant computational benefits over the use of a single, centralized control formulation. This allows both long-term planning and fast disturbance rejection to be performed in real time despite limited availability of computational resources. This is particularly important for aircraft, where on-board computational resources may be constrained by size, weight, cost, power consumption, or limitations of legacy hardware.

The proposed hierarchical MPC framework in this chapter is demonstrated on a scaled HIL testbed representative of an aircraft electro-thermal system. As shown on the right in Fig. 3.2, the testbed consists of a simulated air and electrical system coupled to the fluid-thermal testbed described in Section 2.4.1. The electrical system, simulated in real time, consists of a generator, converter bank, battery, and three loads. The waste thermal energy generated by the converters and batteries is transferred to the electronics bay. Heat loads from other electrical components and the bay are physically imposed on the experimental fuel thermal management system (FTMS) consisting of heat exchangers, tanks, pumps, valves, and a thermal sink. The measured temperatures of these physical components are communicated to the simulated electrical system, where they affect the efficiency of electrical components.

The graph-based modeling approach of Chapter 2 captures coupling within and between the thermal and electrical domains using a common model formulation. The model of the complete system can readily be constructed by assembling component-wise graphs. The graph of the complete system can then be decomposed into sub-graphs to produce models for controllers throughout a hierarchical framework. Given a predefined mission profile and state feedback from both systems, a desktop computer executes the hierarchical MPC framework in real time, issuing actuator commands to both the simulated and experimental plants.

3.1.1 Advantages of Hierarchical MPC

Hierarchical MPC frameworks provide significant advantages over centralized MPC in application to multi-timescale systems such as aircraft energy systems. A centralized MPC controller has a fixed update interval ΔT, and number of steps in the prediction horizon N. To execute the controller in real time, the choice of ΔT and N must be coordinated such that the time required to solve the optimization program given the available computational resources is less than ΔT. In general, a sufficiently long time horizon, $\Delta T \cdot N$, enables anticipatory action by the controller, allowing it to leverage preview of upcoming disturbances and operational requirements to improve performance. Additionally, decreasing ΔT improves the

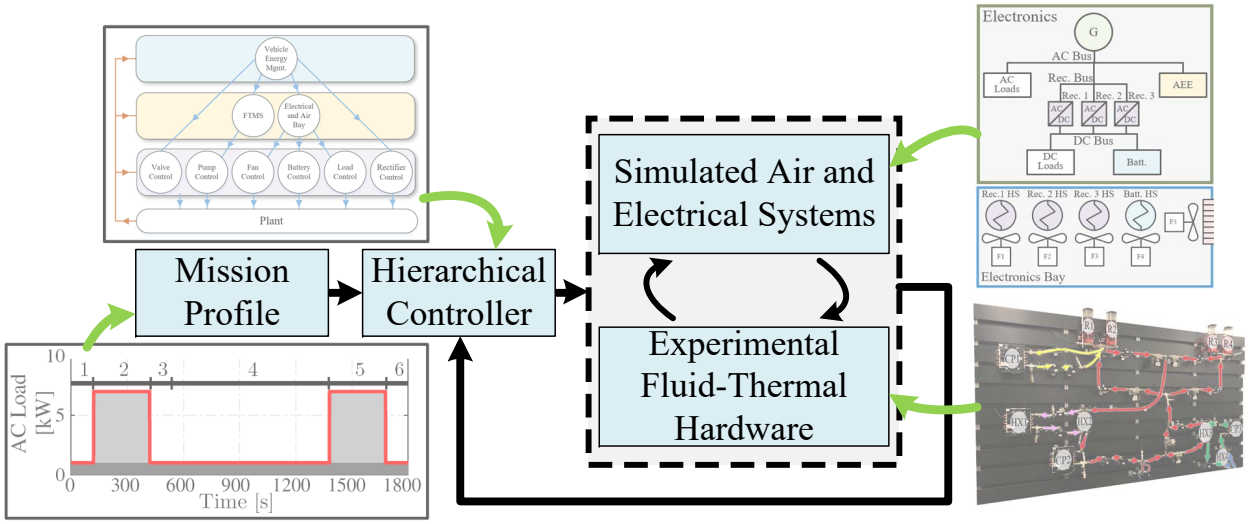

Figure 3.2: Schematic of the key elements of the HIL testbed used to represent an aircraft electro-thermal system [101].

bandwidth of the closed-loop system, improving its ability to perform disturbance rejection and compensate for model error. However, achieving a long time horizon with a small update interval requires a large N, which is constrained by the available computational power. As a result, the design of a centralized control approach must balance between the benefits of governing long timescale behavior and governing short timescale behavior in the selection of a single value for ΔT. By comparison, decomposition of the control problem into multiple levels of a hierarchical MPC framework provides more degrees of freedom in the control design. High level controllers can achieve long time horizons by using a large ΔT, while low level controllers can achieve high bandwidth by using a small ΔT. With communication between controllers enabling their coordination across timescales, the prediction horizon N for each controller can be relatively small, enabling high performance and real time implementation under computational limitations for which a centralized MPC controller would be ineffective.

The remainder of this chapter is organized as follows. Section 3.2 discusses the candidate electro-thermal system and HIL configuration used in this case study. Section 3.3 details the graph-based modeling of thermal and electrical components of the candidate system. Section 3.4 presents the hierarchical control formulation, while Section 3.5 presents a baseline

Table 3.1: Aircraft components associated with fundamental power and energy mechanisms in the electrical and thermal domains.

Mechanism	Electrical domain	Thermal domain
Generation	Generators	Electronics, passengers, hydraulics, engines, aerodynamic heating
Transport	Buses, switches, wires	Pipes, heat exchangers
Conversion	Converters, transformers	Air cycle machines, vapor compression cycles
Storage	Batteries, capacitors	Fuel tanks, phase change materials, cabins, bays
Utilization/ dissipation	Avionics, motors, anti-icing, environmental control systems	Air (ambient, engine bypass, ram, third stream)

control approach used for comparison. The controllers are implemented experimentally and compared in Section 3.6. Lastly, Section 3.7 provides a concluding summary of the chapter.

3.2 Candidate Electro-Thermal System

The architecture, components, and sizing of aircraft electro-thermal systems can vary significantly across vehicle classes and platforms. However, despite this variation, most modern systems reflect the same fundamental power and energy mechanisms, and include the same general component types, as summarized in Table 3.1. The candidate electro-thermal system used for demonstration this chapter, while not intended to emulate any specific platform, includes these key mechanisms and captures coupling between components in the electrical and thermal domains. In accordance with the discussion of Section 1.1.4, a simplified and reduced-scale experimental testbed achieves these features at an economical and laboratory-safe power level.

Figure 3.3 diagrams the system architecture used in this chapter. The electrical system, boxed in green, consists of an AC generator, an AC/DC converter bank with three parallel converters, a battery, and three types of loads. The AC loads and Advanced Electrical

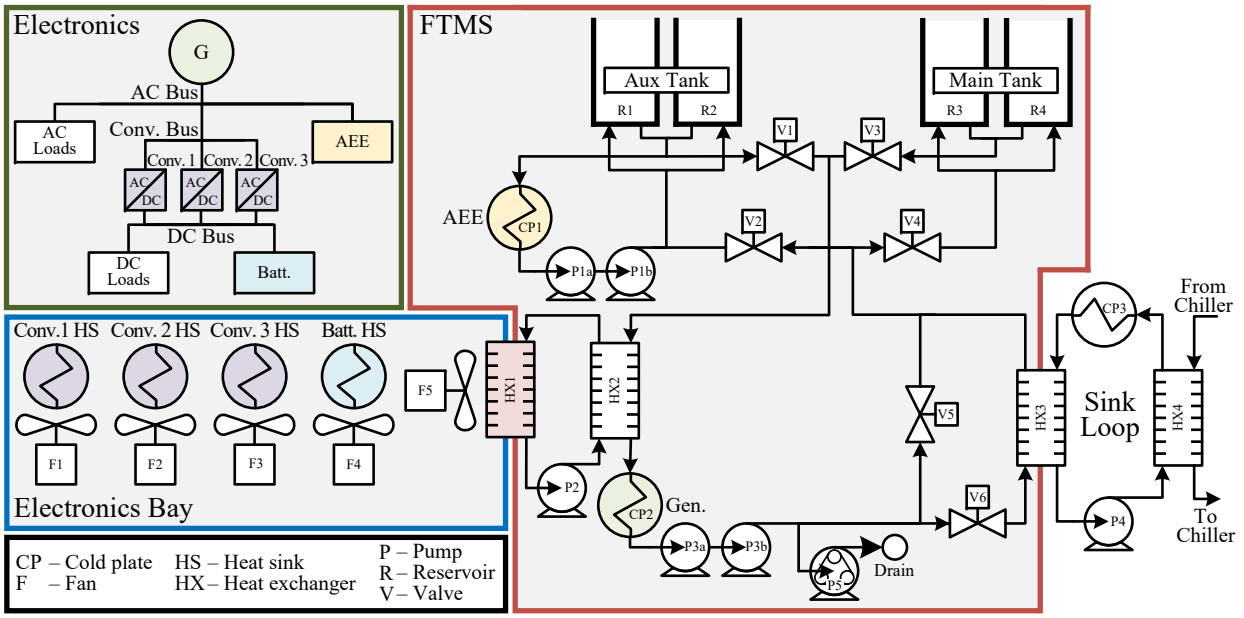

Figure 3.3: Schematic of the electro-thermal system architecture used to demonstrate the proposed hierarchical control framework. Components are color-coded to highlight the coupling created by thermal energy generation in the electrical system [101].

Equipment (AEE) are powered directly from the AC bus. The DC loads are powered from the DC bus, which is supplied by the converter bank and the battery.

While thermal energy generation by the AC and DC loads is assumed to be negligible, the generator, AEE, converters, and battery generate thermal energy in accordance with their electrical efficiencies. In this case study, the generator is assumed to have a constant electrical efficiency, while the AEE and battery electrical efficiencies are functions of their operating temperature, as detailed in Section 3.3.2. The electrical efficiency of each converter is assumed to be a function of its current. Each converter may be switched on or off, therefore the overall efficiency of the converter bank also depends on the number of active converters, as shown in Fig. 3.4. Note that these efficiency curves are notional and, while similar to those of commercial devices, are not intended to represent any specific commercial device or manufacturer specification. For simplicity in this example, the time-varying efficiencies are assumed to be a function of either temperature or current. In practice, efficiency can be a function of both of these variables, as well as other factors. The general methodology of this chapter does not preclude a more complex treatment of component efficiencies.

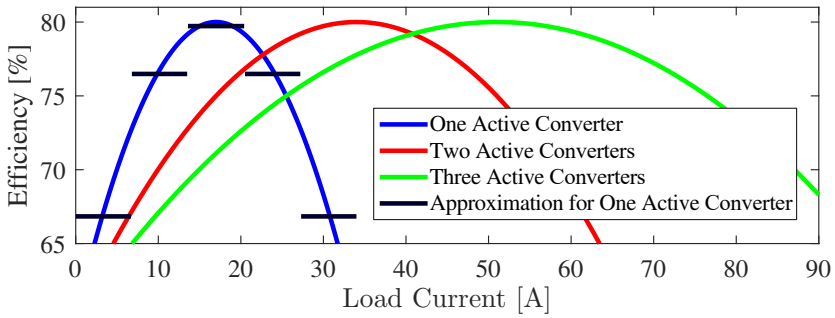

Figure 3.4: The converter bank efficiency is a function of the load current and the number of active converters. In the hierarchical control formulation, these nonlinear efficiencies are approximated by a piecewise constant function, plotted in black [101].

The three converters and battery are housed in an electronics bay, boxed in blue in Fig. 3.3. Each of these components has a heat sink (HS) and a variable-speed fan to transfer thermal energy from the component to air in the bay. The bay air is cooled by the FTMS using an additional variable speed fan and an air-to-liquid heat exchanger (HX), labeled as HX1 in Fig. 3.3.

Thermal energy from the generator, AEE, and electronics bay is transferred to the FTMS, outlined in red in Fig. 3.3, consisting of main and auxiliary fluid tanks, cold plate (CP) and liquid-to-liquid heat exchangers, pipes, splits, junctions, variable-speed pumps, and on/off solenoid valves. The AEE cold plate is cooled by fluid from the auxiliary tank. By controlling valves V1 and V3, the inlet fluid to the right side of HX2 can be sourced from either or both of the tanks. Similarly, valves V2 and V4 can be used to return fluid to either or both tanks. After passing through HX2, the fluid passes through the generator cold plate. The drain pump then allows fluid to be removed from the system, analogous to the burning of fuel in an aircraft engine. This fluid drain also results in a continual decrease of the total thermal energy storage capacity of the system.

Thermal energy is removed from the FTMS by a liquid-to-liquid heat exchanger, labeled as HX3 in Fig. 3.3, which is cooled by a vapor compression-based chiller. The Sink Loop emulates the thermal sinks available to aircraft. With the chiller set to a fixed temperature set point, pump P4 and a resistive heater attached to CP3 are used to control the thermal sink fluid mass flow rate and temperature, respectively, which are treated as exogenous disturbances to the FTMS. By closing valve V6 and opening valve V5, the sink loop can be

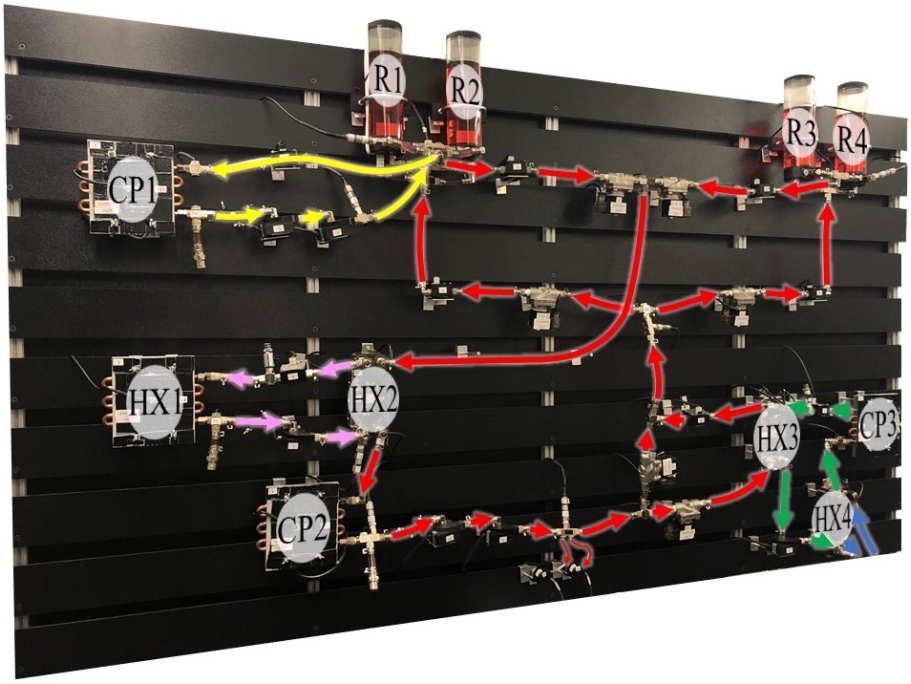

Figure 3.5: Experimental testbed configured to represent the FTMS [101].

bypassed. In practice, this should be commanded by a controller if the temperature of the thermal sink exceeds that of the fluid in the FTMS.

The experimental testbed from Section 2.4.1 has been configured to match the topology of the candidate FTMS, as shown in Fig. 3.5. The electrical system and electronics bay are modeled in MATLAB/Simulink using the Aerospace Power Systems (APS) toolbox detailed in [1], and deployed for real time HIL simulation to a dSPACE DS1005 processor board. The physical FTMS and the simulated electrical system and bay are bidirectionally coupled in real time. The measured temperatures of the CP1, CP2, and HX1 walls of the physical FTMS are communicated to the simulated electrical system, where they affect the efficiency of electrical components. The simulated electrical system in turn calculates a heat load to be applied to each of the three walls based on these efficiencies and other states of the simulated electrical system and bay. These heat loads are physically imposed by actuating resistive heaters attached to the walls of CP1, CP2, and HX1 in Fig. 3.5. Note that while HX1 is modeled as an air-to-liquid heat exchanger between the bay and FTMS, the liquid-side and heat exchanger wall are represented in the physical system by a cold plate heat exchanger, while the air-side of the heat exchanger is represented in the simulated bay model.

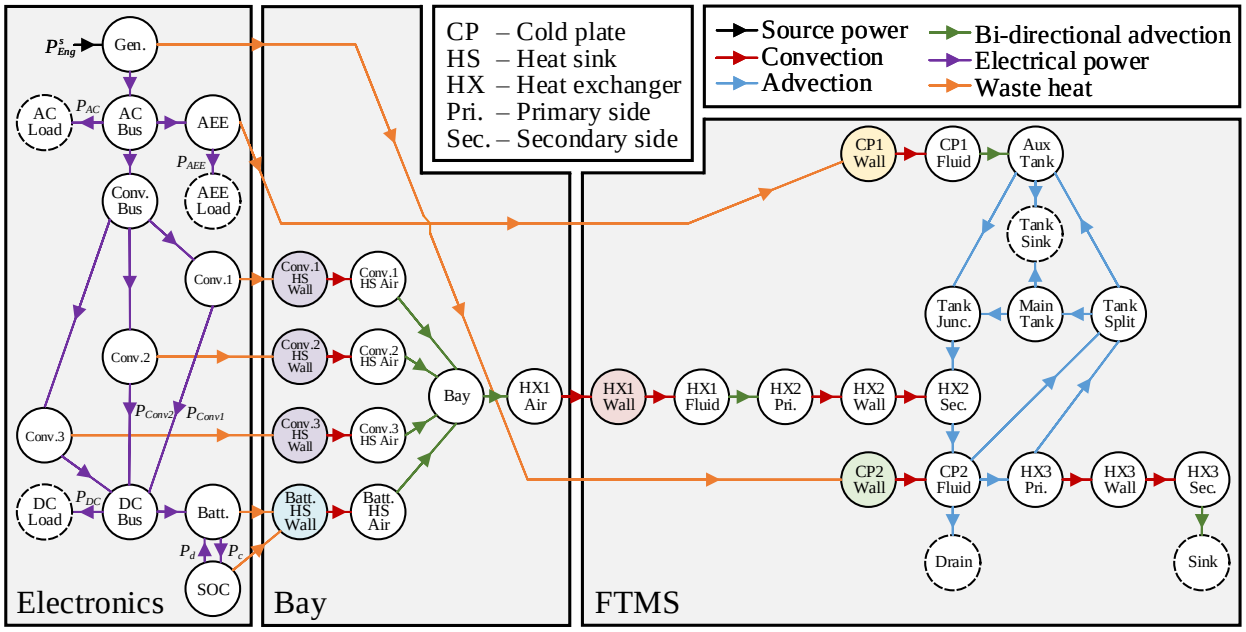

Figure 3.6: Graph-based model of the electro-thermal system in Fig. 3.3. Modified from [101].

3.3 Graph-Based System Modeling

This section applies the graph-based modeling approach of Chapter 2 to capture the dynamics of energy storage and power flow throughout the multi-domain system. As detailed in Chapter 2, vertices of the thermal domain represent thermal capacitances with dynamic states of temperature, while edges represent thermal power flow. In this chapter, graph-based modeling of the electrical domain is introduced. In this case, vertices either represent energy storage, as in a battery or capacitor, or an algebraic junction satisfying conservation of electrical power. To distinguish between thermal and electrical power flows, thermal power flows are denoted by Q, while electrical power flows are denoted by P. Figure 3.6 shows the electro-thermal graph corresponding to the system architecture from Fig. 3.3.

3.3.1 FTMS and Bay Model

The FTMS and bay are modeled using the thermal graph-based component models of Section 2.3.2. The graph of the FTMS is shown on the right in Fig. 3.6. Thermal energy is transferred to the FTMS from the AEE (via CP1), the generator (via CP2) and the bay (via

HX1). The vertex labeled as the "Tank Sink" captures the energy gained or lost due to a change in the fluid mass stored in either tank, as discussed in Section 2.3.2.

The switching of valves can be captured by the first method described in Section 2.6.4. A switching signal σ_v selects from among constraint sets for the edge inputs $\dot{m}$, associated with different combinations of valve positions (i.e., valve modes). Within each set, edges j in line with a closed valve are constrained such that $\dot{m}_j = 0$. The remaining mass flow rates are constrained to fall within a mode-dependent envelope of achievable values, found using the method described in Section 2.4.5. In the notation of that section, this can be expressed as:

$$\sigma_v = i \;\rightarrow\; \dot{m} \in E_i, \tag{3.1}$$

Where E_i is the envelope corresponding to the i^{th} valve mode.

The FTMS model is parameterized in accordance with the specifications of the testbed components detailed in Section 2.4.1. Open-loop validation with the testbed configuration in this chapter achieved similar accuracy as the graph-based model validated with a different configuration in Section 2.4.4, with errors typically less than 2°C. This validation was performed assuming constant heat transfer coefficients, with $\beta_{2,j} = 0$ in (2.17). As a result, $c_{5,j} = 0$ in (2.29). The remaining terms in (2.29) are either linear functions of temperature or bilinear functions between temperatures and mass flow rates serving as edge inputs $\dot{m}$. Therefore, the vector of thermal edge power flows, excluding the waste heat edges treated in the next section, can be written as:

$$Q_T = F_T\left(\dot{m}\right) \begin{bmatrix} T \\ T^t \end{bmatrix}, \tag{3.2}$$

where F_T is linear in $\dot{m}$.

The graph of the bay consists of vertices for the wall of each heat sink and the HX1 wall, vertices for the air flow across these walls, and a vertex for the thermal capacitance of the air in the bay. Using appropriate material properties for air, the vertex for the air in the bay is modeled identically to a reservoir in Section 2.3.2, while the vertices for air flow are modeled identically to pipes. The bay is parameterized to match the scale of the

fluid-thermal testbed, with an air volume of approximately 9.5 m^3, with heat sinks of mass 1 kg.

3.3.2 Electrical System Model

As indicated by the green triangle in Fig. 3.1, the fast dynamics associated with some control tasks in the electrical domain, such as voltage regulation and motor control, fall outside the scope of the case study in this chapter. Instead, the focus here is on high-level decision-making to coordinate electrical and thermal dynamics. Therefore, the only dynamic vertex in the graph-based electrical system model is that of the battery, with C_{SOC} as the energy capacity of the battery and $x_{SOC}(t)$ as the battery state-of-charge (SOC). These are associated with the vertex labeled as "SOC" in Fig. 3.6. All other vertices of the electrical domain are modeled as algebraic, with $C_i \dot{x}_i = 0$.

Validation of similar graph-based electrical system models has been presented in [1, 77], where simulations of graph-based models are compared to simulations using the higher fidelity PowerFlow toolbox [39]. The graph-based approach is shown to accurately match the PowerFlow models for a system consisting of a generator, AC and DC buses, and a set of constant impedance, power, and current loads. Specifically focusing on electro-thermal systems, the graph-based modeling of a multi-level DC/AC inverter is experimentally validated in [88]. The graph-based model matches transient temperature measurements to within 3°C, subject to conductive heat transfer through a multi-layered printed circuit board with thermal energy generated by inefficiencies in 12 gallium nitride (GaN) transistors. While a more thorough treatment of the fast electrical dynamics governed in the bottom level of Fig. 3.1 falls beyond the scope of this case study, it should be noted that these can be captured using the graph-based modeling approach, as detailed in [66, 89, 92].

Edges in the electrical domain (the purple edges in Fig. 3.6) represent the flow of electrical power between vertices. The electrical domain graph consists of 15 edges representing electrical power and one source power flow representing the power extracted from the engine by the generator, P^s_{Eng}. While source power flows are typically treated as exogenous signals in the graph-based approach, P^s_{Eng} is calculated within the model in this case study. Due

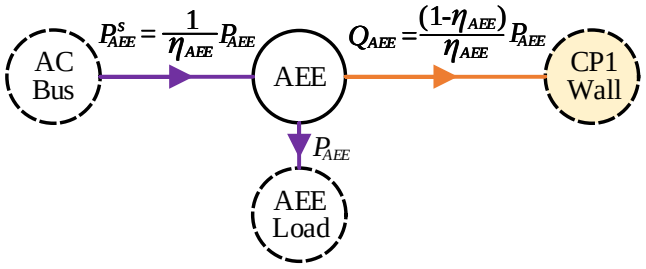

Figure 3.7: Edge equations governing the AEE.

to the algebraic relationships imposed by the graph-based modeling framework for each of the nine algebraic vertices, specifying power flows for seven of the 16 total edges is sufficient to fully define them all. The electrical power edges labeled with signal names in Fig. 3.6 indicate the particular selection for these seven power flows used in this case study. These power flows are considered the controllable inputs to the electrical system. The vector of these inputs is denoted as P_u, while the vector of the remaining nine electrical power flows is denoted as P_{alg}. Note that this does not preclude imposing constraints or objectives on P_{alg} in closed-loop control, as they are still modeled within the graph-based framework. For example, bounds or rate limits on the power extracted from the engine by the generator can still be applied.

A portion of the electrical power through the generator, AEE, battery, and converters goes to waste heat as a function of each component's calculated efficiency. Figure 3.7 exemplifies this for the AEE, where a portion of the source power flow P^s_{AEE} results in thermal energy generation $Q_{AEE} = (1 - \eta_{AEE})\, P^s_{AEE}$, where $\eta_{AEE} \in [0, 1]$ is the electrical efficiency of the component. The remaining source power goes to the AEE load, $P_{AEE} = \eta_{AEE} P^s_{AEE}$. These equations are expressed as a function of P_{AEE} in Fig. 3.7, reflecting the choice of this signal as the associated controllable input in P_u.

As shown in Fig. 3.6, unique edges represent the charging (labeled as P_c) and discharging (labeled as P_d) of the battery, which are treated as distinct operating modes. Unique edges also represent the thermal energy generated in each mode. Figure 3.8 shows which edges are active in each mode, as well as the edge equations used to calculate the thermal energy generation in each case. When discharging, the power flow associated with the edge between the DC bus and battery is assigned a negative sign, indicating power flow in the opposite

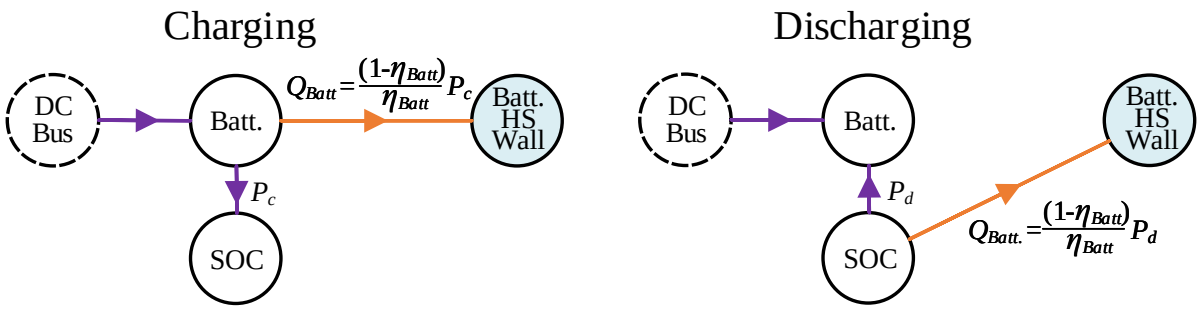

Figure 3.8: Active edges when the battery is charging (left) and discharging (right).

direction of the orientation of the edge, such that the battery is supplying power to the DC bus.

The inactive edges in each mode are zeroed by setting $P_d = 0$ when charging and $P_c = 0$ when discharging. In addition, the battery can be put in an idle mode by setting $P_c = P_d = 0$. Switching of the battery mode can therefore be captured by the first method described in Section 2.6.4. A switching signal σ_{Batt} selects from among three constraint sets for the electrical system inputs, given by:

$$\begin{aligned} \sigma_{Batt} = 1 &\rightarrow \{P_u \mid P_c > 0, P_d = 0\}, && (Charging) \\ \sigma_{Batt} = 2 &\rightarrow \{P_u \mid P_c = 0, P_d > 0\}, && (Discharging) \\ \sigma_{Batt} = 3 &\rightarrow \{P_u \mid P_c = P_d = 0\}. && (Idle) \end{aligned} \tag{3.3}$$

As discussed in Section 3.2, component efficiencies are modeled as either constant, temperature dependent, or current dependent. While this is done to simplify exposition in this example, the general methodology does not preclude a more complex treatment. The generator is assumed to have a constant efficiency. The AEE and battery efficiencies are a function of their heat sink wall temperatures. These efficiencies are modeled as a quadratic function of temperature, $\eta_j(T_j) = c_{1,j} + c_{2,j}T_j + c_{3,j}T_j^2$, where the coefficients $c_{i,j}$ are constant for each i, j. Similarly, the efficiency of each converter is modeled as a quadratic function of its current output to the DC bus. Each current is calculated as a function of the edge power flow assuming a constant voltage, V_j. Therefore, each current is given by $i_j(t) = P_j(t)/V_j$. As shown in Fig. 3.4, the overall efficiency of the converter bank depends on the number of active converters, which are assumed to share the load current evenly. Switching of the

number of active converters can be captured by the first method described in Section 2.6.4. A switching signal σ_{Conv} selects from among three constraint sets for the electrical system inputs, given by:

$$\begin{aligned}
\sigma_{Conv} = 1 &\rightarrow \{P_u \mid P_{Conv,1} \geq 0, P_{Conv,2} = P_{Conv,3} = 0\}, && (\textit{One active converter}) \\
\sigma_{Conv} = 2 &\rightarrow \{P_u \mid P_{Conv,1} = P_{Conv,2} \geq 0, P_{Conv,3} = 0\}, && (\textit{Two active converters}) \\
\sigma_{Conv} = 3 &\rightarrow \{P_u \mid P_{Conv,1} = P_{Conv,2} = P_{Conv,3} \geq 0\}. && (\textit{Three active converters})
\end{aligned} \tag{3.4}$$

As exemplified in Figs. 3.7-3.8, the vector of thermal edge power flows due to waste heat generated by electronics can be written as:

$$Q_w = F_w(\eta) P_u, \tag{3.5}$$

where F_w is linear in the vector of component efficiencies η. However, the efficiencies themselves are in general given by nonlinear functions of electrical power flows or temperature states.

The electrical system is parameterized to match the scale of the fluid-thermal testbed and air bay, such that the DC bus operates at 48 V and the battery has a capacity of 5.8 A·h.

3.3.3 Model Summary and Timescale Analysis

The following summarizes the graph-based modeling effort and resulting system model used to develop the hierarchical control formulation in Section 3.4.

In a similar form as (2.7), the dynamics of the complete electro-thermal system model can be represented as:

$$\begin{bmatrix} C_T \dot{T} \\ C_{SOC} \dot{x}_{SOC} \\ 0 \end{bmatrix} = -\bar{M} \begin{bmatrix} Q_T \\ Q_w \\ P_u \\ P_{alg} \end{bmatrix} + D P^s. \tag{3.6}$$

On the left in (3.6), C_T is the diagonal matrix of capacitances of thermal vertices, T is vector of temperature states of the system, C_{SOC} is the storage capacitance of the battery with

state-of-charge x_{SOC}, and the bottom entry is a vector of zeros representing the nine algebraic electrical vertices. On the right in (3.6), $\bar{M}$ is the upper block of the system incidence matrix defined in (2.3)-(2.4), Q_T is the vector of thermal power flows within the FTMS and bay given by (3.2), Q_w is the vector of waste heat generated by electronics given by (3.5), P_u is the vector of controlled electrical power flows, and P_{alg} is the vector of remaining electrical power flows, which are fully defined by solution of the algebraic rows of (3.6), as discussed in Section 3.3.2. Lastly, D is the mapping defined in (2.5) and $P^s = P^s_{Eng}$ is the lone source power flow to the graph. Switching of the valves, battery, and converters is captured by switching signals σ_v, σ_{Batt}, and σ_{Conv} that select among constraint sets for $\dot{m}$ and P_u.

To put the model equations in a more conventional form, (3.2) and (3.5) can be substituted into (3.6) and the dynamics can be reformulated into an equivalent nonlinear state space representation:

$$\dot{x} = A\left(u\right) x + B\left(x, u\right) u + V\left(u\right) d. \tag{3.7}$$

Here, x, u, and d are the states, inputs, and disturbances, respectively, with

$$x = \begin{bmatrix} T \\ x_{SOC} \end{bmatrix}, \quad u = \begin{bmatrix} \dot{m} \\ P_u \end{bmatrix}, \quad d = T^t. \tag{3.8}$$

The dependence of A and V on u arises from (3.2), in which power flows are a function of $\dot{m}$ while linear and separable in x and d. The dependence of B on x and u arises from (3.5), in which efficiencies η can be nonlinear functions of T and P_u.

In addition to the dynamics of (3.7), the algebraic rows of (3.6) giving P_{alg} can be reformulated as the equality:

$$P_{alg} = Z\left(x, u\right) u, \tag{3.9}$$

where the nonlinearities in Z also arise from the nonlinear dependence of η in (3.5) on T and P_u.

The vertex capacitances C_T in the nonlinear graph-based model (3.6) can be used to categorize the timescale of the associated thermal states. These capacitances are plotted in Fig. 3.9. Vertices associated with large capacitances, such as those of the tanks and

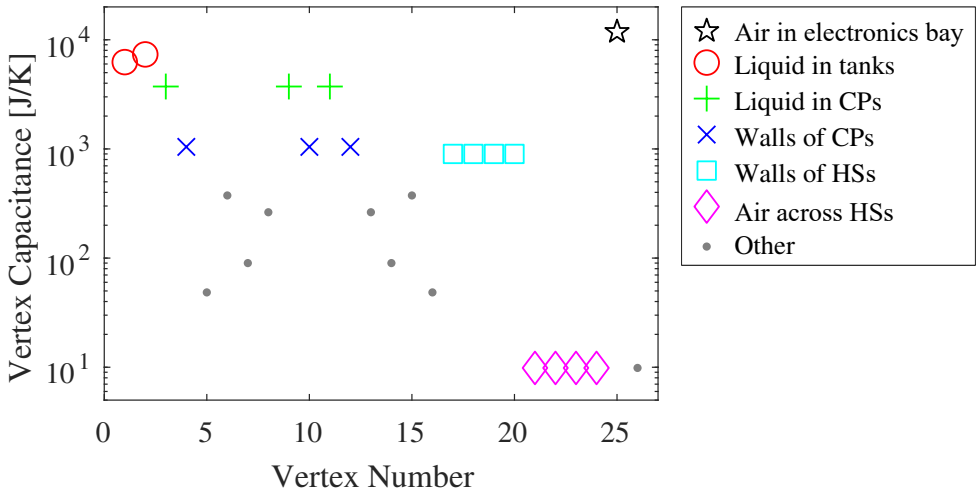

Figure 3.9: For the thermal vertices of the graph in Fig. 3.6, the vertex capacitances identify the slow dynamics associated with the tanks and electronics bay, as well as the fast dynamics associated with the converter and battery heat sinks [101].

electronics bay, evolve over slow timescales. Vertices associated with small capacitances, such as those of the air across the heat sinks, evolve over fast timescales. This timescale separation can be verified by analyzing the eigenvalues of A in (3.7) under nominal operating conditions. The magnitudes of these eigenvalues range from approximately 10^{-4} to 10^{2}.

In addition to governing continuous dynamics across this wide range in timescales, controllers for this system must also select modes of operation for the valves, converters, and battery while maintaining constraints on both inputs and states. Therefore, implementing a centralized MPC design would require formulating and solving a mixed-integer constrained optimization program with a small time step to capture the fast dynamics and a long prediction horizon to capture the slow dynamics. The large computational cost associated with directly solving such programs and the limited computing power on board aircraft may render real time implementation of a centralized controller intractable.

A detailed comparison of the control performance achieved by decentralized PI, centralized MPC, and hierarchical MPC was performed in [64]. This used an alternative topology of the experimental testbed and considers the thermal domain only. With no switching decisions and fewer states and control inputs than the electro-thermal system considered in this chapter, a centralized MPC could be solved in real time. A relatively large 80 second time step was used for the centralized MPC to demonstrate the negative impact of achieving

long time horizons at the cost of requiring large time steps, resulting in poor regulation and disturbance rejection. In accordance with the discussion in Section 3.1.1, hierarchical MPC was shown to be better suited to the control of the multi-timescale fluid-thermal system. By maintaining a long time horizon in the highest level controller and including faster regulation in the lower level controllers, the hierarchical MPC achieved superior closed-loop performance as compared to the decentralized and centralized approaches, with fewer and smaller temperature constraint violations and improved overall system efficiency.

Motivated by the results of [64], the following section provides a hierarchical decomposition of the electro-thermal energy management task such that the resulting control formulations are either quadratic programs (QPs) or mixed-integer quadratic programs (MIQPs) that can be solved in real time. For the particular system and hierarchical control implementation in this case study, the controller update rates and prediction horizons are chosen based on the dynamic timescales observed in [64].

3.4 Hierarchical Control Formulation

3.4.1 Proposed Hierarchical Control Structure

Figure 3.10 shows the three-level hierarchical control structure implemented in this case study. In the "Upper Level" at the top of the hierarchy, a single MPC controller is responsible for coordinating overall vehicle energy management. This controller is designed with a large ($\Delta T = 60$ seconds) time step and a long ($N = 10$ steps) prediction horizon, resulting in the ability to plan electrical and thermal switching, state trajectories, and load shedding over a 10 minute prediction horizon into the future. The planned trajectories are communicated down the hierarchy as references to be tracked.

Below the Upper Level, the "Lower Level" has one controller governing the FTMS, and a second controller governing the electrical system and air bay. This particular choice of decomposition mirrors that in the simulation study of [1], and preserves the strong coupling within the FTMS and between the electrical system and the air bay. While a more thorough treatment of hierarchical model decomposition falls outside the scope of this example, note

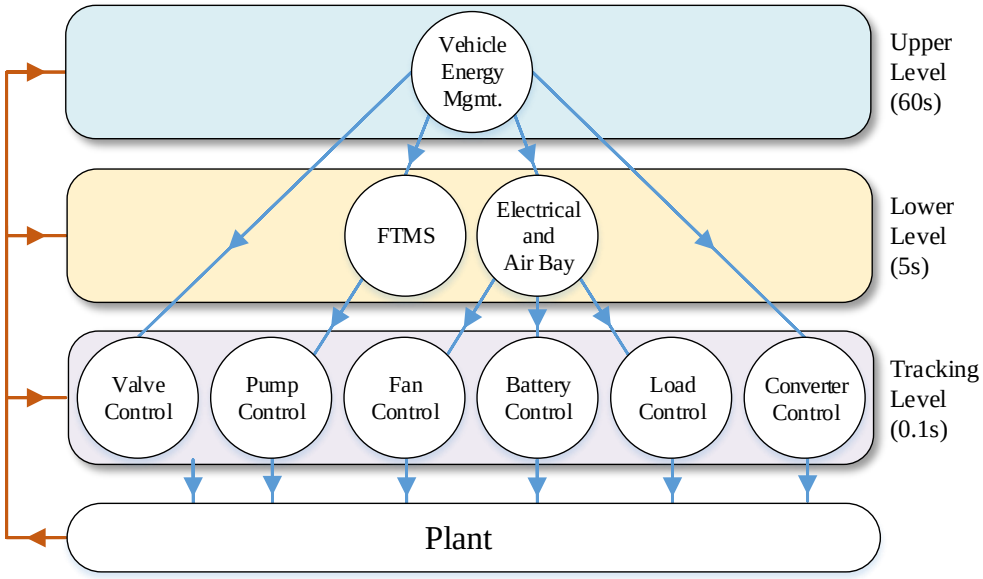

Figure 3.10: Three-level hierarchical control structure implemented in this case study [101].

that an agglomerative hierarchical clustering approach for decomposing graph-based models has been proposed in [92] and subsequently demonstrated in application to hierarchical control of a simulated automotive electric vehicle in [66], as discussed in Section 2.5.1.

Figure 3.11 shows the sub-graphs used to generate a model for each controller of the Lower Level. As indicated in Fig. 3.11, the temperatures of the CP1, CP2, and HX1 walls become sink states of the sub-graph for the electronics and bay, while the heat loads applied to each of these walls become source power flows to the sub-graph of the FTMS. In a decentralized control framework, each Lower Level controller may have to treat these coupling terms as unknown disturbances, incurring significant conservatism. However, in the hierarchical control framework, predicted values for these coupling terms over the time horizon of the Lower Level are included in the information passed from the Upper Level. The resulting mutual knowledge of coupling between the Lower Level controllers is a key enabler of the reduced conservatism and improved performance of the hierarchical framework.

These Lower Level controllers are designed with a smaller ($\Delta T = 5$ s) time step and a shorter ($N = 5$ steps) prediction horizon than the Upper Level, providing improved bandwidth and the ability to compensate for model and signal uncertainty while operating within computational limits [64]. Subject to the valve, battery, and converter mode switching de-

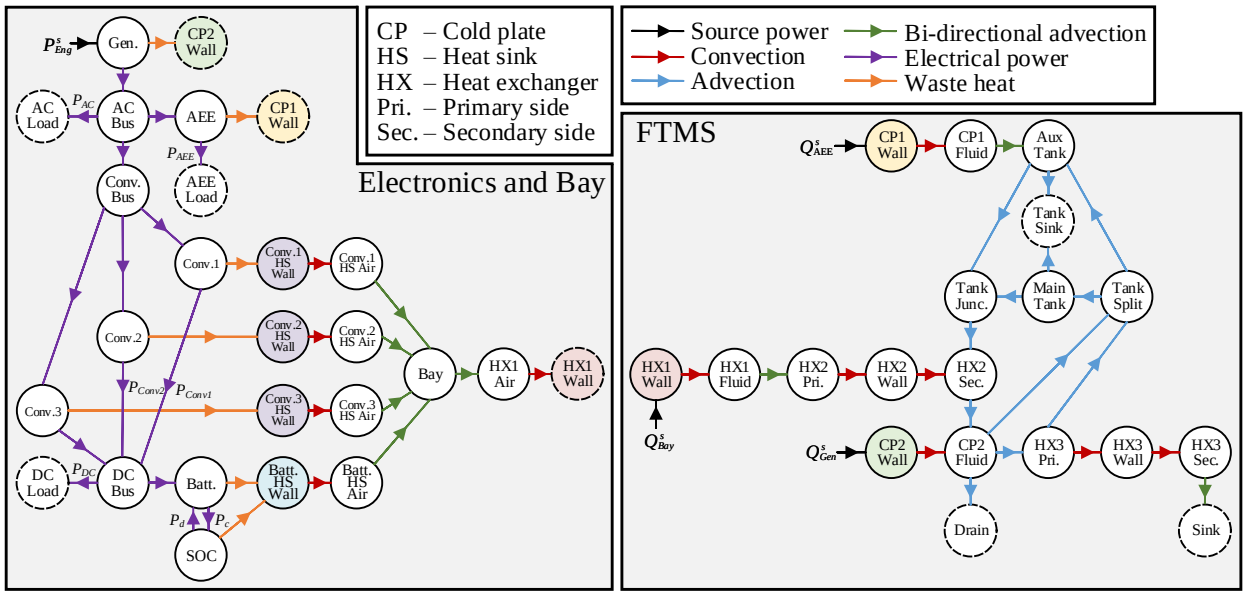

Figure 3.11: Decomposition of the electro-thermal system graph of Fig. 3.6 into an electronics and bay sub-graph (left) and an FTMS sub-graph (right) [101].

cisions made in the Upper Level, the Lower Level controllers make continuous decisions for desired mass flow rates of the pumps and fans, and for the rate of battery charge/discharge. This level can also perform additional load shedding as required to satisfy state constraints.

At the bottom of the hierarchy, the "Tracking Level" consists of decentralized PI and logic-based controllers that determine the pump and fan speeds required to track the desired mass flow rates communicated from the level above, as well as implement the valve, converter, battery, and load control decisions.

The following sections detail the formulations for the Upper Level and Lower Level MPC, including modification of the model in Section 3.3.3 to represent the system dynamics within an MIQP in the Upper Level and QPs in the Lower Level.

3.4.2 Upper Level MPC

3.4.2.1 Switched Linear Representation

While the nonlinear dynamics of (3.7) and algebraic relationships of (3.9) could be discretized and embedded in a nonlinear MPC formulation, solving the resulting optimization program

would incur significant computational cost as compared to that of solving an MIQP or QP of the same order. This would severely limit the update rate and/or prediction horizon that could be used in real time implementation. Therefore, for the Upper Level MPC, (3.7) and (3.9) are approximated by switched linear forms that can be discretized, embedded in an MIQP, and solved in real time with an appropriate update rate and prediction horizon for achieving long-term electro-thermal coordination of continuous and switched behavior.

A primary role of the Upper Level controller is to select modes of operation for the valves, converters, and battery over the prediction horizon. Therefore, it is essential that the model used in this level accurately captures the dynamics of each possible mode. While linearizing about a single operating condition or desired equilibrium is a standard practice in developing MPC controllers for nonlinear systems, the bilinear relationship between temperatures and mass flow rates in the thermal dynamics (e.g., in (2.14)-(2.17)) indicates that this could introduce significant model error. For example, if the mass flow rate in (2.14) is in line with a valve, then closing the valve will zero the mass flow rate, and therefore appropriately zero the associated edge power flows of the nonlinear model. However, if the power flow linearization of (2.30) is performed about an operating condition with the valve open, then zeroing the mass flow rate does not necessarily zero the associated edge power flow of the linear model, as only the third term on the right in (2.30) depends on the mass flow rate. Even in absence of discrete switching of valves, the validation in Section 2.4.4 shows that use of a single linearized model can incur significant error when pump commands are sufficiently far from the values used for linearization.

As a less computationally expensive alternative to using nonlinear optimization in the Upper Level MPC to accurately capture the nonlinear thermal dynamics, this behavior is instead captured by treating the mass flow rate as a switched input. By assuming a fixed nominal command to each pump, $A(u)$ and $V(u)$ in (3.7) are restricted to belong to a finite set of constant matrices $\{A_i, V_i\}$ at each time, where each i corresponds to a given mode of the valves. The active mode is selected by the switching signal for the valves, σ_v. The Upper Level MPC then performs the task of selecting among the valve modes, while the Lower Level MPC performs continuous control of the pumps, subject to the valve topology chosen by the Upper Level. This decomposition of decision-making in the FTMS is supported in the

literature, as optimal decisions for valves are generally governed by long timescale behavior associated with overall vehicle mission segments [19]. For example, heat exchangers between the FTMS of an aircraft and ambient air should be bypassed during the ground hold phase of a flight plan if the ambient temperature on the tarmac exceeds the temperature of the fuel in the tanks. On the timescale of faster thermal dynamics, the continuous range in speeds of pumps and fans can be leveraged to compensate for model and disturbance error, balancing between preventing temperature constraint violations and minimizing the energy consumed to operate these actuators [64].

The remaining nonlinearities in the system model are those in B of (3.7) and Z of (3.9), due in both cases to the nonlinear dependence of the electrical system efficiencies on temperature and current. The dependence on temperature relates specifically to that of the battery HS wall and CP1 wall, which evolve on relatively slow timescales. Therefore, each time the Upper Level MPC is to be solved, the temperature-dependent efficiencies can be calculated using the most recent temperature measurements and assumed with relatively small error to remain constant over the prediction horizon. This strategy aligns with the linear time-varying (LTV) MPC approaches of [46, 102].

While the strategy above provides an appropriate linear approximation of temperature-dependent efficiencies, it is not appropriate for the current-dependent efficiencies of the converters. This is because the converter load currents and number of active converters are control decisions that can change instantaneously, leading to the potential for significant variation in current-dependent efficiencies across the Upper Level MPC prediction horizon. Therefore, for each mode of the converter bank (i.e., for each possible number of active converters), the efficiency is approximated as a piecewise constant function of the load current, as plotted for the mode with one active converter in Fig. 3.4. Similar to the valve modes, this can be represented in the model as a family of matrices $\{B_i, Z_i\}$, where each i corresponds to a given mode of the converter bank σ_{Conv} and range in the commanded Converter 1 load current $P_{Conv,1}$.

Applying the assumptions above, (3.7) and (3.9) can be approximated by the following switched linear dynamic system and switched linear algebraic expression:

$$\dot{x} = A^o_{\sigma(t)} x + B^o_{\sigma(t)} u + V^o_{\sigma(t)} d, \tag{3.10a}$$

$$P_{alg} = Z^o_{\sigma(t)} u, \tag{3.10b}$$

where for simplicity in notation, all switching signals have been coalesced into a single switching variable $\sigma(t)$. The continuous system in (3.10a) can then be discretized with time step ΔT and time index k, yielding the discrete switched linear system:

$$x_{k+1} = A^d_{\sigma_k} x_k + B^d_{\sigma_k} u_k + V^d_{\sigma_k} d_k. \tag{3.11}$$

A zero-order hold is used in discretizing the thermal dynamics, such that the discrete states of (3.11) exactly match the continuous states of (3.10a) at the update times if the disturbances and inputs are constant between updates.

3.4.2.2 Optimization Program

The Upper Level vehicle energy management MPC controller is formulated using the following MIQP:

$$J^*(x_k) = \min_{\boldsymbol{\sigma_k}, \boldsymbol{u_k}, \boldsymbol{s_k}} \sum_{j=0}^{N-1} J_k, \tag{3.12a}$$

subject to, $\forall j \in [0{:}N-1]$,

$$x_{k+j+1|k} = A^d_{\sigma_{k+j|k}} x_{k+j|k} + B^d_{\sigma_{k+j|k}} u_{k+j|k} + V^d_{\sigma_{k+j|k}} d_{k+j|k}, \tag{3.12b}$$

$$P_{alg,k+j+1|k} = Z^o_{\sigma_{k+j+1|k}} u_{k+j+1|k}, \tag{3.12c}$$

$$\underline{x} - \epsilon - s_{k+j|k} \leq x_{k+j+1|k} \leq \bar{x} + \epsilon + s_{k+j|k}, \quad s_{k+j|k} \geq 0, \tag{3.12d}$$

$$\underline{x}_{SOC} \leq x_{SOC,k+j+1|k} \leq \bar{x}_{SOC}, \tag{3.12e}$$

$$H^u_{\sigma_{k+j+1|k}} u_{k+j+1|k} \leq z^u_{\sigma_{k+j+1|k}}, \tag{3.12f}$$

$$H^{alg}_{\sigma_{k+j+1|k}} P_{alg,k+j+1|k} \leq z^{alg}_{\sigma_{k+j+1|k}}, \tag{3.12g}$$

$$u_{k+i+1|k} = u_{k+M_u|k}, \ \forall i \in [M_u, N-1], \tag{3.12h}$$

$$\sigma_{k+i+1|k} = \sigma_{k+M_\sigma|k}, \ \forall i \in [M_\sigma, N-1], \tag{3.12i}$$

$$x_{k|k} = x_k, \tag{3.12j}$$

$$u_{k|k} = u_{k|k-1}, \tag{3.12k}$$

$$\sigma_{k|k} = \sigma_{k|k-1}. \tag{3.12l}$$

The decision variables of this program are the mode sequence $\boldsymbol{\sigma_k} = [\sigma_k, ..., \sigma_{k+M_\sigma}]$, input trajectory $\boldsymbol{u_k} = [u_k, ..., u_{k+M_u}]$ and slack variables $\boldsymbol{s_k} = [s_k, ..., s_{k+N-1}]$ of the N-step prediction horizon. At each time index k, the cost function J_k is a sum of quadratic terms, given by:

$$\begin{aligned} J_k \ = \quad & \|s_{k+j|k}\|^2_{\Lambda_s} + \|x_{k+j|k} - r^x_{k+j|k}\|^2_{\Lambda_x} \\ & + \|u_{k+j+1|k} - r^u_{k+j|k}\|^2_{\Lambda_u} + \|P_{alg,k+j+1|k} - r^{alg}_{k+j|k}\|^2_{\Lambda_P} \\ & + \|u_{k+j+1|k} - u_{k+j|k}\|^2_{\Lambda_{\Delta u}} + \|\sigma_{k+j+1|k} - \sigma_{k+j|k}\|^2_{\Lambda_{\Delta\sigma}}. \end{aligned} \tag{3.13}$$

The first term on the right of this cost function minimizes the non-negative slack variables $s_{k+j|k}$, which ensure feasibility of the state constraints in (3.12d). The second, third, and

fourth terms minimize deviations of states, inputs, and algebraic power flows from desired references $r^x_{k+j|k}$, $r^u_{k+j|k}$, and $r^{alg}_{k+j|k}$, respectively. These references are exogenous to the controller and reflect the desired operation of the system. In particular, time-varying input references are provided for the desired electrical power to the AC, DC, and AEE loads in accordance with the vehicle's mission or flight plan. The controller may reduce the actual power to these loads from the desired value via load shedding as necessary to decrease thermal energy generation and satisfy temperature constraints. However, the power to these loads must remain above a minimum, flight-critical value defined in (3.12f), even at the cost of temperature constraint violation. A reference of zero can be provided for some signals to incentivize efficient operation. For example, this can be done for P^s_{Eng} to penalize power extraction from the engine by the generator. The fifth term of (3.13) serves to smooth the input trajectory, penalizing the change in the inputs between consecutive steps to reduce wear on continuously variable actuators. The sixth term performs a similar role for switching, ensuring that switched actuators only change modes when doing so significantly improves the closed-loop performance.

Equation (3.12b) imposes the switched linear system dynamics from (3.11), while (3.12c) imposes the linear algebraic power flows of the electrical system from (3.10b). Equation (3.12d) imposes softened box constraints on the system states, with lower bound $\underline{x}$ and upper bound $\bar{x}$, where ϵ represents a constraint tightening term to help ensure constraint satisfaction in the presence of model and disturbance uncertainty. For the experimental results in Section 3.6, $\epsilon = 2°\text{C}$ is used to tighten constraints on temperature states, approximately matching the typical open loop model error observed experimentally in Section 2.4.4. Equation (3.12e) imposes a hard constraint on the battery SOC, with lower bound $\underline{x}_{SOC}$ and upper bound $\bar{x}_{SOC}$. It can be confirmed that this hard constraint cannot render the optimization program infeasible by ensuring that the flight-critical loads can always be provided with the battery idle, even under worst-case efficiencies.

Equation (3.12f) constrains the mass flow rate inputs and controllable electrical power flows to fall within their admissible envelopes for each mode. This includes constraints establishing the achievable mass flow rates for each mode of the valves in (3.1). Note that under the assumption that the pumps operate with a fixed nominal command used to lin-

earize the model for the Upper Level MPC in Section 3.4.2.1, each E_i in (3.1) is singleton. Equation (3.12f) also includes the mode-dependent battery and converter constraint sets of (3.3) and (3.4), respectively. Lastly, (3.12f) includes box constraints on the power flows P_u, representing bounds on the power or current allowable through electrical components, as well as the minimum flight-critical power to the AC, DC, and AEE loads. Equation (3.12g) applies similar box constraints to the algebraic power flows P_{alg}.

To reduce computational cost and enable real time implementation, control horizons $M_u, M_\sigma < N$ are introduced in (3.12h) and (3.12i) to perform move blocking [103]. The inputs and mode selection are held constant after the control horizon, reducing the number of decision variables of the program. For the experimental results in Section 3.6, the Upper Level MPC uses $M_\sigma = 3$ and $M_u = 9$.

To enable real time implementation, a one-step delay is imposed to account for the computation time of the controller and the time taken to communicate its solution to the Lower Level controllers. Equation (3.12j) sets the initial state at time k equal to the measured state at time k, but the optimizer is given until time $k+1$ to compute a solution to the MIQP. To compensate for this delay, (3.12k)-(3.12l) set the initial input and mode equal to those of the second step in the optimal trajectory determined at time $k-1$.

3.4.3 Lower Level MPC

In general, the MPC formulation for the Lower Level controllers is very similar to (3.12)-(3.13). However, the Lower Level controllers are constrained to implement the mode selected by the Upper Level MPC. Thus, σ_k is not a decision variable, allowing the Lower Level controllers to be formulated as QPs. As discussed above, the input and state references to the Lower Level controllers are composed of decision variables and interpolated state predictions from the Upper Level. This coordinates the individual actions of the Lower Level controllers such that their behaviors align with the interests of optimizing the entire system. The vertices and edges used to generate the graph-based model for each Lower Level controller correspond to those in the sub-graphs of Fig. 3.11.

For the Upper Level MPC formulation, the nonlinear thermal dynamics of the graph-based model are approximated as a switched linear system by assuming nominal commands for pumps and fans, and associating modes with different valve positions. With the valve mode selected by the Upper Level, the responsibilities of the Lower Level include determining desired mass flow rates for the continuous operation of the pumps and fans. Therefore, the assumption of fixed nominal mass flow rates is removed, and $A(u)$ and $V(u)$ of (3.7) are instead linearized about the current operating condition and selected mode prior to discretization. This results in a model for the Lower Level MPC controllers of the form:

$$x_{k+j+1|k} - x_{k|k} = A_{k|k}\left(x_{k+j|k} - x_{k|k}\right) + B_{k|k}\left(u_{k+j|k} - u_{k|k}\right) + V_{k|k}\left(d_{k+j|k} - d_{k|k}\right). \quad (3.14)$$

Note that the inputs at time k are known as a result of the one-step delay imposed to account for the computation time of the Lower Level controllers, similar to (3.12k) in the formulation for the Upper Level controller.

3.5 Baseline Control Formulation

The baseline controller is representative of the decentralized, reactive control strategies currently implemented in practice for aircraft energy management, as discussed in Section 1.1.3.2. This controller consists of 11 proportional-integral (PI) loops for output regulation, three purely logic-based controllers for mode selection, and a controller regulating the battery charge and discharge that incorporates both logic-based switching and a PI loop. Each PI controller is saturated with anti-windup using upper and lower bounds specific to its actuator. Table 3.2 shows the actuator and measurement pairings for each controller. The specified setpoints and thresholds were designed by iteration in simulation to satisfy all constraints under the nominal load profile in Section 3.6.

The logic-based controllers, in which switching decisions are updated every 60 s, are summarized as follows.

- Valve switching for tanks:
  - If initially closed, V1 and V2 are opened if the Auxiliary Tank temperature exceeds 35°C.
  - If initially open, V1 and V2 are closed if the Auxiliary Tank temperature is less than 30°C.
  - V3 and V4 are always open, allowing fluid flow to and from the Main Tank.
- Valve switching for Sink Loop bypass:
  - If V5 is initially closed and V6 is initially open, their positions are reversed if the Sink Loop temperature exceeds the tank return temperature. This bypasses HX3 to prevent thermal energy from entering the system via the Sink Loop.
  - If V5 is initially open and V6 is initially closed, their positions are reversed if the tank return temperature exceeds the Sink Loop temperature by more than 2°C. This allows thermal energy to leave the system via the Sink Loop.
- Converter selection:
  - The load current is divided equally among all active converters.
  - When all active converter temperatures are below 55°C, the number of active converters is selected to maximize the converter bank efficiency, per Fig. 3.4.
  - An additional converter is activated if any active converter temperature exceeds 55°C, unless all converters are already active.
- Battery charge, discharge, or idle:
  - The battery cannot be charged beyond an upper bound on its SOC, or discharged below a lower bound on its SOC.

– The battery is charged or discharged, within its rate limits and SOC bounds, to maximize the efficiency of the converter bank. This is consistent with the behavior observed of hierarchical MPC in [1].

– To prevent the battery from overheating, its maximum charge and discharge rate limits are scaled linearly between 1 and 0 as a function of temperature when operating between 55°C and 60°C.

Table 3.2: PI and logic-based control loops of the baseline controller.

Actuator	Measurement (temps. unless otherwise specified)	Type	Setpoint/ threshold
P1 - AEE pump	CP1 - AEE wall	PI	40°C
P2 - Bay HX pump	HX1 - Bay HX wall	PI	30°C
P3 - Main pump	CP2 - Generator wall	PI	35°C
F1 - Conv. 1 fan	HS1 - Conv. 1 heat sink	PI	40°C
F2 - Conv. 2 fan	HS2 - Conv. 2 heat sink	PI	40°C
F3 - Conv. 3 fan	HS3 - Conv. 3 heat sink	PI	40°C
F4 - Battery fan	HS4 - Battery heat sink	PI	40°C
F5 - Bay HX fan	Bay air	PI	30°C
V1-4 - Tank valves	Main and Aux. tank	Logic	35°C
V5-6 - Sink valves	ΔT across HX3 inlet flows	Logic	2°C
No. of active conv.	Converter currents and temps.	Logic	17Amps, 55°C
Battery power	Battery SOC, HS4 - Battery heat sink temp., Conv. 1 current	Logic/PI	[0,1], [55,60]°C, 17Amps
AC load shedding	Main Tank	PI	38°C
AEE load shedding	Aux. Tank	PI	38°C
DC load shedding	HS1 - Conv. 1 heat sink	PI	58°C

3.6 Experimental Results

3.6.1 Testing Scenario

For clarity of exposition, the performance of the hierarchical and baseline controllers is compared for a 30 minute mission segment consisting of sheddable and flight-critical (i.e., non-sheddable) profiles for the AC, AEE, and DC loads, as shown in Fig. 3.12. This is intended to represent only a portion of a complete flight profile of an aircraft energy system, such as that simulated in [1].

The loads in Fig. 3.12 represent a nominal load profile. To compare the capability of the hierarchical control framework and baseline controller to achieve higher loading, each approach is applied for five scaled load profiles, where the magnitudes of the nominal sheddable loads are scaled by 1, 1.5, 2, 2.5, and 3, while the flight-critical loads are left unscaled. In accordance with the HIL testbed schematic of Fig. 3.2, these load profiles represent the mission of the electro-thermal energy management system.

Each optimization program within the hierarchical MPC framework is solved in real time with the update interval indicated in Fig. 3.10 and Table 3.3. All programs were formulated using the YALMIP toolbox [104] and solved with the Gurobi optimization suite [105] using a desktop computer with a 3.40 GHz Intel i7 processor and 16 GB of RAM. The top-down structure of communication within the hierarchy allows controllers of the framework to be solved in parallel. Table 3.3 also provides the average and peak computation times of the Upper Level and Lower Level controllers from the implementation in this section. With a maximum computation time of 50.4 s, the Upper Level can require the majority of its 60 s update interval to solve the MIQP. As discussed in Section 3.4.2.1, approximation of the nonlinear system dynamics by a switched linear system proved to be a key enabler of real time implementation, as solving a nonlinear program of the same order with the same hardware would likely require longer than the update interval. Similarly, as discussed in Section 3.3.3, a centralized MPC approach would not be able to simultaneously achieve the 10 minute time horizon of the Upper Level and 5 s update interval of the Lower Level while solving in real time on the same hardware. Compromising on either of these these parameters

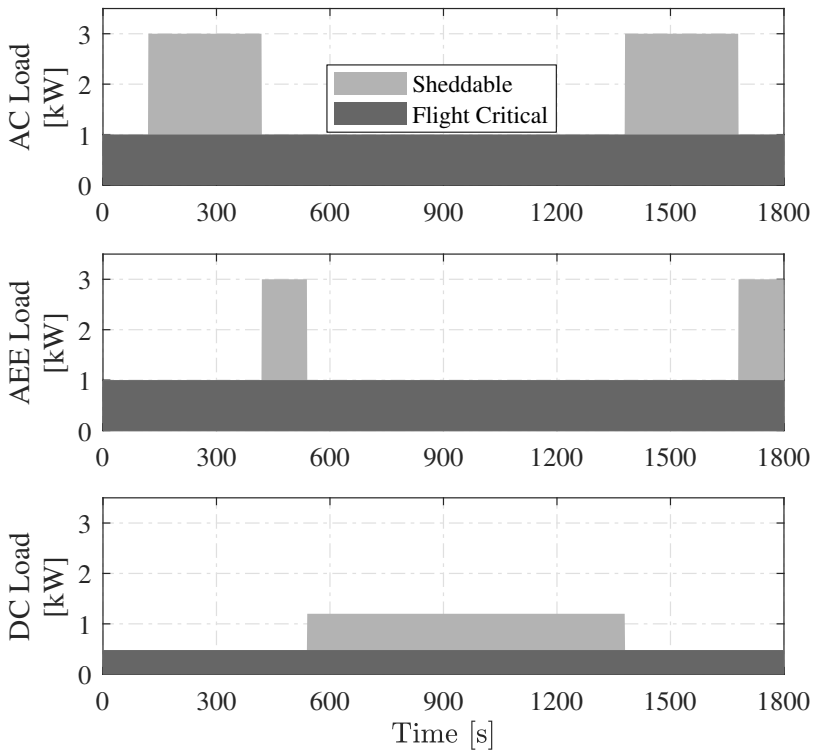

Figure 3.12: Nominal sheddable and flight critical load profiles for the AC, AEE, and DC loads [101].

to enable real time solution would significantly degrade the controller performance. While the modern desktop computer used in this demonstration has greater computational power than what is available on board many legacy aircraft, future research efforts will include the development of design methods for hierarchical MPC that account for the available computational resources in selecting update rates and prediction horizons throughout the hierarchy.

Table 3.3: Update intervals and computation times for all controllers of the hierarchical MPC framework.

Level	Controller	Update interval (s)	Average (s)	Maximum (s)
Upper (MIQP)	Vehicle Energy Mgmt.	60	28.1	50.4
Lower (QP)	FTMS	5	0.03	0.07
Lower (QP)	Electrical and Air Bay	5	0.03	0.09

3.6.2 Performance Comparison

This section provides an overall comparison of the performance of the hierarchical and baseline controllers across the five scaled sheddable load profiles. A more detailed discussion of the behavior of each controller at the largest scale factor is provided for the baseline controller in Section 3.6.3, and for the hierarchical controller in Section 3.6.4.

As defined in Section 3.1, three high-level control objectives for aircraft energy management are to maximize capability, satisfy constraints, and minimize fuel consumption. To evaluate the relative performance of the hierarchical and baseline controllers, these objectives have been translated into four quantitative figures of merit (FoM). In this case study, maximizing capability corresponds to minimizing load shedding of the AC, AEE, and DC loads over the flight segment. Therefore, the first FoM is quantified as the percentage:

$$\text{FoM}_1 = 100\frac{\int_t \text{Achieved Sheddable Load Power } dt}{\int_t \text{Desired Sheddable Load Power } dt}. \tag{3.15}$$

The second and third FoM quantify the ability to satisfy thermal constraints. The second FoM is calculated as:

$$\text{FoM}_2 = \sum_{i\in[1:N_{v,temp}]} \int_t \max\left(0, x_i(t) - \bar{x}_i\right) dt, \tag{3.16}$$

and measures the integral of all temperature constraint violations over the mission. The third FoM equals the sum of the peak violations of all constrained temperatures:

$$\text{FoM}_3 = \sum_{i\in[1:N_{v,temp}]} \max_t \left(\max(0, x_i(t) - \bar{x}_i)\right). \tag{3.17}$$

Lastly, as engine modeling falls outside the scope of this case study, the objective of minimizing fuel consumption for efficiency is instead quantified by the total energy extracted from the engine by the generator over the mission:

$$\text{FoM}_4 = \int_t P^s_{Eng}(t) dt. \tag{3.18}$$

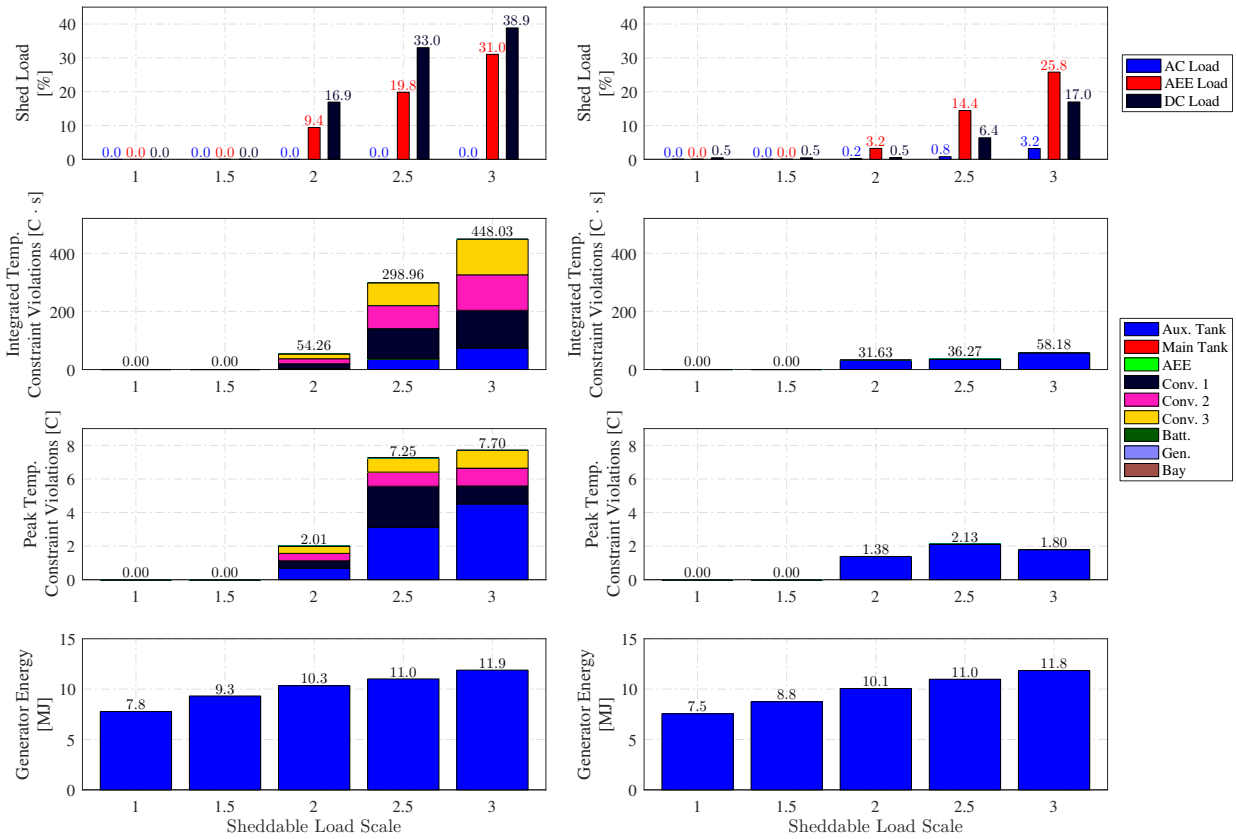

(a) Baseline (b) Hierarchical

Figure 3.13: Comparison of the four FoM used to evaluate the performance of the baseline and hierarchical controllers across five sheddable load scale factors. At high scale factors, the hierarchical controller exhibits significantly reduced load shedding, violates fewer constraints, and extracts slightly less energy from the engine using the generator [101].

Figure 3.13 compares these four FoM for the baseline and hierarchical controllers over the five scaled sheddable load profiles. The first row of subplots shows that the hierarchical approach sheds much less of the loads at high sheddable load scale factors. Note that the small degree of load shedding by the hierarchical controller at low scale factors occurs because load shedding is part of a weighted cost function in (3.13). Despite highly weighted penalties on load shedding, the controller may choose to shed loads to help minimize other terms in the cost function. If desired, the MPC formulations can be modified to prevent this behavior by associating load shedding with a binary decision variable, as in [23].

The second and third rows of subplots in Fig. 3.13 show the FoM corresponding to integrated and peak temperature constraint violations, respectively. As mentioned previously

for the baseline controller in Section 3.5, both controllers are designed to satisfy constraints under the nominal sheddable load profile. However, at higher sheddable load scale factors, the baseline controller violates temperature constraints significantly in several components. By comparison, the hierarchical controller maintains all but one of the temperature constraints by strategically shedding some of the sheddable loads and by better coordinating the conversion, storage, and transport of energy throughout the electro-thermal system. The lone constraint violation under the hierarchical controller occurs in the temperature of the Auxiliary Tank. This can again be attributed to the competing objectives between cost function terms, as well as error in the linearized models used by the controllers, and remains relatively small across all the power levels. For many components, small and short duration temperature constraint violations may impact the longevity and reliability of the component but do not necessarily cause immediate failure. For critical components that cannot tolerate even minor constraint violations, additional constraint tightening could be employed to prevent these, at the cost of increased conservatism.

Lastly, the fourth row of subplots in Fig. 3.13 shows the total energy extracted from the engine by the generator. At all power levels, the hierarchical controller extracts slightly less energy than the baseline controller. At high sheddable load scale factors, this occurs despite the hierarchical controller providing much more of the desired the sheddable load power, highlighting the ability of the hierarchical framework to greatly improve the overall system efficiency.

The comparison of the four FoM between the two controllers for the sheddable load scale factor of three is summarized by the radar chart shown in Fig. 3.14. Each axis is normalized by the value for the baseline controller, and closer proximity of each axis to the origin represents improved performance. With similar generator energy consumption, the hierarchical controller sheds only 69% of what the baseline sheds while simultaneously having integrated and peak temperature violations of 13% and 23% those of the baseline controller, respectively.

In summary, the comparison of controllers in this section demonstrates the ability of coordinated, model-based control of electrical and thermal systems to improve energy management as compared to traditional approaches. While the setpoints for the PI and logic-based

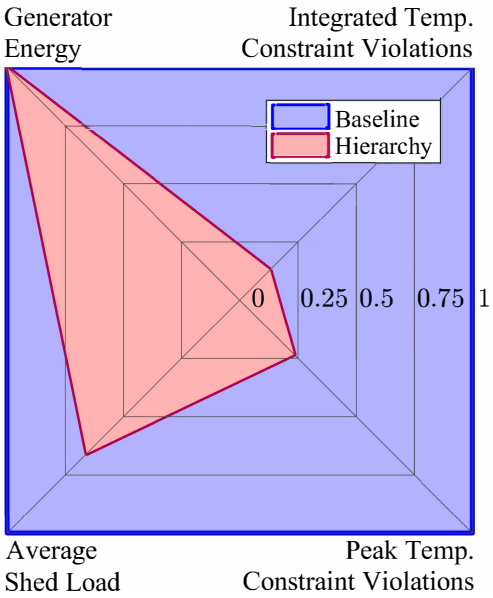

Figure 3.14: Radar chart showing the relative performance of the proposed hierarchical controller as compared to the baseline controller for a sheddable load scale factor of three [101].

controllers of the baseline approach were effective at nominal power levels, this approach became inefficient and ineffective at more taxing, off-nominal operating conditions. In comparison, the model-based hierarchical controller is less sensitive to the particular operating condition and is able to achieve superior performance under both nominal and off-nominal operation. In this way, hierarchical control can facilitate the accelerating electrification of vehicle energy systems, as discussed in Section 1.1.1, allowing energy density and power density to be increased.

3.6.3 Baseline Controller Results

This section provides a more detailed discussion of the performance of the baseline controller for the sheddable load scale factor of three. Figure 3.15(a) shows the thermal actuator commands of the baseline controller, consisting of the valve, pump, and fan commands. Figure 3.15(b) shows a selection of the electrical actuator commands associated with the converter bank and battery. Figure 3.16(a) shows a selection of temperatures under the baseline controller, and Fig. 3.16(b) shows the sheddable, flight critical, and applied loads.

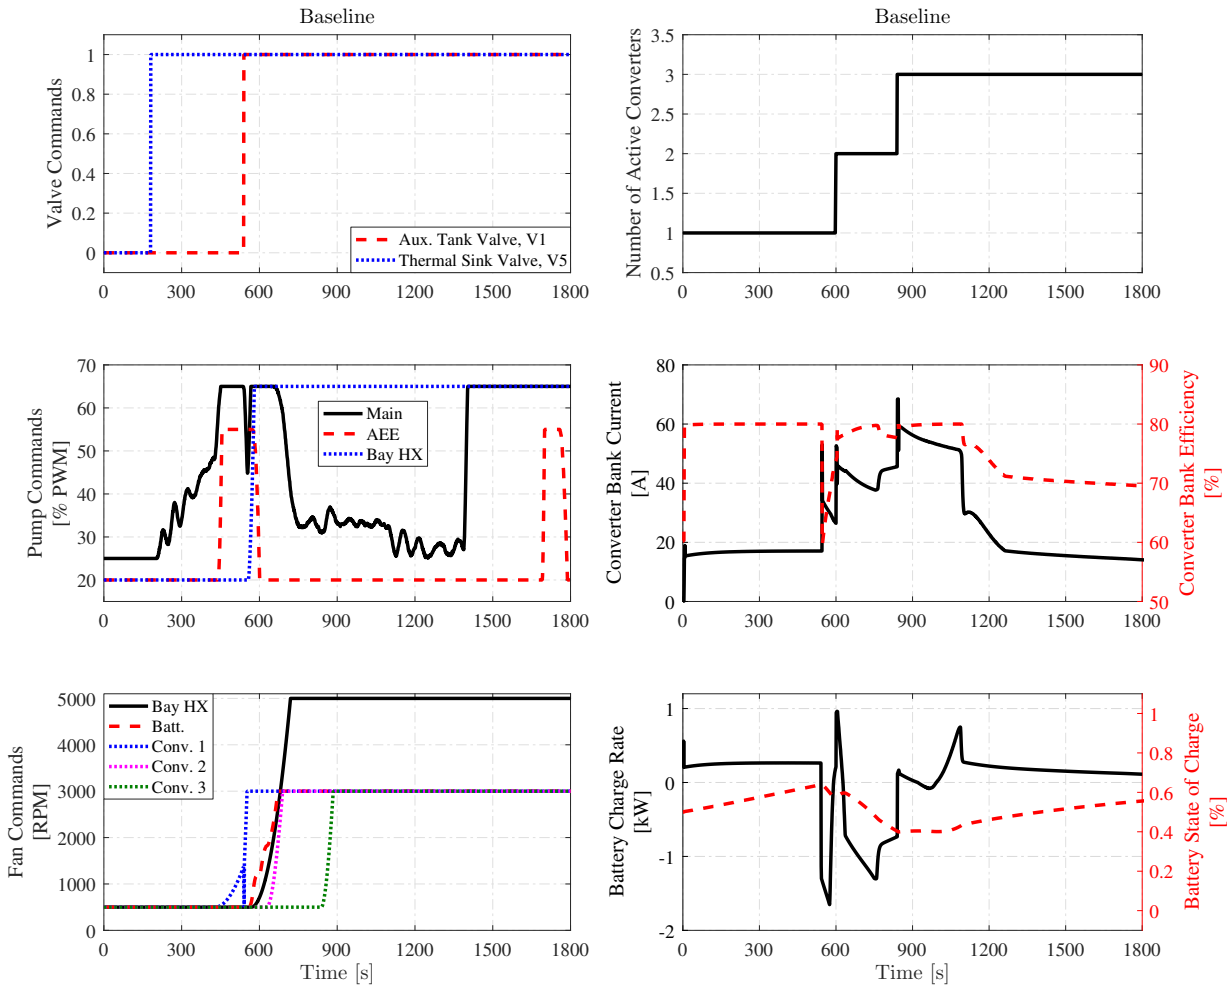

(a) Thermal actuator commands

(b) Selected electrical actuator commands

Figure 3.15: Selected actuator commands under the baseline controller for a sheddable load scale factor of three.

The top subplot of Fig. 3.15(a) shows that the valve allowing thermal energy exchange with the Sink Loop is opened early in the profile, as the fluid in the Sink Loop is colder than that in the return line to the Main Tank. Similarly, once the Auxiliary Tank begins to increase in temperature due to an increase in the AEE load, the valves allowing fluid from the Auxiliary Tank to mix with fluid in the Main Tank are opened. This occurs after the fluid in the Auxiliary Tank is hotter than that in the Main Tank, as shown in the top subplot of Fig. 3.16(a). The middle and bottom subplots of Fig. 3.15(a) show that the AEE pump generally runs high when the AEE load is high, while the Bay HX pump and all fans saturate after the DC load increases and causes the battery, converters, and bay air to

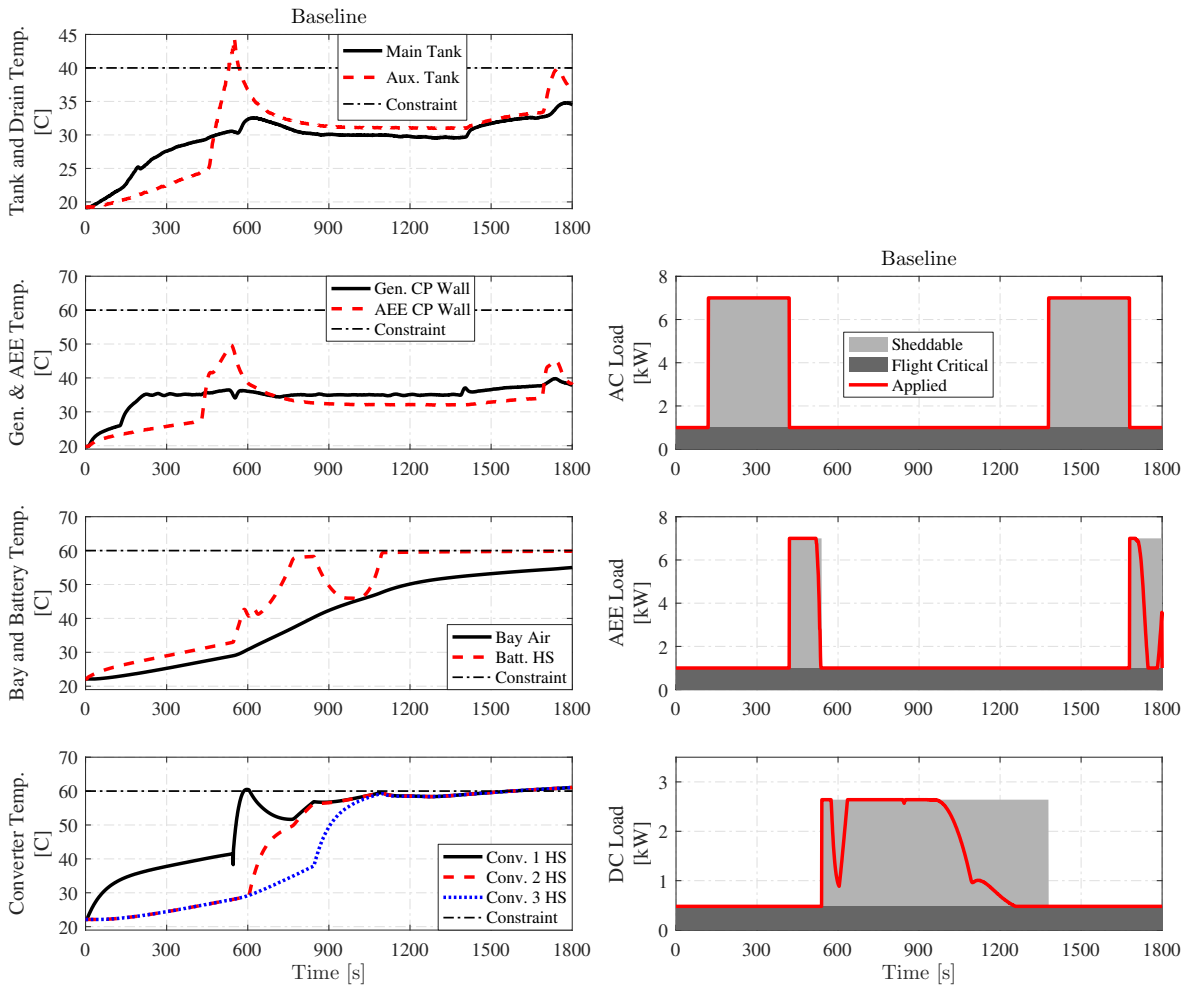

(a) Selected temperatures

(b) Sheddable, flight critical, and applied loads

Figure 3.16: Selected temperatures and loads under the baseline controller for a sheddable load scale factor of three.

increase in temperature, as shown in the bottom two subplots of Fig. 3.16(a). Note that the oscillations in the Main pump could be reduced by using a less aggressive tuning in the PI controller for this actuator.

The top subplot of Fig. 3.15(b) shows that additional converters are activated with time to reduce the load current on active converters as they approach their temperature limits. This results in the cascading temperature increases in the bottom subplot of Fig. 3.16(a). The middle subplot of Fig. 3.15(b) shows that the converter bank current is high during the step in the DC load. The converter bank generally achieves its peak efficiency of 80% for the first two-thirds of the profile, aided by supplemental charging or discharging of the battery,

as shown in the bottom subplot of Fig. 3.15(b). However, the converter bank efficiency decreases to approximately 70% by 1200 s, when three converters must be active due to their high operating temperature and the battery is unable to significantly help bring the converters to a more efficient current due to being thermally limited as well.

Figure 3.16(a) shows that the most significant temperature violation occurs in the Auxiliary Tank, which violates its constraint by approximately 5°C. This constraint violation occurs 10 s after load shedding begins in the AEE and 10 s before the conclusion of the step in the sheddable AEE load, but continues to grow even after the step in sheddable load has concluded as the AEE CP wall transfers thermal energy to the Auxiliary Tank. This illustrates the drawback of a purely reactive control approach, in which waiting to take corrective action until a constraint violation is imminent may not afford enough time to avoid detrimental behavior. At the end of the profile, all three converter HSs exceed their temperature constraints by more than 1°C, and the extent of this violation is increasing even though all thermal actuators of the bay and the Main pump are saturated at their upper limits and there is no sheddable DC load applied. Therefore, continued operation of the system beyond this point would lead to even more significant violations of these constraints.

Lastly, Figure 3.16(b) shows that the most significant shedding occurs in the second half of the step in the DC sheddable load, when the converters and battery are thermally limited beyond their ability to achieve the desired load profile. A significant portion of the second step in the AEE sheddable load is also shed in the last 100 s of the profile.

3.6.4 Hierarchical Controller Results

This section provides a more detailed discussion of the performance of the hierarchical controller for the sheddable load scale factor of three. Figure 3.17(a) shows the thermal actuator commands of the hierarchical controller, consisting of the valve, pump, and fan commands. Figure 3.17(b) shows a selection of the electrical actuator commands associated with the converter bank and battery. Figure 3.18(a) shows a selection of temperatures under the hierarchical controller, and Fig. 3.18(b) shows the sheddable, flight critical, and applied loads.

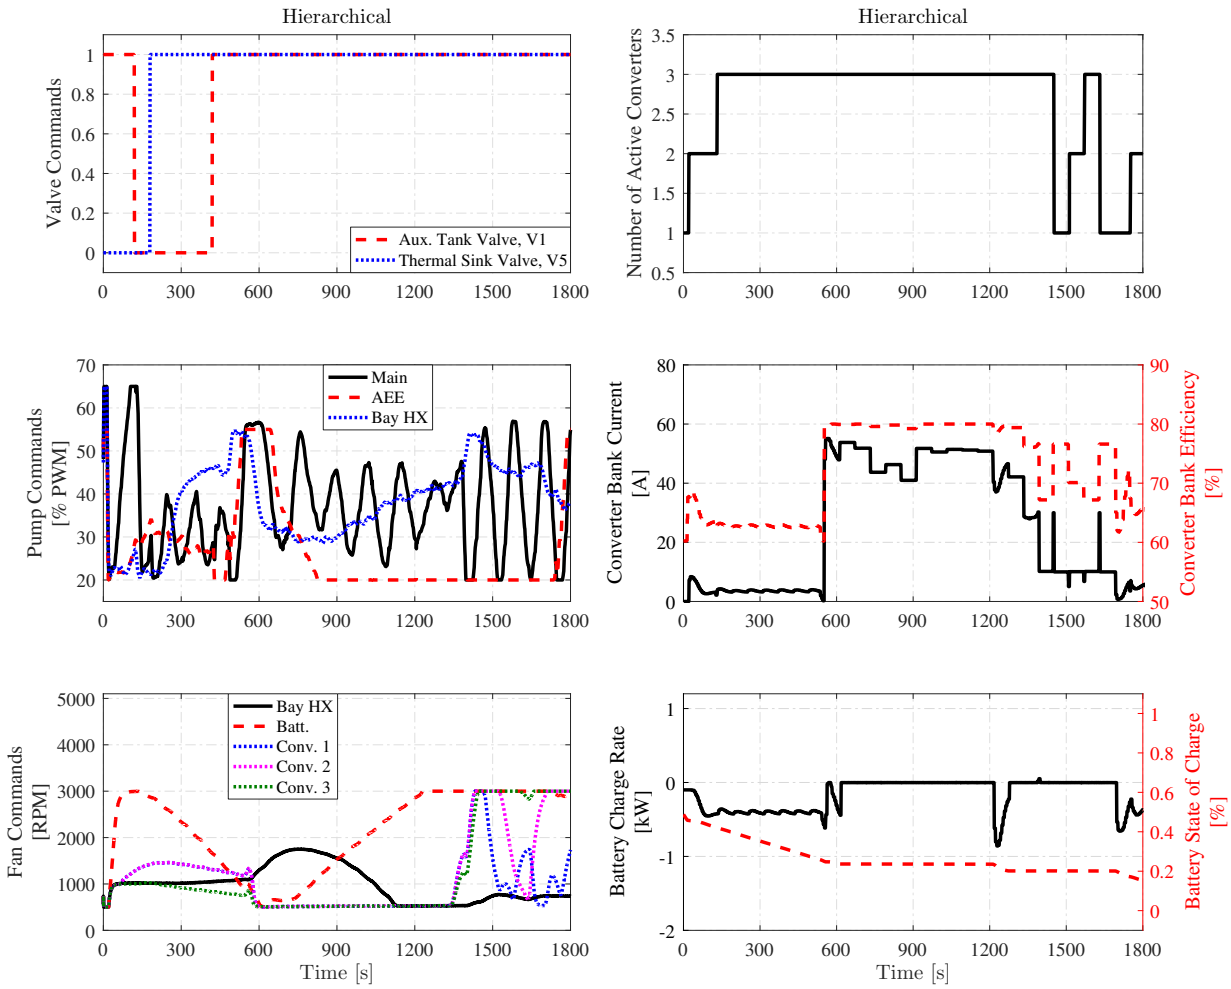

(a) Thermal actuator commands

(b) Selected electrical actuator commands

Figure 3.17: Selected actuator commands under the hierarchical controller for a sheddable load scale factor of three.

Modulo a differing initial choice for the valves associated with the Auxiliary Tank, the hierarchical controller makes similar decisions for the valves as the baseline controller, as shown in the top subplot of Fig. 3.17(a). However, prediction in the hierarchical MPC allows the spike in the Auxiliary Tank temperature that occurs with the first step in the sheddable AEE load to be mitigated. The valves allowing fluid from the Auxiliary Tank to mix with fluid in the Main Tank are opened more than 100 s earlier than under the baseline controller. This, together with a lower fluid temperature in the Main Tank, results in a peak violation of the Auxiliary Tank temperature constraint that is approximately 40% of that under the baseline controller, as shown in the top subplot of Fig. 3.18(a). The extent

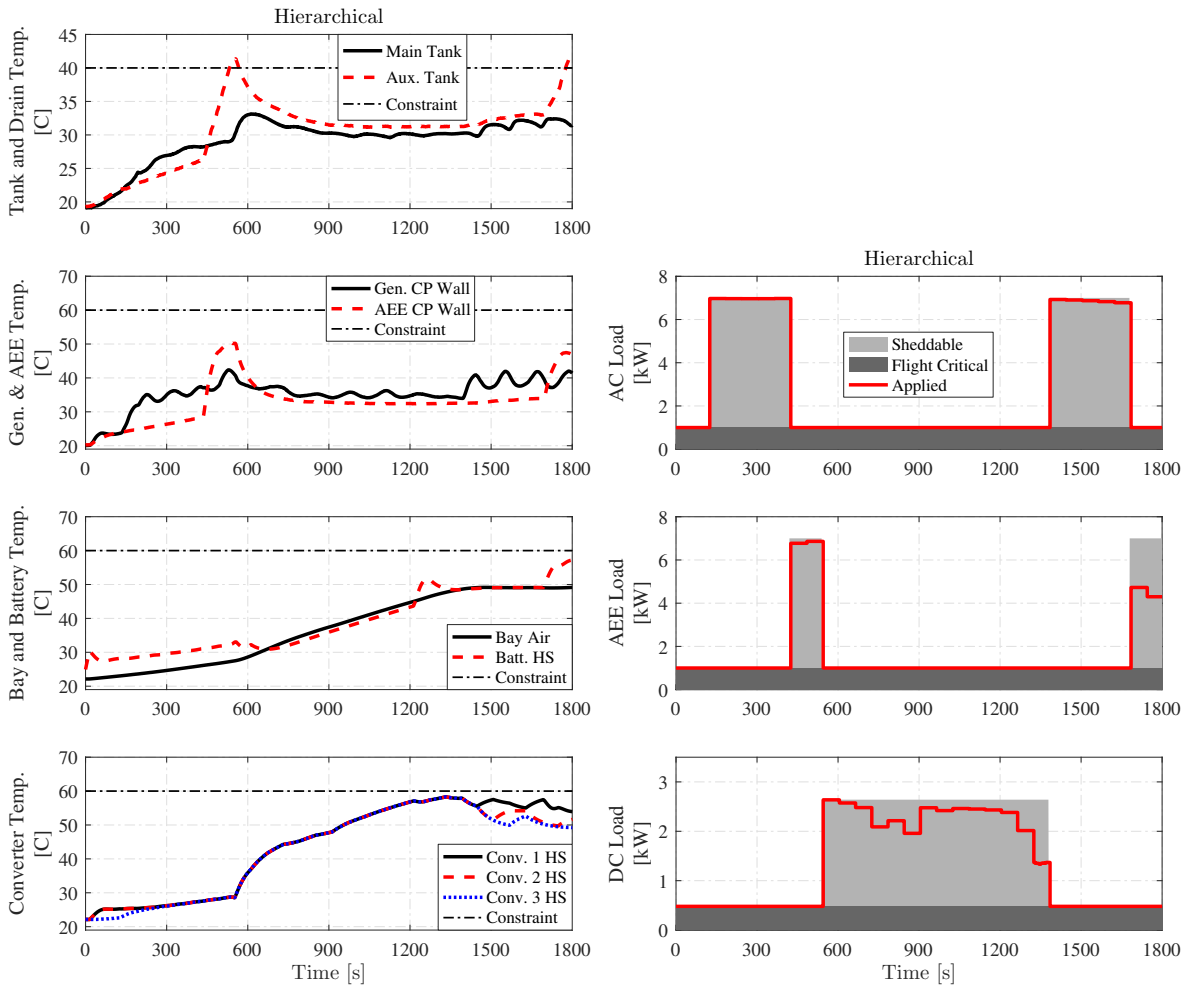

(a) Selected temperatures

(b) Sheddable, flight critical, and applied loads

Figure 3.18: Selected temperatures and loads under the hierarchical controller for a sheddable load scale factor of three.

of this violation could be decreased further by performing additional constraint tightening to compensate for model error or by increasing the weighting on the corresponding slack variables of the MPC formulations.

As shown in the middle subplot Fig. 3.17(a), the AEE pump behaves similar in the hierarchical controller as for the baseline controller. However, while the other pumps and fans are more active early in the profile than for the baseline controller, they generally are operated at much lower values throughout the final two-thirds of the profile. This is especially true for the Bay HX fan. While minimizing energy consumption by these actuators was not a focus of this case study, these results suggest that the hierarchical controller could do so sig-

nificantly better than the baseline controller, indicating additional advantages in efficiency beyond those discussed in Section 3.6.2. The oscillations in the Main pump, which also occurred under the baseline controller, could be reduced by using a less aggressive tuning for the corresponding PI loop in the Tracking Level of the hierarchy.

The top subplot of Fig. 3.17(b) shows that the hierarchical controller switches the number of active converters more frequently than the baseline controller near the end of the profile. As shown in the middle subplot of Fig. 3.17(b) and bottom subplot of Fig. 3.18(a), this switching strikes a balance between operating the converter bank efficiently and reducing the load current on individual converters for the purposes of thermal management. The converter bank operates at peak efficiency of 80% when under high current due to the step in the sheddable DC load. Lower efficiencies only occur when the bank current is small. For example, this happens in the first third of the profile, during which the battery is discharged to provide much of the power to the flight critical DC load, as shown in the bottom subplot of Fig. 3.17(b). The battery is discharged occasionally later in the profile to supplement the power of the converters, improving their efficiency or reducing their current to reduce the heat load on their HSs. Due to the hierarchical controller's more effective leveraging of the energy stored in the battery, the SOC of the battery at the end of the profile under the hierarchical controller is about a third of that under the baseline controller. In performing energy management for a complete flight plan, the objective functions of the MPC formulations for the hierarchical controller could be augmented to encourage charging of the battery when loading elsewhere is relatively small so that the battery has sufficient charge to provide supplementary power when most effective throughout the flight.

Figure 3.18(a) shows that all temperatures except that of the Auxiliary Tank are maintained within their constraints, and that the Auxiliary Tank has peak and integrated constraint violations much smaller than under the baseline controller. The thermal state of the system at the end of the profile is improved as compared to that under the baseline controller, with the converter HSs well under their constraints.

Figure 3.18(b) shows that the hierarchical controller performs significant shedding of the second step in the sheddable AE load and the step in the sheddable DC load, as also done by the baseline controller. However, while the baseline controller tends to supply the full

sheddable load early in the steps and is later forced by thermal limitations to taper this to zero, the hierarchical controller performs shedding earlier in the steps and does not have to perform as significant tapering later on. A slight tapering does occur towards the end of the step in the sheddable DC load, likely because the duration of the step is longer than the prediction horizon of the Upper Level MPC. However, much more of the total sheddable power is still provided under the hierarchical controller than under the baseline controller. Furthermore, the ability to use prediction to provide more consistent shedding across the duration of steps is an important feature of the hierarchical controller. For some systems, load shedding may not be a continuous decision, but instead a discrete decision in which either no power or the full desired power must be provided. In such cases, the prediction in the hierarchical controller can provide an assessment of whether providing the full desired power will fall within the predicted electrical and thermal capabilities of the vehicle over the duration of the load.

3.7 Chapter Summary

This chapter presents and experimentally demonstrates a switched hierarchical control framework for vehicle energy management. The candidate system for this demonstration consists of an HIL experimental testbed that embodies key dynamic characteristics of aircraft electro-thermal systems. The graph-based modeling approach of Chapter 2 allows complex dynamics in both the thermal and electrical domains to be represented using a common framework, and captures the coupling between these domains. The resulting switched nonlinear model is then approximated using a switched linear representation. This enables real time implementation of the MPC controllers of the hierarchy, formulated as MIQPs or QPs. As compared to a baseline controller representative of the traditional control approaches discussed in Section 1.1.3.2, the hierarchical controller achieves improved performance in figures of merit quantifying the capability, safety, and efficiency of the closed-loop system. While the case study used for demonstration in this chapter focuses on a specific candidate electro-thermal system and hierarchical control architecture, the underlying methodology is generalizable. Application to other architectures, vehicle classes, and physical domains is a focus of ongo-

ing and future research. For example, application to an automotive electric vehicle energy system has been performed in simulation in [66].

Chapter 4

Passivity-Based Stability Under Switching

4.1 Motivation

The case study in Chapter 3 demonstrates that hierarchical MPC frameworks can increase the capability, safety, and efficiency of vehicle energy systems as compared to traditional approaches. While the aggressive decision-making of hierarchical MPC frameworks enables high performance, this also brings a need for formal analysis guaranteeing closed-loop stability in safety-critical applications. The necessity of accounting for potentially destabilizing interactions among controllers of the hierarchy and sub-models of the plant can make this a challenging task [40]. While MPC enables significant performance benefits, its lack of a closed-form solution results in further challenges to analysis. The presence of switching brings a third complication. Even when the continuous dynamics under each mode of operation can be proven stable under closed-loop control, this does not guarantee that the overall switched system is stable [26]. In this chapter, the structural information embedded in the graph-based modeling approach of Chapter 2 is shown to facilitate stability analysis of interconnected switched nonlinear systems and their closed-loop control under decentralized, distributed, or hierarchical switched MPC frameworks. This is achieved by analyzing the passivity of graph-based systems and augmenting their controllers with a local passivity-based constraint.

4.2 Background

The nonlinear graph-based modeling approach of Chapter 2 is derived by applying conservation equations to a dynamic system. As a result, this modeling approach naturally pairs with

analysis tools from dissipativity theory, which was developed to characterize the storage of energy in a system as a function of its inputs and outputs, and can be used to investigate the Lyapunov stability of nonlinear systems [106–111]. Passivity, as a special case of dissipativity, has been applied to support the analysis and control of many energy systems. Examples include building thermal systems [84, 112], liquid level systems [113, 114], electrical power systems [60, 115], biochemical reaction networks [116], and mechanical systems [117].

Passive behavior under closed-loop control can be ensured using Model Predictive Control (MPC), as has been done with centralized approaches in [118–122] and distributed approaches in [57, 123]. However, these distributed approaches, as in the analysis of [116, 124], require satisfaction of global condition on the structure and gain of couplings among subsystems to ensure stability. This is not required in the decentralized MPC approach for passive graph-based models of [125], which shows that the interconnection of subsystems in this class forms negative feedback loops, preserving passivity under this coupling.

Passivity-based analysis and control has also been investigated for switched systems, and can be broadly classified by the use of a common storage function [126–128] or multiple storage functions [129–132]. While the latter approach can be employed in cases where all modes in a family of switched systems do not share a common storage function, this can introduce additional complexity to the analysis, such as the cross-supply rates of [130] that characterize the change in energy of inactive modes under the trajectory governed by the active mode. As shown in this chapter, the use of a common storage function is sufficient to prove the passivity of a class of switched graph-based systems. While the literature cited above indicates that distributed passivity-based MPC and passivity under switching have been studied independently, this chapter presents a unique result that contributes to both these areas, including a passivity-based guarantee of stability for switched systems controlled by decentralized MPC. As achieved in the absence of switching in [125], no global constraints are required. Instead, a local passivity-based constraint within each controller is proven sufficient to guarantee closed-loop stability.

The remainder of this chapter is organized as follows. Section 4.3 describes the switched graph-based modeling approach. Section 4.4 proves that a common storage function can be employed to certify passivity under switching for this class of systems, shows that pas-

sivity holds under the interconnection of multiple graph-based subsystems, and presents a decentralized MPC formulation that stabilizes the closed-loop system without requiring any global analysis of passivity to be performed. Section 4.5 demonstrates the proposed approach in simulation on a fluid tank system similar to that introduced in Chapter 2. Section 4.6 provides a concluding summary of the chapter.

4.3 Class of Systems

4.3.1 Nominal (Non-Switched) Graph-Based Systems

Consistent with the graph-based model formulation introduced in Chapter 2, the dynamic associated with each vertex of a graph is governed by the conservation equation:

$$C_i \dot{x}_i = \sum_{\{j | e_j \in e_i^{head}\}} P_j - \sum_{\{j | e_j \in e_i^{tail}\}} P_j, \tag{4.1}$$

where $C_i > 0$ is the storage capacitance of the vertex and each P_j is the power flow associated with an edge incident to the vertex.

Also as in Chapter 2, the dynamics of the complete system $\mathbf{S}$ can be written as:

$$\mathbf{S}: \quad C\dot{x} = -\bar{M}P + DP^s, \tag{4.2}$$

where $C = diag\left([C_i]\right)$, $\bar{M}$ is a partition of the incidence matrix of the graph, P^s is the vector of source power flows, and D is a mapping from the source edges to the vertices to which they are incident.

In a slight departure from the graph-based model formulation of Chapter 2, a more restricted form of the edge power flow equation (2.2) is employed in this chapter. Specifically, each power flow is assumed to follow the input-affine form below.

Assumption 4.1:

For all power flows of $\mathbf{S}$, the power flow P_j along edge e_j is given by:

$$P_j = f_j\left(x_j^{head}, x_j^{tail}\right) + g_j\left(x_j^{head}, x_j^{tail}\right) u_j, \tag{4.3}$$

where $f_j, g_j \colon \mathbb{R} \times \mathbb{R} \to \mathbb{R}$, $\left(x_j^{tail}, x_j^{head}\right)$ are the states associated with the vertices to which e_j is incident, and u_j is an associated actuator input.

The following continuity assumptions are made to formalize the analysis. For simplicity of exposition, the dynamics are also assumed to have an equilibrium at the origin.

Assumption 4.2:

For all j, f_j is twice continuously differentiable (and therefore locally Lipschitz [110]) with $f_j(0,0) = 0$, and g_j is continuous with $g_j(0,0) = 0$.

Following from (4.3), the vector of power flows P in $\mathbf{S}$ is given by:

$$P = F\left(x, x^t\right) + G\left(x, x^t\right) u, \tag{4.4}$$

where $F\left(x, x^t\right) = \left[f_j\left(x_j^{tail}, x_j^{head}\right)\right]$, $G\left(x, x^t\right) = diag\left(\left[g_j\left(x_j^{tail}, x_j^{head}\right)\right]\right)$, and $u = [u_j]$.

4.3.2 Switched Graph-Based Systems

This section extends the above modeling approach to a switched system in which edges internal and external to the graph are rendered active or inactive by switching. To describe this switching, indicator matrices $\Phi = diag\left(\left[\phi_j\right]\right), j \in [1{:}N_e]$ and $\Gamma = diag\left(\left[\gamma_j\right]\right), j \in [1{:}N_s]$ are defined by:

$$\phi_j = \begin{cases} 1 & \text{if } e_j \text{ is active,} \\ 0 & \text{if } e_j \text{ is inactive,} \end{cases} \tag{4.5a}$$

$$\gamma_j = \begin{cases} 1 & \text{if } e_j^s \text{ is active,} \\ 0 & \text{if } e_j^s \text{ is inactive.} \end{cases} \tag{4.5b}$$

The switching signal $\sigma(t)\colon [0,\infty) \to [1{:}N_w]$ selects from among a set of N_w possible indicator matrices $\{\Phi_n, \Gamma_n\}, n \in [1{:}N_w]$. The system corresponding to the n^{th} set of indicator matrices is said to be in the n^{th} mode. This mode is said to be active at times when $\sigma(t) = n$, and inactive at times when $\sigma(t) \neq n$. Note that an equivalent representation could employ a unique switching signal for each of the indicator matrices. However, for simplicity of notation in this chapter, a common switching signal is assigned to both.

Incorporating the indicator matrices into the nominal system of (4.2), the switched system under switching signal $\sigma(t)$ is given by:

$$\mathbf{S}_\sigma\colon \quad C\dot{x} = -\bar{M}\Phi_\sigma P + D\Gamma_\sigma P^s, \tag{4.6}$$

where P is given by (4.4).

For a given Φ_n, $\Phi_n F(x, x^t)$ and $\Phi_n G(x, x^t)$ retain the properties assumed of $F(x, x^t)$ and $G(x, x^t)$ in Assumption 4.2. However, the following additional assumption on the autonomous dynamics of each mode will be necessary to Theorem 4.1. First, note that because $F(0,0) = 0$ by Assumption 4.2, for each Φ_n the autonomous dynamics given by:

$$\begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix} \dot{\bar{x}} = -M\Phi_n F(\bar{x}), \quad \bar{x} = \begin{bmatrix} x \\ x^t \end{bmatrix} \tag{4.7}$$

admit the equilibrium $\bar{x}^* = 0$, where the argument (x, x^t) of F in (4.4) has been replaced with the equivalent argument $\bar{x}$.

Assumption 4.3:
For each Φ_n, the Jacobian matrix:

$$A_{\Phi_n} = -\left.\frac{\partial M\Phi_n F(\bar{x})}{\partial \bar{x}}\right|_{\bar{x}=0} \tag{4.8}$$

has eigenvalues λ_i such that $\mathrm{Re}(\lambda_i(A_{\Phi_n})) < 0$, $\forall i$.

In accordance with Lyapunov's indirect method, Assumption 4.3 ensures that (4.7) is asymptotically stable within a neighborhood of the origin. Alternatively, if $\mathrm{Re}(\lambda_i(A_{\Phi_n})) = 0$

for any i, then a center manifold analysis [110] can be used to check this stability requirement, facilitated by the requirement in Assumption 4.2 that each f_j is twice continuously differentiable.

4.4 Main Results

4.4.1 Passivity Under Switching

The classical notion of passivity is presented in Definition 4.1 below. This is extended to the case of switched systems in Definition 4.2.

Definition 4.1: *Passivity [110]*

Consider a system $\mathbf{H}$ given by $\dot{x} = f(x,u)$ and $y = h(x,u)$. Let $f\colon \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^n$ be locally Lipschitz, $h\colon \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^p$ be continuous, $f(0,0) = 0$, and $h(0,0) = 0$. $\mathbf{H}$ is *passive* if there exists a continuously differentiable positive semidefinite function $V(x)$ such that:

$$u^T y \geq \dot{V} = \frac{\partial V}{\partial x} f(x,u), \quad \forall (x,u). \tag{4.9}$$

$\mathbf{H}$ is locally passive if (4.9) holds in a neighborhood of the origin.

In (4.9), V is called the storage function and $s(u,y) = u^T y$ is called the supply rate. In extending Lyapunov stability to switched systems, it is typical to employ a common Lyapunov function that decreases along the solutions of all modes [26]. Passivity can be extended to switched systems by a similar approach, employing a common storage function whose rate of decrease along the solutions of all modes is upper bounded by the supply rate [128, 133]. Performing this modification to Definition 4.1 leads to the following.

Definition 4.2: *Passivity of Switched Systems*

Consider a switched system $\mathbf{H}_\sigma$ given by $\dot{x} = f_\sigma(x,u)$, $y = h_\sigma(x,u)$, and $\sigma(t)\colon [0,\infty) \to [1{:}N_w]$, where each mode satisfies the requirements in Definition 4.1. $\mathbf{H}_\sigma$ is *passive* if there exists a continuously differentiable positive semidefinite function $V(x)$ such that:

$$u^T y \geq \dot{V} = \frac{\partial V}{\partial x} f_\sigma(x,u), \quad \forall (x,u,\sigma). \tag{4.10}$$

$\mathbf{H}_\sigma$ is locally passive if (4.10) holds in a neighborhood of the origin.

Existence of the common storage function for all modes in Definition 4.2 ensures passivity of the switched system under arbitrary switching [128]. It is next demonstrated that such a common storage function exists for the switched graph-based system $\mathbf{S}_\sigma$.

Theorem 4.1: *Passivity of Switched Graph-Based Systems*
The system $\mathbf{S}_\sigma$ given by (4.6) is locally passive from inputs $\bar{u}$ to outputs $\bar{y}$ with:

$$\bar{u} = \begin{bmatrix} \Gamma_\sigma P^s \\ u \\ -x^t \end{bmatrix}, \quad \bar{y} = \begin{bmatrix} x^s \\ y \\ \Phi_{2,2} P^t \end{bmatrix}. \tag{4.11}$$

Here, $\Phi_{2,2} = diag\left([\phi_{(N_v - N_t + 1)}, ..., \phi_{N_v}]_\sigma\right)$ is the lower right block of Φ_σ indicating the switching of P^t and

$$y = -G(\bar{x})\Phi_\sigma M^T \bar{x}, \tag{4.12}$$

where $\bar{x}$ is as defined in (4.7).

Proof: Consider the storage function $V = \frac{1}{2} x^T C x$, which satisfies the requirements of Definition 4.2. Taking the derivative and substituting in (4.6) gives:

$$\dot{V} = x^T C \dot{x} = -x^T \bar{M} \Phi_\sigma P + x^T D \Gamma_\sigma P^s. \tag{4.13}$$

Noting that $x^s = D^T x$, (4.13) can be written as:

$$\dot{V} = -x^T \bar{M} \Phi_\sigma P + (x^s)^T \Gamma_\sigma P^s. \tag{4.14}$$

Adding and subtracting $(x^t)^T \Phi_{2,2} P^t$ and substituting $\Phi_{2,2} P^t = -\underline{M} \Phi_\sigma P$ into the added term gives:

$$\dot{V} = -x^T \bar{M} \Phi_\sigma P - (x^t)^T \underline{M} \Phi_\sigma P + (x^s)^T \Gamma_\sigma P^s - (x^t)^T \Phi_{2,2} P^t. \tag{4.15}$$

Noting that $x^T\bar{M} + (x^t)^T\underline{M} = \bar{x}^T M$ and substituting in the right side of (4.4) for P gives:

$$\dot{V} = -\bar{x}^T M\Phi_\sigma F(\bar{x}) - \bar{x}^T M\Phi_\sigma G(\bar{x})u + (x^s)^T\Gamma_\sigma P^s - (x^t)^T\Phi_{2,2}P^t. \tag{4.16}$$

Note that the argument (x, x^t) of F and G in (4.4) has been replaced in (4.16) with the equivalent argument $\bar{x}$ as defined in (4.7). Noting also that $G = G^T$ and $\Phi_\sigma = \Phi_\sigma^T$, the second term of (4.16) equals $u^T y$, with y as defined in (4.12). Equation (4.16) can therefore be written as:

$$\dot{V} = -\bar{x}^T M\Phi_\sigma F(\bar{x}) + \bar{u}^T\bar{y}, \tag{4.17}$$

with $\bar{u}$ and $\bar{y}$ as defined in (4.11). By Assumption 4.3 and following similar steps as the proof of Lyapunov's indirect method in Theorem 4.7 of [110], it can be shown that within a neighborhood of the origin there exists a negative semidefinite matrix $\tilde{A}_{\Phi_n}$ for each Φ_n such that:

$$-\bar{x}^T M\Phi_n F(\bar{x}) \leq \bar{x}^T\tilde{A}_{\Phi_n}\bar{x} \leq 0. \tag{4.18}$$

Therefore, $\dot{V} \leq \bar{u}^T\bar{y}$, completing the proof. □

4.4.2 Zero-State Detectability

The following definition is found in [109] for non-switched systems and extended to the case of switched systems in [128].

Definition 4.3: *Zero-State Detectability and Observability [128]*
Consider a switched system $\mathbf{H}_\sigma$ with zero input, given by $\dot{x} = f_\sigma(x, 0),\ y = h_\sigma(x, 0),\ \sigma(t) \in [1{:}N_w]$, and let $Z \subset \mathbb{R}^n$ be the largest positively invariant set contained in $\bigcap_{i=1}^{N_w}\{x \in \mathbb{R}^n | y_i = 0\}$. Then $\mathbf{H}_\sigma$ is zero-state detectable (ZSD) if $x = 0$ is asymptotically stable conditionally to Z. If $Z = 0$, then $\mathbf{H}_\sigma$ is zero-state observable (ZSO).

To guarantee that the system of switched graph-based models is ZSD, the following assumption is required.

Assumption 4.4:

For each Φ_n, the intersection of the zero sets of $\phi_j g_j, \forall j \in [1{:}N_e]$ is the origin, i.e.,

$$\bigcap_j \mathcal{N}\left(\phi_j g_j\left(x_j^{tail}, x_j^{head}\right)\right) = \{0\}. \tag{4.19}$$

Lemma 4.1:

The system $\mathbf{S}_\sigma$ given by (4.6) is ZSO with respect to the outputs $\bar{y}$ as defined in (4.11). Therefore it is also ZSD.

Proof: ZSO can be interpreted as a requirement that when $\bar{u} = 0$, $\bar{x} = 0$ if $\bar{y} = 0$. This can in turn be ensured by proving that $y \neq 0$ if $x \neq 0$ and $\bar{u} = 0$. From the definition of M in (2.3), the vector y in (4.12) is given by:

$$y = [y_j] = -[\phi_j g_j\left(x_j^{tail}, x_j^{head}\right)\left(x_j^{tail} - x_j^{head}\right)]. \tag{4.20}$$

By Assumption 4.4, there exists j such that $\phi_j g_j\left(x_j^{tail}, x_j^{head}\right) \neq 0$ if $\left(x_j^{tail}, x_j^{head}\right) \neq 0$. Furthermore, $x_j^{tail} - x_j^{head} = 0\ \forall j$ only at the equilibrium $\bar{x}^* = 0$. This is true because $\bar{u} = 0$ implies that $x^t = 0$, and x^t is a component of $\bar{x}$. Therefore, $y \neq 0$ if $x \neq 0$ and $\bar{u} = 0$ as required, completing the proof. $\square$

4.4.3 Passivity of Interconnected Subsystems

In this section it is demonstrated that passivity is preserved under the interconnection of multiple graph-based subsystems that are each passive under switching. The k^{th} subsystem, subject to its switching signal $\sigma_k(t)$, is denoted as $\mathbf{S}_{k,\sigma_k}$. This subsystem's associated signals are similarly identified by the subscript k, for example, $P_{k,1}^s$ is the first power flow source of subsystem $\mathbf{S}_{k,\sigma_k}$.

Figure 4.1 shows a notional interconnection of subsystem $\mathbf{S}_{k,\sigma_k}$ with upstream subsystem $\mathbf{S}_{k-1,\sigma_{k-1}}$ and downstream subsystem $\mathbf{S}_{k+1,\sigma_{k+1}}$. Note that $(\Phi_{2,2}P^t)_{k-1} = (\Gamma_\sigma P^s)_k$ and $x_{k-1}^t = x_k^s$, where source states x_k^s are the states internal to $\mathbf{S}_{k,\sigma_k}$ incident to a source edge of the system's graph. This interconnection can be seen as a negative feedback connection of

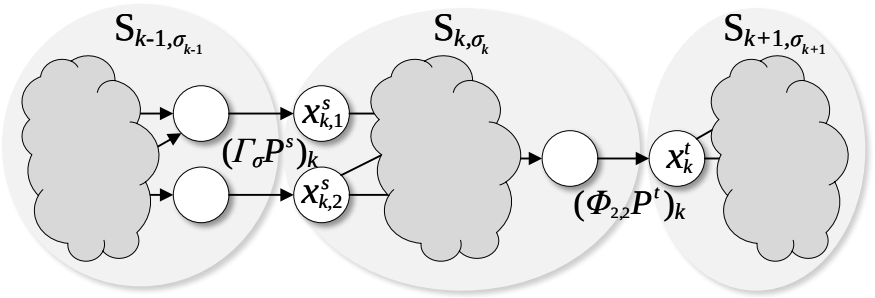

Figure 4.1: Notional interconnection of three subsystems [134].

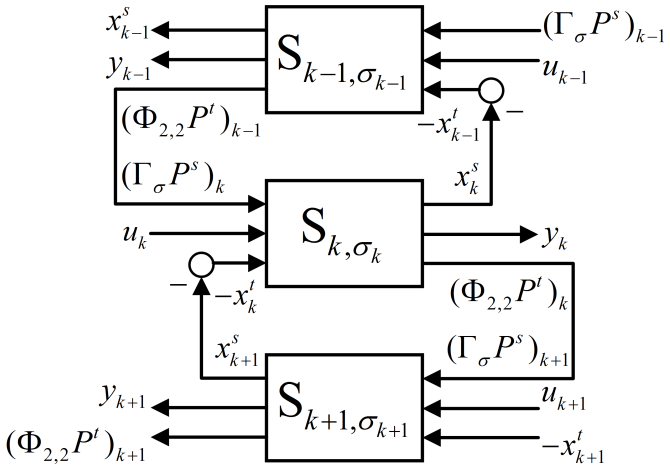

Figure 4.2: Block diagram representation of subsystems from Fig. 4.1 [134].

subsystems, as shown in Fig. 4.2, where the inputs and outputs of the k^{th} subsystem are $\bar{u}_k$ and $\bar{y}_k$, respectively, as defined in (4.11).

The overall system given by an interconnection of K subsystems is denoted as $\tilde{\mathbf{S}}_\sigma$, with $\tilde{u} = [u_k]$ and $\tilde{y} = [y_k]$, $k \in [1{:}K]$. Furthermore, $\tilde{P}^s$, $\tilde{P}^t$, $\tilde{x}^s$, and $\tilde{x}^t$ denote the source power flows, sink power flows, source states, and sink states, respectively, of the overall system. Similarly, $\tilde{\Gamma}_\sigma$ and $\tilde{\Phi}_{2,2}$ indicate switching of the source and sink power flows of the overall system. For example, in the interconnection of Fig. 4.2, $\tilde{\Gamma}_\sigma \tilde{P}^s = (\Gamma_\sigma P^t)_{k-1}$, $\tilde{\Phi}_{2,2}\tilde{P}^t = (\Phi_{2,2}P^t)_{k+1}$, $\tilde{x}^s = x^s_{k-1}$, and $\tilde{x}^t = x^t_{k+1}$.

Note that a cascaded structure of subsystems is shown in Figs. 4.1 and 4.2 to clearly demonstrate the nature of the interconnections, but the forthcoming results of this section are not limited to a cascaded structure, and instead apply to any subsystem interconnection structure satisfying the following assumption.

Definition 4.4:

A v_i, v_j*–path* on a graph $\mathbf{G}$ is a sequence of edges (regardless of orientation) connecting two distinct vertices $v_i, v_j \in v$, not including any sink vertices as intermediate vertices. $\mathbf{G}$ is *connected* if it has a v_i, v_j*–path* for each $v_i, v_j \in v$.

Assumption 4.5:

The overall graph given by an interconnection of subsystem graphs is connected for all allowable modes of all subsystems.

The above assumption can be seen as a constraint on both the nominal topology of interconnections of subsystems and on the allowable switching modes. The role of Assumption 4.5 is to ensure that there always exists at least one edge along which power can flow between neighboring subsystems, forming a negative feedback connection. A classical result in passivity theory is that a system formed by negative feedback connection of two passive subsystems is itself passive [110]. This leads to Theorem 4.2 below, proved in [125] by induction on the number of subsystems. While relaxations of Assumption 4.5 fall outside the scope of this dissertation, future work could explore a *weak* connectivity assumption [135], requiring connectivity to hold only over the union of active edges over sufficiently long time intervals. Alternately, when the graph is not connected, it is possible to analyze the decoupled components independently.

Theorem 4.2: *Passivity Under Interconnection [125]*

Consider a set of K subsystems, each given by (4.6) and locally passive in accordance with Theorem 4.1. Let $\tilde{\mathbf{S}}_\sigma$ be a system given by the interconnection of these subsystems satisfying Assumption 4.5. Then $\tilde{\mathbf{S}}_\sigma$ is locally passive from inputs $\bar{u}$ to outputs $\bar{y}$ with:

$$\bar{u} = \begin{bmatrix} \tilde{\Gamma}_\sigma \tilde{P}^s \\ \tilde{u} \\ -\tilde{x}^t \end{bmatrix}, \quad \bar{y} = \begin{bmatrix} \tilde{x}^s \\ \tilde{y} \\ \tilde{\Phi}_{2,2} \tilde{P}^t \end{bmatrix}. \tag{4.21}$$

where $\tilde{u} = [u_k]$ and $\tilde{y} = [y_k]$, $k \in [1{:}K]$.

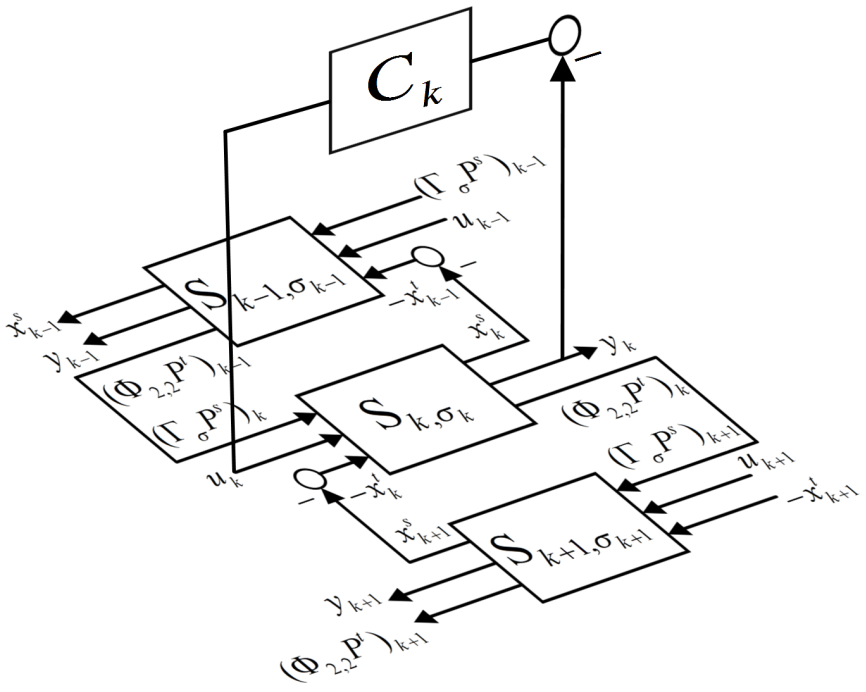

Figure 4.3: Block diagram showing negative feedback connection between subsystem $\mathbf{S}_{k,\sigma_k}$ and its controller $\mathbf{C}_k$. The additional controllers $\mathbf{C}_{k-1}$ and $\mathbf{C}_{k+1}$ for subsystems $\mathbf{S}_{k-1,\sigma_{k-1}}$ and $\mathbf{S}_{k+1,\sigma_{k+1}}$ are not depicted.

4.4.4 Stabilization by Decentralized MPC

This section presents the design of decentralized MPC controllers that stabilize the closed-loop dynamics of an interconnection of switched graph-based models. Switching occurrences are assumed to be known to these controllers at the switching instant. It is also assumed that the discretization of continuous dynamics performed to implement MPC preserves passivity and stability. While not guaranteed in all cases [136], this is a common assumption in the analysis of nonlinear MPC that can be relaxed by explicit consideration of the sampled-data nature of implementation [137, 138].

The controller for the k^{th} subsystem, denoted as $\mathbf{C}_k$, is inserted between the signals u_k and y_k of that subsystem, as shown in Fig. 4.3. Just as with the interconnection among subsystems, the connection of each subsystem to its controller can be viewed as a negative feedback connection. Therefore, passivity of the closed-loop system is achieved by ensuring that each controller is passive.

Each controller $\mathbf{C}_k$ solves the following nonlinear MPC optimization problem:

$$\min_{\boldsymbol{u_k}, \boldsymbol{s_k}} \quad \int_0^T \ell\left(x_k(\tau), u_k(\tau), r_k(\tau), s_k(\tau)\right) d\tau \tag{4.22a}$$

$$\text{subject to} \quad C_k \dot{x}_k = -\bar{M}_k \Phi_{k,\sigma_k} P_k + D_k \Gamma_{k,\sigma_k} P_k^s, \tag{4.22b}$$

$$P_k = F_k\left(\bar{x}_k\right) + G_k\left(\bar{x}_k\right) u_k, \tag{4.22c}$$

$$y_k = -G_k\left(\bar{x}_k\right) \Phi_{k,\sigma_k} M_k^T \bar{x}_k \tag{4.22d}$$

$$x_k^{min} - s_k(\tau) \leq x_k(\tau) \leq x_k^{max} + s_k(\tau), \quad s_k(\tau) \geq 0, \tag{4.22e}$$

$$u_k(\tau) \in \mathbb{U}_k, \tag{4.22f}$$

$$\dot{z}_k = u_k^T y_k, \quad z_k(\tau) \leq \beta_k, \tag{4.22g}$$

$$\tau \in [0, T], \tag{4.22h}$$

$$x_k(0) = x_{k,init}, \quad z_k(0) = z_{k,init}, \tag{4.22i}$$

where the stage cost $\ell(\cdot)$ is a positive definite function, $r_k(\tau)$ is a set of references to be tracked, $s_k(\tau)$ is a slack variable ensuring feasibility of the state constraints as defined in (4.22e), and $\mathbb{U}_k = \{u_k | u_{k,j}^{min} \leq u_{k,j} \leq u_{k,j}^{max}, \forall j\}$ with $u_{k,j}^{min} < 0 < u_{k,j}^{max}, \forall j$.

Equation (4.22g) enforces a passivity-based constraint, where $z_k \in \mathbb{R}$ represents the accumulation of passivity and is used to reduce the conservatism associated with more conventional passivity-based MPC, such as [120]. Similar to [118], when $u_k^T y_k < 0$, the excess passivity is stored by decreasing z_k. This stored passivity gives the controller the flexibility to operate the system non-passively for a finite amount of time in minimizing the integrated stage cost, until $z_k = \beta_k$, where β_k is a predetermined constant for each subsystem. When $z_k = \beta_k$, the controller must operate the system passively. The following theorem proves that this passivity constraint guarantees stability of the closed-loop system.

Theorem 4.3: *Closed-Loop Stability*

Let $\tilde{\mathbf{S}}_\sigma$ be a switched system formed by the interconnection of switched subsystems, each given by (4.6), satisfying Assumptions 4.1-4.5. If each subsystem is paired with an MPC controller of the form (4.22), then the overall closed-loop system is stable.

Proof: It is first shown that the optimization problem (4.22) is recursively feasible. With feasibility established, closed-loop stability can then be proved. *Feasibility:* As proved for non-switched systems in [109] and extended to the case of switched systems in [128], a subsystem $\mathbf{S}_{k,\sigma_k}$ meeting the assumptions in this chapter is asymptotically stable about the origin under any control law of the form $u_k = -K_k(y_k)$, where $K_k(y_k)$ is a continuous function such that $K_k(0) = 0$ and $y_k^T K_k(y_k) > 0,\ \forall y_k \neq 0$. A requirement for this to hold is ZSD, which has already been established in Lemma 4.1. Because $\mathbb{U}_k$ contains a neighborhood of the origin, for any y_k there exists $\alpha_k > 0$ such that $u_k = -\alpha_k y_k \in \mathbb{U}_k$. Under this α_k, $\dot{z}_k = u_k^T y_k = -\alpha y_k^T y_k \leq 0$. Therefore, if $z_k(0) \leq \beta_k$ and $x_k(0)$ is sufficiently close to the origin, the optimization problem (4.22) is recursively feasible.

Stability: As in [118], let $V_{C_k}(z_k) = -z_k + \beta_k$ be a storage function for the controller $\mathbf{C}_k$. Then the closed-loop subsystem formed by $\mathbf{S}_k$ and $\mathbf{C}_k$ has a storage function given by $V_k + V_{C_k}$ with:

$$\begin{aligned} \dot{V}_k + \dot{V}_{C_k} &\leq \bar{u}_k^T \bar{y}_k - \dot{z}_k, \\ &\leq \bar{u}_k^T \bar{y}_k - u_k^T y_k, \\ &\leq \begin{bmatrix} (\Gamma_\sigma P^s)_k^T & -(x_k^t)^T \end{bmatrix} \begin{bmatrix} x_k^s \\ (\Phi_{(2,2)} P^t)_k \end{bmatrix}. \end{aligned} \tag{4.23}$$

The above proves that the closed-loop subsystem is passive with respect to its interconnections to neighboring subsystems, while Theorem 4.2 proves that these interconnections preserve passivity. Lemma 4.1 establishes that each subsystem is ZSD, and a passive ZSD system is stable [128]. Therefore, the closed-loop system is stable. □

This result differs from many approaches in the literature that require satisfaction of a *global* condition on the structure and gain of couplings among subsystems to guarantee stability [57, 116, 123]. By comparison, the result here is achieved using *local* analysis of each subsystem and a *local* passivity-based constraint within each controller. This is enabled by the structure with which the graph-based models interconnect.

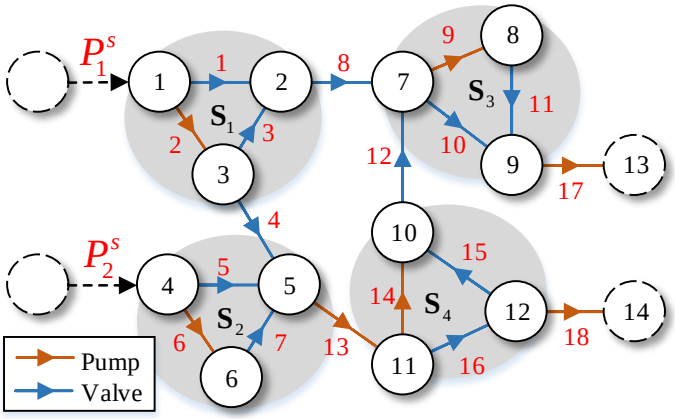

Figure 4.4: Graph with four subsystems for the example fluid tank system [134].

4.5 Simulation Example

This section presents a simulation example to demonstrate the efficacy of the passivity-based control approach. Figure 4.4 shows the graph of a fluid tank system, which has the same structure as the example system from [139]. Each vertex v_i corresponds to a fluid tank and has a state x_i representing the deviation in height of fluid in the tank, in meters, from an equilibrium value. In the graph-based model of this hydraulic system, the dynamics of (4.1) correspond to the conservation of fluid mass for each tank:

$$\rho A_i \dot{x}_i = \sum_{\{j | e_j \in e_i^{head}\}} \dot{m}_j - \sum_{\{j | e_j \in e_i^{tail}\}} \dot{m}_j, \tag{4.24}$$

where $\rho = 1000\ kg/m^3$ is the density of the fluid and $A_i = \pi d_i^2/4$ is the cross-sectional area of the tank, with diameter $d_i = 0.1\ m$, $\forall i$. In this example, the states are initialized using a pseudorandom uniform distribution centered at the origin.

The power flows from (4.3) correspond to mass flow rates between the tanks, which are each assumed to be controlled by either a pump or a variable aperture valve. For flows controlled by pumps, the mass flow rate along edge e_j is given by:

$$\dot{m}_j = Disp \cdot \omega_j - k_{leak} \left(x_j^{head} - x_j^{tail} \right), \tag{4.25}$$

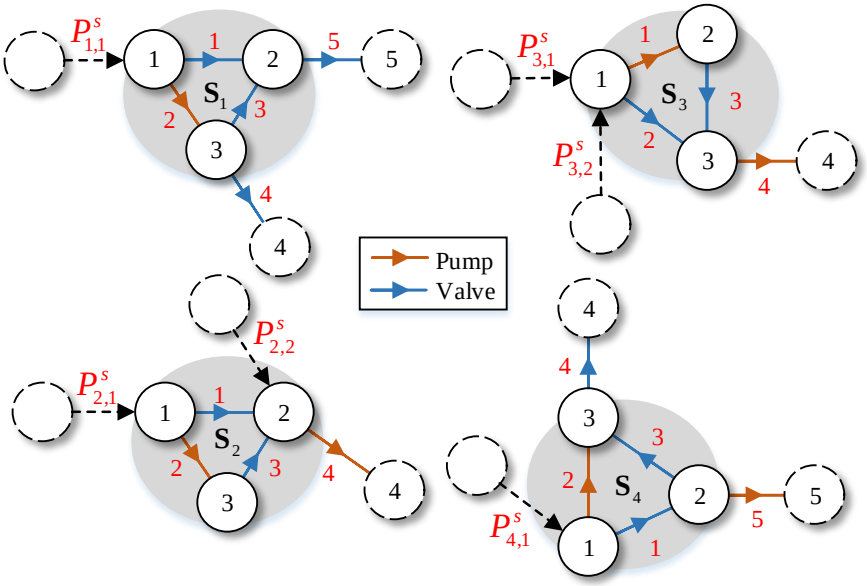

Figure 4.5: Decomposition of the example system graph into four subsystem graphs, used to develop four decentralized MPC controllers [134].

where $Disp = 1\ kg/rev$ is the displacement of the pump, ω_j is the variable pump speed in revolutions per second, and $k_{leak} = 0.005\ kg/(m \cdot s)$ is the leakage coefficient. For flows controlled by valves, the mass flow rate along edge e_j is given by:

$$\dot{m}_j = C_D^{max} \left(x_j^{tail} - x_j^{head} \right) a_j, \tag{4.26}$$

where $C_D^{max} = 0.05\ kg/(m \cdot s \cdot \%)$ is the maximum discharge coefficient and a_j is the variable aperture by which the valve is controlled, in % open. These equations for the mass flow rate along each edge can be shown to satisfy the conditions in Assumptions 4.1-4.2.

The overall system is composed of four dynamically coupled subsystems, as indicated in Fig. 4.4. The four subsystem graphs are shown in Fig. 4.5. Under the assumption that some edges may be switched on or off, the corresponding switched subsystem dynamics of (4.6) can be derived from the graphs of Fig. 4.5 and (4.24)-(4.26). Additionally, the passivity outputs y_k for each subsystem can be determined from (4.12). For a pump edge $e_{k,j}$ in subsystem $\mathbf{S}_k$, the corresponding passivity output is given by:

$$y_{k,j} = -\phi_{k,j} Disp \cdot \left(x_{k,j}^{tail} - x_{k,j}^{head} \right). \tag{4.27}$$

Similarly, for a valve edge, the corresponding passivity output is given by:

$$y_{k,j} = -\phi_{k,j} C_D^{max} \left(x_{k,j}^{tail} - x_{k,j}^{head} \right)^2 . \tag{4.28}$$

It can be shown that each nominal (non-switched) subsystem is stable in open loop. For simplicity of exposition in this example, only one edge is switched. This is edge e_4 of subsystem $\mathbf{S}_3$ in Fig. 4.5, which can equivalently be referred to as edge e_{17} of the overall system in Fig. 4.4. Subsystem $\mathbf{S}_3$ can be shown to be stable in open loop with this edge removed. Therefore, Assumption 4.3 is satisfied. Assumptions 4.4 and 4.5 can be easily verified to hold as well. In this example, both the source flows and sink states of the overall system are set to zero, and the control task is to stabilize the system from its initial condition to equilibrium (i.e., the origin).

For each subsystem $\mathbf{S}_k$, an MPC-based controller $\mathbf{C}_k$ is designed based on the optimization problem from (4.22) with:

$$\ell \left(x_k(\tau), u_k(\tau), r_k(\tau) \right) = \|x_k - r_k\|_2^2 + 0.01 \|u_k\|_2^2. \tag{4.29}$$

In this example, $T = 2$ s, $\mathbb{U}_k = \{u_k | u_{k,j} \in \mathbb{R}, \forall j\}$, and $\beta_k = 0$. The optimization problem is discretized and solved with an update rate of 1 Hz, which in this case is sufficiently fast to render the model error due to discretization negligible. YALMIP [104] and IPOPT [140] are used to formulate and solve the optimization problem for each controller.

The local nature of the guarantee in Theorem 4.3 allows the subsystem controllers to ensure stability while operating in a purely decentralized fashion, or to improve performance they can be coordinated within a distributed or hierarchical control framework by receiving references communicated from each other or additional controllers. In this example, a hierarchical framework is employed to demonstrate the ability of the subsystem controllers to form a supervisory stability assurance layer at the bottom of a hierarchical framework, similar to that in [141]. The role of this layer is to ensure that decisions made above in the hierarchy, including switching commands, do not destabilize the plant.

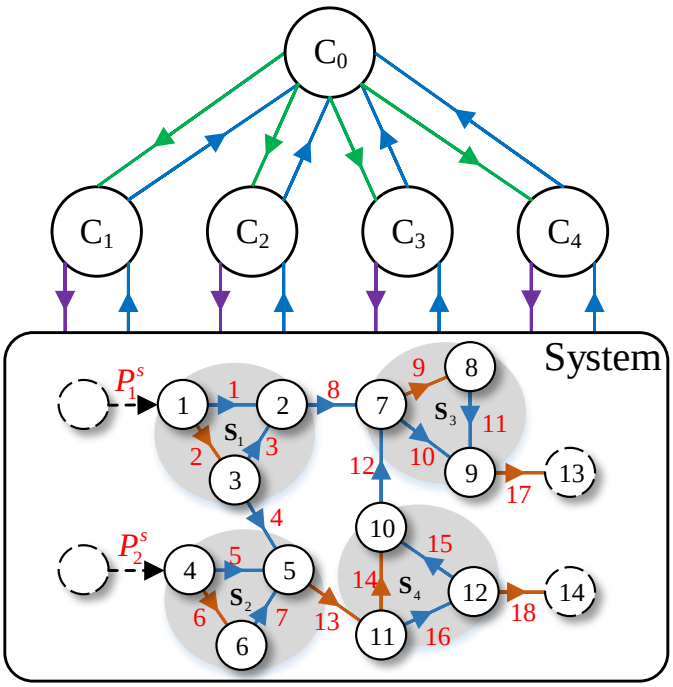

Figure 4.6: Control hierarchy for the example system, where $\mathbf{C}_{1-4}$ are passivity-based decentralized MPC controllers and $\mathbf{C}_0$ is a centralized reference generator [134].

The two-level hierarchy used for this example is shown in Fig. 4.6. Upper-level controller $\mathbf{C}_0$ sends references r_k to be tracked by each subsystem, with the goal of coordinating their actions. This controller was designed using a linearization of the system about the initial state $x(0)$ and has the form $\mathbf{C}_0 \colon r = -Zx$. This reference generator is used for both the nominal (non-switched) system and the system under switching.

To demonstrate the stabilizing effect of the decentralized passivity-based constraints, edge e_{17} is periodically switched on and off as shown in Fig. 4.7. Without the passivity-based constraint (4.22g) in each controller, this switching has a destabilizing effect on the system, despite the individual stability of each mode. Figure 4.8 shows the closed-loop state trajectory for vertex v_2 of subsystem $\mathbf{S}_1$ both with and without the passivity-based constraint under non-switched and switched conditions. As shown by the red trace of Fig. 4.8a, without the passivity-based constraint and with no switching, the system is stabilized to the origin, although this is not guaranteed in general. However, when switching occurs in subsystem $\mathbf{S}_3$, this results in highly oscillatory behavior in subsystem $\mathbf{S}_1$ that goes unstable just before 900 seconds, as shown by the blue trace of Fig. 4.8a. This instability is captured by the

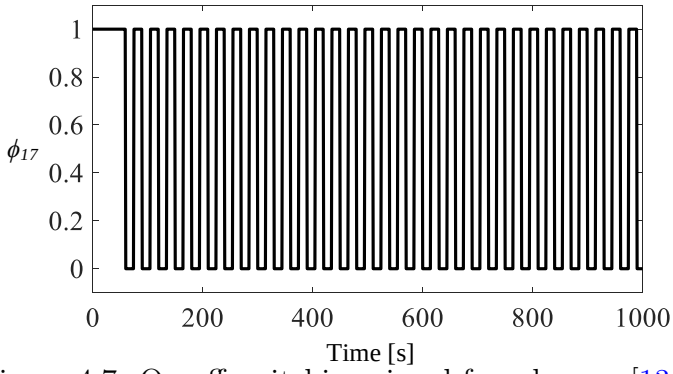

Figure 4.7: On-off switching signal for edge e_{17} [134].

increasing stored passivity z_1 for subsystem $\mathbf{S}_1$, as shown in Fig. 4.9a. By comparison, when the passivity-based constraint is included in the subsystem control formulations, the system is stabilized to the origin regardless of the switching, as shown in Fig. 4.8b.

In the passivity-constrained cases, $\beta_k = 0$, which necessitates that $z_k \leq 0$ as shown in Fig. 4.9b. However, for MPC without the passivity constraint and without switching, the stable system response is not instantaneously passive, with $\dot{z}_1 > 0$ for the initial transient, as shown by the red trace of Fig. 4.9a. This highlights the potential conservatism associated with a passivity-based stability approach. This conservatism can be reduced by increasing the value of β_k and allowing the system to act non-passivity temporarily. In practice, β_k should be designed based on the application-specific tradeoff between the benefit of aggressive control and the cost of potentially, yet temporarily, following an unstable trajectory.

4.6 Chapter Summary

This chapter presents a decentralized approach to certifying closed-loop passivity for a class of nonlinear graph-based dynamic models in which paths of power flow are switched on and off. A common storage function is shown to guarantee passivity of graph-based models under arbitrary switching. Passivity is then shown to be preserved under the interconnection of multiple graph-based systems, allowing for the formation of passive "systems of systems." Decentralized Model Predictive Controllers paired with each system can then be formulated with a local passivity-preserving constraint to ensure closed-loop stability. This

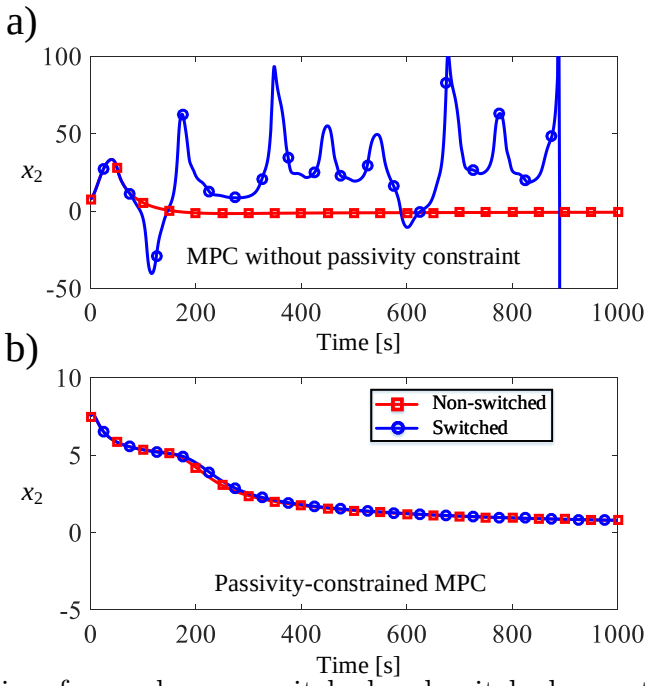

Figure 4.8: Trajectories of x_2 under non-switched and switched operation for a) MPC without the passivity constraint (4.22g) and b) passivity-constrained MPC [134].

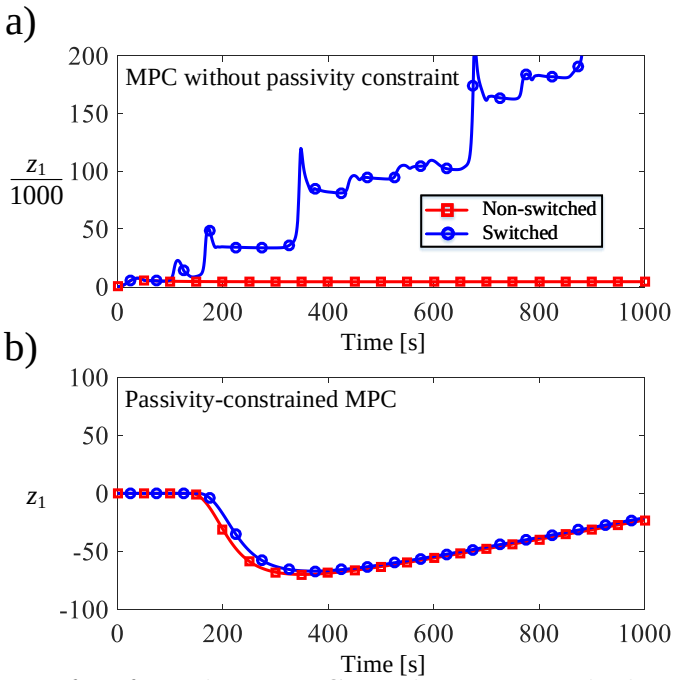

Figure 4.9: Trajectories of z_1 for subsystem $\mathbf{S}_1$ under non-switched and switched operation for a) MPC without the passivity constraint (4.22g) and b) passivity-constrained MPC [134].

allows complex energy systems to be stabilized with decentralized, distributed, or hierarchical control architectures where centralized control may not be possible due to the inherent computational complexity or communication bandwidth limitations. An simulation example demonstrates the efficacy of the proposed approach on a fluid tank system controlled by a hierarchical framework.

Chapter 5

Cooperativity-Based Hierarchical MPC for State-Constrained Switched Systems

5.1 Motivation

While the passivity-based approach in Chapter 4 can ensure the stability of a switched energy system about an equilibrium, no claim is made about the ability of the closed-loop system to satisfy hard constraints on states. Instead, the slack variable in (4.22e) imposes soft constraints on states, with constraint violations penalized in the cost function to avoid their occurrence. Softened state constraints are also used in the experimental case study of Chapter 3, where state constraint violations occurred under the hierarchical control framework despite a tightening of constraints by 2°C in MPC formulations to account for errors in the disturbance preview and model. While the constraint violations in this case were small relative to those under the baseline controller and within acceptable margins for many applications, some applications may require more strict enforcement. For example, laser diode arrays must be thermally managed within strict temperature limits to prevent excessive wavelength shift and spectral broadening [142]. This motivates the need for theoretical methods applicable to energy systems that guarantee the satisfaction of hard constraints on states. These methods must be robust in the sense that guarantees hold under the presence of bounded uncertainties in disturbances, and must also support systems governed by both continuous and switched behavior.

This chapter presents a two-level switched hierarchical control framework that guarantees satisfaction of state and input constraints for a class of graph-based dynamic systems. Conditions are provided under which switched energy systems belong to this class, in which each mode is a cooperative system. These conditions relate to physical phenomena such as satisfaction of conservation equations, and therefore are inherently met by many of the

energy systems within the scope of this dissertation. The overall objective of the control design is to track a reference trajectory for the rate of energy transfer to a system, subject to bounded uncertainty in the exogenous disturbances associated with its ability to dissipate energy to sinks. The hierarchical framework manages slow system dynamics over a long time horizon in the upper level while simultaneously leveraging fast dynamics in the lower level to improve performance and reject disturbances. The applicability of this approach is demonstrated in both simulation and experimental application using a fluid-thermal system with dynamic behavior representative of aircraft fuel thermal management systems. Note that the results in this chapter extend the previous work in [143] by permitting multiple controllers within the lower level of the hierarchical framework and including experimental demonstration.

5.2 Background

Monotone systems preserve partial orderings on their state trajectories, as explored in detail for autonomous dynamics in [144] and controlled systems in [145]. They appear in applications including molecular biology [146], traffic flow models [147], data transmission systems [85], building thermal dynamics [148, 149], and fluid-thermal management systems, as demonstrated in this chapter. Cooperative systems are a subclass of monotone systems in which the ordering on the states is specified as elementwise inequalities. In this chapter, as in [85], cooperativity is ensured by conditions on the structure of a graph representing a dynamic system. However, inputs enter additively in the models of [85], whereas in this chapter inputs affect the coupling between dynamic states through multiplicative nonlinear functions, allowing a different set of physical systems and energy domains to be represented. This chapter also considers systems that may be affected by coupling to multiple environmental disturbances or exogenous neighboring systems, whereas the networks in [85] are assumed to be coupled to a single external node.

Monotone systems exhibit several properties that can be exploited in control design. This includes the aforementioned ordering of state trajectories, which allows state trajectories to be bounded by their solutions under maximal and minimal disturbances. In addition,

robust positively invariant (RPI) sets can be generated from equilibrium values [150]. In [85, 148], decentralized linear feedback control laws with saturation are applied to achieve asymptotic stability and robustness of cooperative systems. In [149], a finite-state abstraction is performed and a receding-horizon controller that can be computed offline is applied. The separability of Lyapunov functions for monotone systems has been an area of focus in their study, including in [151–154]. Analysis and control design tools for the related class of positive systems, the linear counterpart of cooperative systems, are also relevant. These often seek to reveal or leverage computational advantages of control synthesis for positive systems as compared to linear systems in general. For example, [155] proves that stabilizing distributed state feedback controllers for positive systems can be synthesized using linear programming. Motivated by these results in distributed stability, this chapter demonstrates that the task of operating a cooperative system within hard constraints can be distributed as well.

A robust hierarchical control approach for graph-based systems has been proposed in [63, 156], where feedback integralization ensures that states do not violate constraints between updates of the controllers at each level of the hierarchy. This approach is is applicable to hierarchical frameworks with any number of levels, but requires that known disturbances only change at the rate of the slowest controller of the hierarchy and does not consider switching. The results in this chapter address switched systems and allow disturbances to change at the rate of the fastest controller of the hierarchy. While the focus here is on two-level hierarchies, extension to an arbitrary number of levels is a topic of ongoing study.

The remainder of this chapter is organized as follows. Section 5.3 describes the class of switched graph-based systems used in this chapter. Section 5.4 proves that, under appropriate assumptions, each mode is a stable cooperative system. This section also formulates a two-level switched hierarchical MPC framework that ensures satisfaction of state constraints subject to bounded uncertainties in disturbances. Section 5.5 demonstrates the applicability and efficacy of the approach using a simulation example. The demonstration is extended to experimental application in Section 5.6 using the fluid-thermal testbed described in Section 2.4.1. Section 5.7 provides a concluding summary of the chapter.

5.3 Class of Systems

5.3.1 Graph-Based System

Consistent with the graph-based model formulation introduced in Chapter 2, the dynamic associated with each vertex of a graph $\mathbf{G}$ is governed by the conservation equation:

$$C_i \dot{x}_i = \sum_{\{j | e_j \in e_i^{head}\}} P_j - \sum_{\{j | e_j \in e_i^{tail}\}} P_j, \tag{5.1}$$

where $C_i > 0$ is the storage capacitance of the vertex and each P_j is the power flow associated with an edge incident to the vertex.

Also as in Chapter 2, the dynamics of the complete system $\mathbf{S}$ can be written as:

$$\mathbf{S}: \quad C\dot{x} = -\bar{M}P + DP^s, \tag{5.2}$$

where $C = diag\left([C_i]\right)$, $\bar{M}$ is a partition of the incidence matrix of the graph, P^s is the vector of source power flows, and D is a mapping from the source edges to the vertices to which they are incident.

In this Chapter, stability analysis is facilitated by the assumption of a state-affine form for the edge power flow equations, as stated next. This represents a key point of departure from the passivity-based analysis in Chapter 4, in which the edge power flow equations are restricted to an input-affine form.

Assumption 5.1:
For all power flows of $\mathbf{S}$, the power flow P_j along edge e_j is given by:

$$P_j = f_j^{tail}\left(u_j\right) x_j^{tail} - f_j^{head}\left(u_j\right) x_j^{head}, \tag{5.3}$$

where x_j^{tail} and x_j^{head} are the states associated with the vertices to which e_j is incident, and u_j is an associated actuator input to the scalar functions $f_j^{head}, f_j^{tail} \colon \mathbb{R} \to \mathbb{R}_+$.

Remark 5.1:

Assumption 5.1 restricts the class of power flows to those affine in the state and restricts the signs of the coefficients on these states. For many energy systems, this assumption is inherently satisfied. For example, thermal power flow due to advection is modeled in (2.14) as $P(t) = \dot{m}(t)c_p T^{tail}(t)$, where c_p is the specific heat capacitance of the fluid, $\dot{m}$ is the fluid mass flow rate input, and T^{tail} is the upstream fluid temperature state. In this case, $f^{tail}(\dot{m}(t)) = \dot{m}(t)c_p$ and $f^{head}(\dot{m}(t)) = 0$. This is similar to the 'h-type' flows used to model compartmental flow systems in [85]. In (2.17), the heat transfer coefficient associated with convective heat transfer is a function of both mass flow rate and temperature. However, assumption of constant heat transfer coefficients in the case study of Chapter 3 did not significantly degrade the model accuracy, suggesting that this may not be necessary in all cases. If the heat transfer coefficient is assumed to be a function of mass flow rate only, the thermal power flow due to convective heat transfer in a heat exchanger from (2.16) can be written as $P(t) = h(\dot{m}(t))\, A_s \left(T^{tail}(t) - T^{head}(t)\right)$, where $h(\dot{m}(t))$ is the heat transfer coefficient, A_s is the convective surface area, and $T^{tail}(t) - T^{head}(t)$ is the temperature gradient. In this case, $f^{tail}(\dot{m}(t)) = f^{head}(\dot{m}(t)) = h(\dot{m}(t))\, A_s$. This is similar to the 'g-type' flows in [85].

Assumption 5.2:

For all vertices i of graph $\mathbf{G}$ and all t,

$$\sum_{\{j|e_j \in e_i^{tail}\}} \left(f_j^{tail}(u_j(t)) - f_j^{head}(u_j(t))\right) - \sum_{\{j|e_j \in e_i^{head}\}} \left(f_j^{tail}(u_j(t)) - f_j^{head}(u_j(t))\right) = 0. \quad (5.4)$$

Remark 5.2:

As with Assumption 5.1, Assumption 5.2 has physical relevance. For example, in modeling the temperature of a fluid junction or split, Assumption 5.2 requires that a steady-state form of conservation of mass holds among the associated advective power flows. This is exemplified in Fig. 5.1, where for constant c_p, Assumption 5.1 is equivalent to the requirement that $\dot{m}_1(t) = \dot{m}_2(t) + \dot{m}_3(t),\ \forall t$;

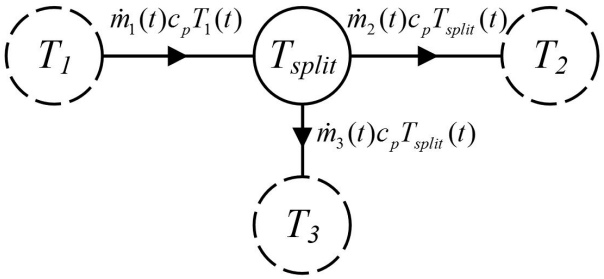

Figure 5.1: Adjective heat transfer at a fluid split, where steady-state conservation of mass necessitates that $\dot{m}_1(t) = \dot{m}_2(t) + \dot{m}_3(t)$.

5.3.2 State Space Representation

From the state-affine equation for each power flow in (5.3), it is clear that the vector of power flows $P = [P_j]$ in $\mathbf{S}$ can be written as:

$$P = F(u) \begin{bmatrix} x \\ x^t \end{bmatrix}, \tag{5.5}$$

where $F(u) = [f(u_j)_{ij}] \in \mathbb{R}^{N_e \times N_v}$ such that:

$$f(u_j)_{ji} = \begin{cases} f_j^{tail}(u_j) & \text{if } v_i \text{ is the tail of } e_j, \\ -f_j^{head}(u_j) & \text{if } v_i \text{ is the head of } e_j, \\ 0 & \text{else.} \end{cases} \tag{5.6}$$

Modulo the transposition of its indices, (5.6) bears strong resemblance to the weighted incidence matrix of (2.33) that results from linearization of the thermal graph-based model in Section 2.3.5. However, as a result of Assumption 5.1, the linear representation in this case is exact rather than approximate. Note also that $F(u)$ has at least the same sparsity as M^T, possibly with additional entries equal to zero, which will be of relevance to Lemma 5.1.

A partition similar to that used for M in (2.4) can be applied to $F(u)$, yielding:

$$F(u) = \begin{bmatrix} \bar{F}(u) & \underline{F}(u) \end{bmatrix} \text{ with } \bar{F}(u) \in \mathbb{R}^{N_e \times (N_v - N_t)}. \tag{5.7}$$

Substituting (5.5) and (5.7) into (5.2), the dynamics of system **S** can be written as:

$$\mathbf{S}: \quad \dot{x} = A(u)x + BP^s + E(u)x^t, \tag{5.8}$$

where:

$$A(u) = -C^{-1}\bar{M}\bar{F}(u), \tag{5.9a}$$

$$B = \quad C^{-1}D, \tag{5.9b}$$

$$E(u) = -C^{-1}\bar{M}\underline{F}(u). \tag{5.9c}$$

5.3.3 Switched Linear System Representation

As exemplified by the case study in Chapter 3, many energy systems are governed by discrete actuators, such as solenoid valves and electrical switches. To capture switching with a representation that supports predictive control under limited computational resources, the following assumption transforms the nonlinear system (5.8) into a switched linear system. When both continuously variable and switched actuators are present, it is well known that nonlinear systems can be approximated arbitrarily well by switched linear systems using a sufficient number of modes [157, 158].

Assumption 5.3:
The vector of system inputs u is restricted to a finite set of values $\mathbb{U} = \{u_n\}$, $u_n \in \mathbb{R}^{N_e}$, $n \in [1{:}N_m]$ such that $u(t) \in \mathbb{U}, \ \forall t$.

The system under input u_n is said to be in the n^{th} mode. The active mode at each time is selected by a switching signal $\sigma(t)\colon [0, \infty) \to [1{:}N_m]$, with $u(t) = u_{\sigma(t)}$. Under Assumption 5.3, (5.8) can be written as:

$$\mathbf{S}_\sigma: \quad \dot{x} = A_{\sigma(t)}x + BP^s + E_{\sigma(t)}x^t, \tag{5.10}$$

where:

$$A_{\sigma(t)} = A(u_{\sigma(t)}),$$
$$E_{\sigma(t)} = E(u_{\sigma(t)}). \tag{5.11}$$

5.4 Main Results

5.4.1 Stability Conditions

This section proves that A_σ is Hurwitz in each mode under appropriate assumptions on the system graph structure.

Definition 5.1:

A v_i, v_j*–path* on a graph $\mathbf{G}$ is a sequence of edges (regardless of orientation) connecting two distinct vertices $v_i, v_j \in v$. An *admissible path* is a path such that $f_k^{tail}(u_k) > 0$ for each edge k such that the path enters the tail of the edge, and $f_k^{head}(u_k) > 0$ for each edge k such that the path enters the head of the edge. $\mathbf{G}$ is *strongly connected* if it has a v_i, v_j*–admissible path* for each $v_i, v_j \in v \backslash v^t$. $\mathbf{G}$ is *externally connected* if at least one vertex in v^t is the tail of at least one edge k such that $f_k^{tail}(u_k) > 0$, or is the head of at least one edge k such that $f_k^{head}(u_k) > 0$.

Assumption 5.4:

For each mode of system $\mathbf{S}_\sigma$, the associated graph $\mathbf{G}_\sigma$ is strongly connected and externally connected.

Remark 5.3:

Assumption 5.4 performs a similar role as the analogous conditions in [85]. Specifically, this ensures that no subset of the states are decoupled from the others and ensures that there is coupling to at least one of the exogenous sink states. It is possible for the power flow along individual edges to necessarily equal zero in some modes (i.e., for rows of $F(u_\sigma)$ to have all zero entries in some modes), provided that the overall graph structure satisfies Assumption 5.4. For example, the power flow along an edge could equal zero in a mode where

a fluid valve is completely closed or an electrical switch is opened, removing the mechanism for power to flow directly between the incident energy storage elements.

Definition 5.2: *Diagonal Dominance [150]*

A square matrix $H = [h_{ij}]$ is *weakly row-diagonally dominant* if

$$|h_{ii}| \geq \sum_{j \neq i} |h_{ij}|, \ \forall i. \tag{5.12}$$

H is *strictly row-diagonally dominant for at least one row* if it is weakly row-diagonally dominant and

$$\exists i \text{ s.t. } |h_{ii}| > \sum_{j \neq i} |h_{ij}|. \tag{5.13}$$

Lemma 5.1:

The following hold for each mode n of the system matrices in (5.10):

(a) A_n is weakly row-diagonally dominant.

(b) A_n is strictly row-diagonally dominant for at least one row.

(c) A_n has only strictly negative diagonal entries and only non-negative off-diagonal entries.

(d) B and E_n have only non-negative entries.

Proof: (a) The left side of (5.4) includes all nonzero terms in the i^{th} row of $MF(u)$. Therefore by Assumption 5.2, $MF(u)$ has a row sum of zero for all rows at all times. From (2.4) and (5.7):

$$MF(u) = \begin{bmatrix} \bar{M}\bar{F}(u) & \bar{M}\underline{F}(u) \\ \underline{M}\bar{F}(u) & \underline{M}\underline{F}(u) \end{bmatrix}, \tag{5.14}$$

with $\bar{M}\bar{F}(u) \in \mathbb{R}^{(N_v - N_t)\times(N_v - N_t)}$ and $\bar{M}\underline{F}(u) \in \mathbb{R}^{(N_v - N_t)\times N_t}$. In forming $A(u)$ in (5.9a), only the $\bar{M}\bar{F}(u)$ block of $MF(u)$ is used. Deleting the N_t columns of the block to the right can only reduce the sum of the absolute value of the off-diagonal entries in each row, while the diagonal entries remain unchanged. Therefore, $\bar{M}\bar{F}(u)$ is weakly row-diagonally dominant. This property is preserved under multiplication by the diagonal matrix $-C^{-1}$ in (5.9a).

(b) Under the external connectivity of Assumption 5.4, at least one entry of $\bar{M}\underline{F}(u)$ must be nonzero. $MF(u)$ has a row sum of zero for all its rows as established above, therefore deleting the columns of $\bar{M}\underline{F}(u)$ necessitates that $\bar{M}\bar{F}(u)$ is strictly row-diagonally dominant for at least one row. This property is again preserved under multiplication by the diagonal matrix $-C^{-1}$ in (5.9a).

(c) As noted following (5.6), $F(u)$ has at least the same sparsity as M^T. The entries also have matching signs, in the sense that $m_{ij} = 1 \iff f_{ji}(u_j) \geq 0$, $m_{ij} = -1 \iff f_{ji}(u) \leq 0$, and $m_{ij} = 0 \iff f_{ji}(u) = 0$. By matrix multiplication,

$$(MF(u))_{ij} = \begin{bmatrix} M_{i1} & M_{i2} & \cdots & M_{iN_e} \end{bmatrix} \cdot \begin{bmatrix} F_{1j}(u) & F_{2j}(u) & \cdots & F_{N_e j}(u) \end{bmatrix}. \tag{5.15}$$

For diagonal entries of $MF(u)$, where $i = j$, any negative entry on the left of the dot product is multiplied by a non-positive entry on the right, while any positive entry on the left of the dot product is multiplied by a non-negative entry on the right. Furthermore, by Assumption 5.4 at least one of the entries on the right of the dot product must be nonzero, otherwise vertex j cannot be a member of any admissible path. As a result, the diagonal entries in $MF(u)$ must be strictly positive. For the off-diagonal entries, note that each column of M has one positive entry, one negative entry, and zeros elsewhere. Similarly, each row of $F(u)$ has at most one positive and one negative entry. For off-diagonal entries of $MF(u)$, where $i \neq j$ in the dot product above, any positive entry on the left of the dot product must be multiplied by a non-positive entry on the right, while any negative entry on the left of the dot product must be multiplied by a non-negative entry on the right. As a result, the off-diagonal entries of $MF(u)$ must be non-positive. Multiplication of the $\bar{M}\bar{F}(u)$ block of $MF(u)$ by the diagonal matrix $-C^{-1}$ to form $A(u)$ in (5.9a) reverses the signs of all entries.

(d) From (2.5), D clearly has non-negative entries, and this is preserved under multiplication by C^{-1} in (5.9b). From (5.9c), $E(u)$ consists of a subset of the off-diagonal entries of $MF(u)$, which are non-positive as stated above. Multiplication by the diagonal matrix $-C^{-1}$ in (5.9c) reverses these signs. $\square$

Definition 5.3: *Irreducible Matrix [150]*

A square matrix H is *irreducible* if there does not exist a variable permutation such that the corresponding matrix $\tilde{H}$ has the form:

$$\tilde{H} = \begin{bmatrix} \tilde{H}_{11} & \tilde{H}_{12} \\ 0 & \tilde{H}_{22} \end{bmatrix}, \tag{5.16}$$

where $\tilde{H}_{11}$ and $\tilde{H}_{22}$ are square matrices.

Lemma 5.2:

A_n is irreducible for each n.

Proof: If A_n were not irreducible for some n, then for that mode there would not exist an admissible path from any vertices associated with the states in the lower right block of its upper triangular permutation to any of the remaining vertices, contradicting the assumption of strong connectivity for all modes [85]. □

Theorem 5.1:

Each A_n is Hurwitz, i.e, $\text{Re}(\lambda_j(A_n)) < 0,\ \forall j \in [1{:}(N_v - N_t)]$ and $\forall n \in [1{:}N_m]$.

Proof: By Theorem 4.60 in [150], if a square matrix (1) is weakly row-diagonally dominant, (2) is strictly row-diagonally dominant for at least one row, (3) has only strictly negative diagonal entries, and (4) is irreducible, then it is Hurwitz. The first three of these conditions are established by Lemma 5.1, and the fourth is established by Lemma 5.2. Note that this is *not* sufficient to guarantee stability of the switched system. In general, switching between stable modes can result in unstable dynamics [26]. □

5.4.2 Cooperative System

This section proves that each mode of system $\mathbf{S}_\sigma$ is cooperative and specifies several of the properties of cooperative systems that will be leveraged for control design.

Definition 5.4: *Metzler Matrix and Metzler System [145, 150]*

A matrix $H = [h_{ij}]$ is a *Metzler matrix* if $h_{ij} \geq 0,\ \forall i \neq j$. A linear dynamic system is a

Metzler system if the state matrix is a Metzler matrix and the remaining matrices (input, disturbance, etc.) have only non-negative entries. Metzler systems are often called *positive systems* when the state matrices are continuous or constant, named in reference to the fact that the positive orthant is forward invariant for positive input signals [145, 159].

Definition 5.5: *Cooperative System [145]*
Consider a dynamic system given by $\dot{x}(t) = g(x(t), w(t))$ subject to the continuity assumptions in [145], and let $\varphi(t, x_0, w(t))$ be its solution at time t with initial state x_0 at time $t_0 \leq t$ and disturbance $w(t)$. The system is *cooperative* if

$$x_a \leq x_b \text{ and } w_a(t) \leq w_b(t) \implies \varphi(t, x_a, w_a(t)) \leq \varphi(t, x_b, w_b(t)), \ \forall t \geq t_0, \tag{5.17}$$

where the inequalities are applied elementwise at each t.

Theorem 5.2:
Each mode of system $\mathbf{S}_\sigma$ is cooperative.

Proof: From (c) and (d) of Lemma 5.1, each mode of system $\mathbf{S}_\sigma$ is a Metzler system. As proved in [145], an LTI Metzler system is cooperative. □

Definition 5.6:
Consider a dynamic system given by $\dot{x}(t) = g(x(t), w(t))$ and let $\varphi(t, x_0, w(t))$ be its solution at time t with initial state x_0 at time $t_0 \leq t$ and disturbance $w(t) \in \mathbb{W}$. A set $\mathbb{S}$ is a *robust positively invariant (RPI) set* if

$$x_0 \in \mathbb{S}, \ w(t) \in \mathbb{W} \implies x(t) \in \mathbb{S}, \ \forall t \geq t_0. \tag{5.18}$$

The following theorem is a direct consequence of the definition of a cooperative system.

Theorem 5.3: *RPI Sets for Cooperative Systems [150]*
Let $\dot{x}(t) = g(x(t), w(t))$ be a cooperative system subject to the continuity assumptions in [145]. Let $\underline{x}_e$ and $\bar{x}_e$ be equilibrium points corresponding to disturbances $\underline{w}_e$ and $\bar{w}_e$ (i.e. $0 = g(\underline{x}_e, \underline{w}_e) = g(\bar{x}_e, \bar{w}_e)$) such that $\underline{w}_e \leq \bar{w}_e$ elementwise. Then the set $\mathbb{S} = \{x \colon \underline{x}_e \leq$

$x \leq \bar{x}_e\}$ is RPI $\forall w(t)$ s.t. $\underline{w}_e \leq w(t) \leq \bar{w}_e$. In this chapter, such sets are referred to as *equilibrium-based* RPI sets.

Remark 5.4:

A key feature of the equilibrium-based RPI sets above is that they are axis-aligned boxes in the state space, termed *box invariant* sets in [160]. Box invariant sets can be defined by $2N$ inequality constraints, where N is the state dimension. By comparison, the complexity of representation of general polyhedral RPI sets is not fixed by the state dimension [150]. In addition, because the equilibrium points are necessarily reachable, this approach yields the *minimal* robust invariant interval with respect to the disturbances [161]. Minimal RPI sets are desirable in many robust MPC techniques, but in general cannot be computed exactly, and are often instead approximated using iterative methods [162]. In this chapter, solving for two equilibrium points provides a computationally convenient method for calculating RPI sets online within control formulations, improving the scalability of the proposed approach.

Theorem 5.4:

Consider a Metzler system $\dot{x} = Ax + Ew$ composed of N subsystems. The system dynamics can be written as a function of the subsystem state vectors $x_i \in \mathbb{R}^{M_i},\ i \in [1{:}N]$ as:

$$\underbrace{\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \vdots \\ \dot{x}_N \end{bmatrix}}_{\dot{x}} = \underbrace{\begin{bmatrix} A_{11} & A_{12} & \cdots & A_{1N} \\ A_{21} & A_{22} & & \\ \vdots & & \ddots & \\ A_{N1} & & & A_{NN} \end{bmatrix}}_{A} \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{bmatrix}}_{x} + \underbrace{\begin{bmatrix} E_{11} & E_{12} & \cdots & E_{1N} \\ E_{21} & E_{22} & & \\ \vdots & & \ddots & \\ E_{N1} & & & E_{NN} \end{bmatrix}}_{E} \underbrace{\begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_N \end{bmatrix}}_{w}. \tag{5.19}$$

Note that the dynamics for the i^{th} subsystem can be written as:

$$\dot{x}_i = A_{ii}x_i + \sum_{j \neq i} A_{ij}x_j + \sum_j E_{ij}w_j. \tag{5.20}$$

Let $\mathbb{S}$ be the RPI set defined by equilibrium points $\underline{x}_e,\ \bar{x}_e$ with respect to disturbance bounds $\underline{w}_e,\ \bar{w}_e$, and let $\mathbb{S}_i$ be the projection of $\mathbb{S}$ onto the subspace corresponding to the

i^{th} subsystem. Then $\mathbb{S}_i$ is an RPI set for the i^{th} subsystem with respect to disturbances $\underline{w}_e \leq w(t) \leq \bar{w}_e$ and $\underline{x}_{e,j} \leq x_j(t) \leq \bar{x}_{e,j},\ j \neq i$.

Proof: In (5.20), each A_{ii} must be a Metzler matrix because it is a sub-matrix from the block diagonal of Metzler matrix A. Similarly, each $A_{ij},\ j \neq i$ is an off-diagonal block of A, and therefore must have non-negative entries. Each E_{ij} must also have non-negative entries, as a block of the disturbance matrix of a Metzler system. Therefore, each subsystem is a Metzler system, and as stated in the proof of Theorem 5.2, an LTI Metzler system is cooperative. □

Remark 5.5:

Theorem 5.4 is leveraged in Level 2 of the hierarchical control formulation in this chapter to distribute the task of operating the overall system within an RPI set among multiple controllers, without requiring them to directly exchange information. Note that, if $A_{ij} = 0$ for some j in (5.20), then x_i will still remain within $\mathbb{S}_i$ even if the inequality $\underline{x}_{e,j} \leq x_j(t) \leq \bar{x}_{e,j}$ does *not* hold for that j but does hold over all j for which $A_{ij} \neq 0$. In the next section, similar sparsity in the coupling between subsystems will be leveraged to allow controllers of a hierarchy the flexibility to govern a subset of states beyond their local RPI sets, while still ensuring robustness of the overall system.

The open-loop system analysis in this chapter concludes with with a final observation, leading to a relaxation of Assumption 5.4. Rather than requiring the entire system graph to be strongly and externally connected, it is sufficient for the graph to consist of multiple disconnected components, provided that each component is itself strongly and externally connected. Clearly, the analyses above can be applied individually to each component to reach the same conclusions about the properties of the full system. Therefore, the relaxation of Assumption 5.4 is as follows:

Assumption 5.5:

For each mode of system $\mathbf{S}_\sigma$, the associated graph $\mathbf{G}_\sigma$ is composed exclusively of components that are both strongly connected and externally connected.

5.4.3 Hierarchical Control

The overall design of the hierarchical control framework proposed in this chapter draws from the goals in Chapter 3 of achieving high capability, efficiency, and safety of vehicle energy systems. Capability is represented as the objective of achieving a desired source power flow $P^{s,des}$ associated with the desired operation of mission- and flight-critical hardware. The controller has authority over the source power flows P^s, and must ensure that a minimum value for the source power flows $\underline{P}^s$ can always be applied, similar to the flight-critical loads in Chapter 3. When the desired source power flow exceeds the flight-critical loads, the controller can choose to perform load shedding of the excess as necessary to ensure satisfaction of constraints, at the cost of reducing the capability of the system to achieve its mission.

The controller is also subject to an upper bound on the source power flows $\bar{P}^s$. In some applications, the objective associated with a subset of the source power flows could be to minimize their values rather than to achieve a desired value. For example, an objective of thermal management for an automotive electric vehicle in a cold climate could be to minimize the power output to a resistive heater aiding in thermal management of the battery. This would be weighed against additional objectives associated with the temperature-dependent efficiency and life cycle of the battery. Source power flow minimization objectives can be included by setting $P^{s,des} = 0$ for the corresponding signals.

The goal of efficiency is captured by minimizing costs of operation associated with mode selection. As discussed in Section 5.3.3, different modes of the switched systems in this chapter correspond to dynamics of graph-based models under different values of the edge inputs u, which are typically associated with continuously-variable and discrete actuators of the system. For example, in the thermal modeling of Sections 2.3.2 and 3.3.1, these inputs correspond to the mass flow rates achieved by pumps, valves, and fans. The efficiency of pumps and fans is proportional to their mass flow rate [1], while the electrical energy consumption of solenoid valves depends on whether they are open or closed. The specific solenoid valve listed in Table 2.2(c) is closed when unpowered, and requires power supply to lift its plunger.

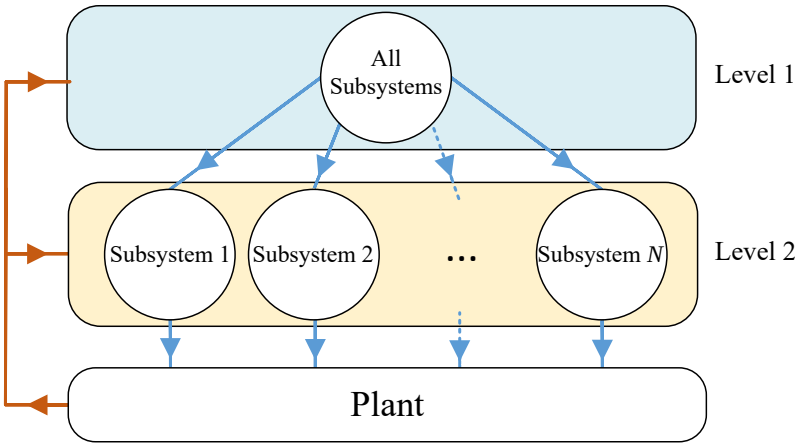

Figure 5.2: Two-level hierarchical control framework used in this chapter.

Safety is captured by enforcing hard box constraints on system states, $\underline{x} \leq x_k \leq \bar{x}$, where the subscript k denotes the time steps of a discretized dynamic system. This is performed subject to uncertainty in the sink states, which are treated as exogenous disturbances also bounded by box constraints, $\underline{x}^t \leq x_k^t \leq \bar{x}^t$.

The proposed hierarchical control framework consists of two levels, as shown in Fig. 5.2. The Level 1 controller operates on a relatively slow timescale and long prediction horizon, overseeing the operation of all subsystems. This level is responsible for governing slow timescale behavior of the system, taking proactive action as necessary to achieve long-term coordination among the subsystems. Equilibrium-based RPI sets are also generated at this level, which are associated with conservative bounds on the capability achievable within state constraints. In Level 2 of the hierarchy, multiple controllers corresponding to specific subsystems each operate on a relatively fast timescale and short prediction horizon, leveraging the fast dynamics of the plant to improve capability and compensate for disturbances. These controllers have the authority to drive states beyond the RPI sets generated in the level above, provided that they can ensure recovery to within the sets by the next update of the Level 1 controller. This significantly reduces the conservatism of the framework as compared to what would occur if only the Level 1 controller were present.

5.4.3.1 Discretization

To implement MPC, a discrete representation of the continuous dynamics for each mode of system $\mathbf{S}_\sigma$ is required. This discretization is performed by assuming a zero-order hold for P^s, x^t, and σ. The resulting discrete system is given by:

$$\mathbf{S}_\sigma^d: \quad x_{k+1} = A_{\sigma_k}^d x_k + B_{\sigma_k}^d P_k^s + E_{\sigma_k}^d x_k^t, \tag{5.21}$$

where for each i,

$$\begin{aligned} A_i^d &= \exp(A_i \tau^d), \\ B_i^d &= A_i^{-1}(A_i^d - I)B, \\ E_i^d &= A_i^{-1}(A_i^d - I)E_i, \end{aligned} \tag{5.22}$$

and τ^d is the sample time in seconds. Note that each A_i is invertible by Theorem 5.1. When the zero-order hold assumption is correct, the discrete states exactly match the continuous dynamics at the sample times. Therefore, the discrete system adheres to the discrete equivalents of the cooperativity and stability properties established previously.

The guarantee of closed-loop state constraint satisfaction in this chapter is made with respect to the discretized dynamics at the update rate of the fastest controller of the hierarchy. In many applications, this update rate can be chosen fast enough such that the error from the continuous evolution of the states is negligible. However, future work will seek to extend the guarantee to encompass the continuous dynamics as well.

5.4.3.2 Level 1 Controller

The Level 1 controller updates every τ^{d_1} seconds with a prediction horizon of N_1 steps. The subscript k denotes steps of this prediction horizon, while the subscript h will denote steps of the faster-timescale prediction horizon of the Level 2 controllers, described in the next section. The Level 1 controller receives the current states of the plant x_k and a profile of the desired source power flows over its horizon, denoted as $P_{k+j|k}^{s,des}$, $j \in [0{:}N_1 - 1]$. This profile does not have to be exact, but instead represents a "best guess" of the upcoming desired

source power flows, sampled at the update rate of the Level 1 controller. For example, in thermal management of a connected or autonomous vehicle as discussed in Section 1.1.3.3, the desired source power flows could be the thermal energy generation associated with anticipated operation of electrical and/or mechanical components. Without loss of generality, it is also assumed that predicted (but possibly inaccurate) preview of the sink states is available for each step in the horizon, denoted as $\hat{x}^t_{k+j|k},\ j \in [0{:}N_1 - 1]$, falling within the box constraints on the sink states $\underline{x}^t, \bar{x}^t$. For example, in a vehicle thermal management system, the sink state preview information could be the ambient temperature predicted using weather data. In absence of the availability of sink state preview, the midpoints of the corresponding box constraints can be used instead.

Theorem 5.3 states that an RPI set can be found for a cooperative system by solving for two equilibrium points. As noted in Remark 5.4, this provides a computationally convenient method for calculating RPI sets online in the Level 1 control formulation. Such sets are found for the first M_1 steps of the horizon, where $1 \leq M_1 \leq N_1$. The set corresponding to the k^{th} step is denoted as: $\mathbb{S}_k = \{x \mid \underline{y}_k \leq x \leq \bar{y}_k\}$, where the equilibrium points $\underline{y}_k, \bar{y}_k$ defining the bounds of the set must fall within the state constraints, i.e., $\underline{x} \leq \underline{y}_k \leq \bar{y}_k \leq \bar{x}$. While the bounds on the sink states are fixed in computing these sets, there is freedom in the mode of operation σ_k and source power flows $\underline{P}^{s,y}_k, \bar{P}^{s,y}_k$ used to compute the equilibrium points at each time step, subject to bounds $\underline{P}^s \leq \underline{P}^{s,y}_k \leq \bar{P}^{s,y}_k \leq \bar{P}^s$.

As shown in Figure 5.3, the first RPI set of the horizon $\mathbb{S}_{k|k}$ is defined by equilibrium points $\underline{y}_{k|k}, \bar{y}_{k|k}$. This set must contain the corresponding state $\hat{x}_{k|k}$, and by definition of an RPI set (Definition 5.6) will contain the next predicted state $\hat{x}_{k+1|k}$, provided that the sink states obey their constraints and the applied source power flows fall between $\underline{P}^{s,y}_k$ and $\bar{P}^{s,y}_k$. The second RPI set $\mathbb{S}_{k+1|k}$ is defined by equilibrium points $\underline{y}_{k+1|k}, \bar{y}_{k+1|k}$, and so on.

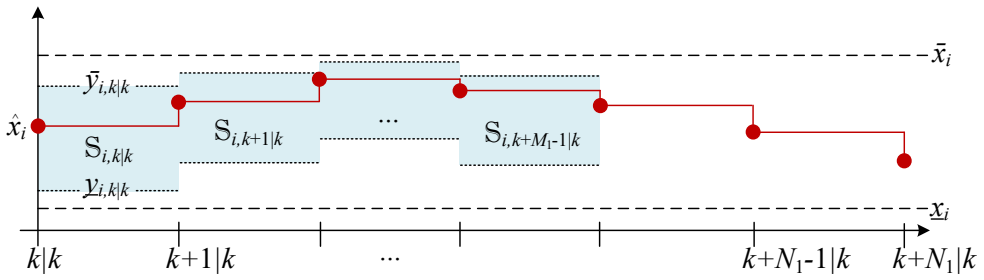

Figure 5.3: Sample predicted trajectory for a state x_i of the Level 1 controller, showing the state constraints and equilibrium-based RPI set bounds.

The Level 1 MPC controller is formulated using the following MIQP:

$$J^*(x_k) = \min_{\boldsymbol{P_k^s}, \boldsymbol{\sigma_k}} \sum_{j=0}^{N_1-1} \ell_1\left(P^{s,des}_{k+j|k}, P^s_{k+j|k}, m_{k+j|k}\right), \tag{5.23a}$$

subject to, $\forall j \in [0{:}N_1-1]$,

$$\hat{x}_{k+j+1|k} = A^{d_1}_{\sigma_{k+j|k}} \hat{x}_{k+j|k} + B^{d_1}_{\sigma_{k+j|k}} P^s_{k+j|k} + E^{d_1}_{\sigma_{k+j|k}} \hat{x}^t_{k+j|k}, \tag{5.23b}$$

$$m_{i,k+j|k} = \begin{cases} 1 & \text{if } \sigma_{k+j|k} = i, \\ 0 & \text{else,} \end{cases} \tag{5.23c}$$

$\forall j \in [0{:}M_1-1]$,

$$\underline{y}_{k+j|k} = (I - A^{d_1}_{\sigma_{k+j|k}})^{-1}(B^{d_1}_{\sigma_{k+j|k}} \underline{P}^{s,y}_{k+j|k} + E^{d_1}_{\sigma_{k+j|k}} \underline{x}^t), \tag{5.23d}$$

$$\bar{y}_{k+j|k} = (I - A^{d_1}_{\sigma_{k+j|k}})^{-1}(B^{d_1}_{\sigma_{k+j|k}} \bar{P}^{s,y}_{k+j|k} + E^{d_1}_{\sigma_{k+j|k}} \bar{x}^t), \tag{5.23e}$$

$$\underline{x} \le \underline{y}_{k+j|k} \le \hat{x}_{k+j|k} \le \bar{y}_{k+j|k} \le \bar{x}, \tag{5.23f}$$

$$\underline{P}^s \le \underline{P}^{s,y}_{k+j|k} \le P^s_{k+j|k} \le \bar{P}^{s,y}_{k+j|k} \le \bar{P}^s, \tag{5.23g}$$

$\forall j \in [M_1{:}N_1-1]$,

$$\underline{x} \le \hat{x}_{k+j+1|k} \le \bar{x}, \tag{5.23h}$$

$$\underline{P}^s \le P^s_{k+j|k} \le \bar{P}^s, \tag{5.23i}$$

$$\hat{x}_{k|k} = x_k. \tag{5.23j}$$

Equation (5.23a) defines the objective function of the controller, for which the stage cost $\ell_1(\cdot)$ is a quadratic function. This will typically include a term penalizing errors in the

tracking of $P^{s,des}$ by P^s, given by $\|P^{s,des}_{k+j|k} - P^s_{k+j|k}\|^2_{\Lambda_P}$. The third argument of $\ell_1(\cdot)$ allows costs associated with mode selection to be applied, where (5.23c) defines $m = [m_i]$ for each j as a vector of binary variables, with unity in the row corresponding to the index of the active mode and zeros elsewhere.

Equation (5.23b) defines the predicted system dynamics at the update rate of the Level 1 controller, subject to the nominal source power flow P^s and predicted sink state $\hat{x}^t$. Equations (5.23d)-(5.23e) solve for the equilibrium points used to generate RPI sets. Note that existence of the inverse in these equations follows from the Schur stability of each $A_i^{d_1}$ [155]. Equation (5.23f) constrains the predicted state to be within the equilibrium points, which must in turn be within the state constraints. Similarly, (5.23g) constrains the nominal source power flows to be within the equilibrium source power flows, which in turn must be within the source power flow constraints. Equations (5.23h) and (5.23i) constrain the predicted states and source power flows of steps for which equilibrium-based RPI sets are not computed. Lastly, Equation (3.12j) sets the initial state of the horizon equal to the current state of the system.

As an intermediate step in proving recursive feasibility and constraint enforcement of the proposed two-level hierarchical control framework, the following Lemma establishes these properties for the Level 1 controller in direct feedback connection with the plant (i.e., in absence of Level 2 in Fig. 5.2).

Lemma 5.3: *Recursive Feasibility of Level 1 Controller*
Consider the feedback connection of the Level 1 controller (5.23) to the discrete system (5.21), both updating every τ^{d_1} seconds. Between step k and step $k+1$, the applied source power flows $P^{s,app}(t)$ are bounded as $\underline{P}^{s,y}_k \leq P^{s,app}(t) \leq \bar{P}^{s,y}_k$, and the sink state disturbances $x^t(t)$ obey their assumed bounds, $\underline{x}^t \leq x^t(t) \leq \bar{x}^t$. If the controller has a feasible solution upon initialization, then it is recursively feasible. Specifically, this implies that the state trajectories always satisfy their their constraints.

Proof: Feasibility upon initialization of the controller at step $k = 0$ requires that there exists an equilibrium-based RPI set $\mathbb{S}_{0|0}$ containing the initial states. If this is the case, by Definition 5.6 the system can remain within an RPI set for all future steps of the horizon

and therefore has a feasible solution for all future iterations of the controller. Note that this does *not* require that the sink states x^t or applied source power flows $P^{s,app}$ remain constant between updates, as equilibrium-based RPI sets hold regardless of whether a continuous or discrete representation of the system model is used because the equilibrium points do not change with the representation. □

Remark 5.6:

Many MPC implementations, especially those for non-switched systems, require that the states enter an invariant set by the end of the prediction horizon. By comparison, the Level 1 control formulation presented here could potentially be conservative and limit the domain of the controller. However, it will be shown that the Level 2 controllers can reduce the conservatism of the overall framework, operating outside the RPI sets by applying different source power flow than that determined by the Level 1 controller. In addition, the RPI sets used by the Level 1 control formulation are computed online, and can therefore be sized to facilitate minimization of the objective function by selecting the mode of operation and/or adjusting $\underline{P}^{s,y}$ and $\bar{P}^{s,y}$ to allow P^s to better track $P^{s,des}$ over the horizon. This is enabled by the computational benefits of using equilibrium-based RPI sets as compared to general RPI sets, discussed in Remark 5.5, and differs from many robust MPC approaches in which RPI sets are computed a priori and do not change during closed-loop operation.

5.4.3.3 Level 2 Controllers

The Level 2 controllers update every τ^{d_2} seconds with a prediction horizon of $N_2 = \tau^{d_1}/\tau^{d_2}$ steps, where it is assumed that τ^{d_1} and τ^{d_2} are designed such that N_2 is an integer. The time horizon of the Level 2 controllers is therefore equal to the update interval of the Level 1 controller. At instances when both the Level 1 and Level 2 controllers are due to update, it is assumed that Level 1 updates first and passes its decision variables to Level 2. Known delays can be built into the control formulations to account for computation times, as done in the case study of Chapter 3. However, doing so falls beyond the scope of the theoretical results in this chapter.

Similar to the Level 1 controller, each Level 2 controller receives a profile of the desired source power flows over its horizon, denoted as $P^{des}_{h+j|h},\ j \in [0{:}N_2 - 1]$. However, in this case the profile is provided at the fast update rate and short horizon of the Level 2 controller, and therefore represents more refined short-term desired behavior of the system rather than the downsampled long-term estimate sent to Level 1.

System Decomposition:

In forming Level 2 of the hierarchical framework, the graph of the full system is first decomposed into sub-graphs corresponding to subsystems, and each subsystem is paired with its own controller. As discussed in Section 1.1.3.4, controllers within the same level of the hierarchical framework do not directly exchange information with each other. This greatly reduces the communication requirements of the hierarchy and precludes the use of the iterative or cascaded methods from distributed control approaches that can be computationally expensive or inhibit parallelization. While it falls beyond the scope of this dissertation to examine the design considerations involved in selection of a specific system decomposition, note that an agglomerative hierarchical clustering approach for decomposing graph-based models has been proposed in [92] and subsequently demonstrated in application to hierarchical control of a simulated automotive electric vehicle in [66], as discussed in Section 2.5.1.

Subsystem Coupling and Sink State Bounds:

The sink states of each sub-graph represent either coupling to exogenous disturbances of the overall system or coupling to states of adjacent subsystems along edges that were cut to decompose the system graph into sub-graphs. For example, one possible decomposition of the notional graph example from Fig. 2.1 is shown in Fig. 5.4. Sub-graph $\mathbf{G}_2$ has four sink vertices. Two of these are also sink vertices of the system graph $\mathbf{G}$, denoted as $v^t_{2,1}, v^t_{2,2}$ in $\mathbf{G}_2$ and corresponding to v^t_1, v^t_2 of $\mathbf{G}$. The remaining two sinks of $\mathbf{G}_2$ represent the coupling to sub-graph $\mathbf{G}_1$. Denoted as $v^t_{2,3}, v^t_{2,4}$ in $\mathbf{G}_2$, these correspond to v_2, v_3 of $\mathbf{G}$, or equivalently, $v_{1,2}, v_{1,3}$ of $\mathbf{G}_1$. Similarly, $\mathbf{G}_1$ has two sink vertices, both of which represent its coupling to $\mathbf{G}_2$.

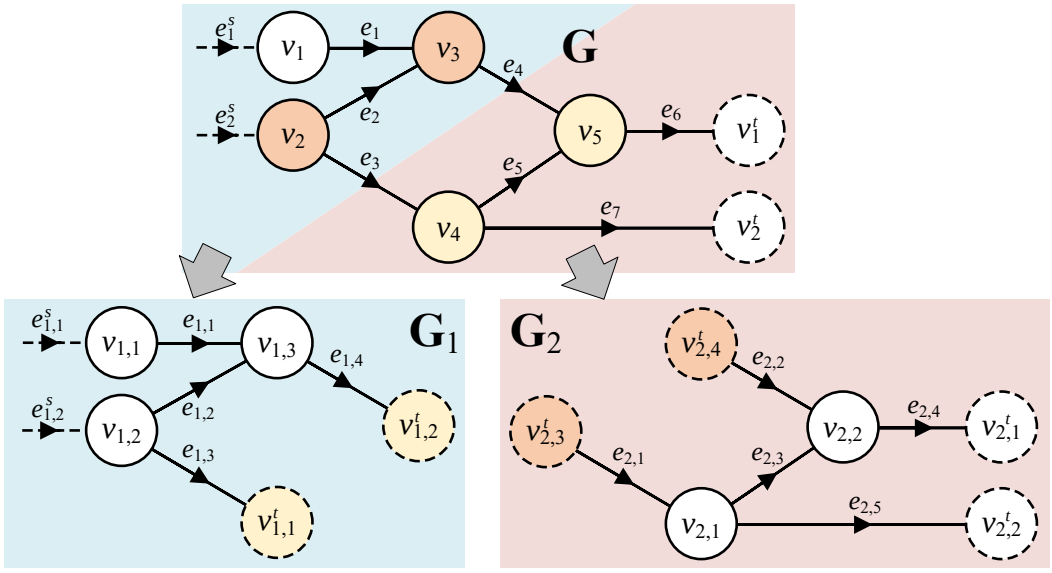

Figure 5.4: Example decomposition of a system graph $\mathbf{G}$ into two sub-graphs, $\mathbf{G}_1$ and $\mathbf{G}_2$.

At the faster timescale and shorter prediction horizon of the Level 2 controllers, it is likely that more refined bounds than the constant box constraints used in the Level 1 controller could be available for sink states that are exogenous to the overall system. For example, in a thermal management system, it might be possible to bound the ambient temperature over a short prediction horizon by measuring the current value and assuming a rate limit on its short-term evolution. To support this, bounds on sink states in the Level 2 control formulation can vary across steps of the prediction horizon, provided that they always fall within their corresponding box constraints in Level 1. In the Level 2 control formulation, sink state bounds are only required for steps of the prediction horizon prior to the next update of the Level 1 controller. Therefore, the bounds on the Level 2 sink state trajectories are denoted as $\underline{x}^t_{h+j|h}, \bar{x}^t_{h+j|h},\ \forall j \in [0{:}N_2 - n_{h|k} - 1]$, where $n_{h|k} \in [0{:}N_2 - 1]$ is the number of times the Level 2 controller has updated since the last update of the Level 1 controller.

While uncertainties in sink states of subsystems that are also sink states of the overall system can be treated as described above, the bounds on sink states associated with states of adjacent subsystems are held constant over the prediction horizon of each Level 2 controller. These bound are given by the corresponding elements of $\underline{y}_{k|k}, \bar{y}_{k|k}$ from the most recent update of the Level 1 controller. For example, let $\mathbf{G}$ in Fig. 5.4 be the graph used for a Level 1 controller, and $\mathbf{G}_2$ in Fig. 5.4 be the graph used for a controller of Level 2. The constant

bounds on the sink states associated with $v^t_{2,3}, v^t_{2,4}$ over the Level 2 prediction horizon then correspond to the elements of $\underline{y}_{k|k}$ and $\bar{y}_{k|k}$ communicated from the Level 1 controller for the states associated with v_4 and v_5 of $\mathbf{G}$. By comparison, the sink states associated with $v^t_{2,1}, v^t_{2,2}$ can have bounds that vary over the prediction horizon of Level 2, subject to the box constraints $\underline{x}^t_1 \leq \underline{x}^t_{2,1,h+j|h} \leq \bar{x}^t_{2,1,h+j|h} \leq \bar{x}^t_1$ and $\underline{x}^t_2 \leq \underline{x}^t_{2,2,h+j|h} \leq \bar{x}^t_{2,2,h+j|h} \leq \bar{x}^t_2$, $\forall j \in [1{:}N_2 - n_{h|k} - 1]$. For simplicity of notation in this chapter, the Level 2 sink state bounds are in general denoted as time-varying over the prediction horizon, where it is understood that the value of some of these signals may be constant over j.

Let $x^{cou}_{h+j|h}$, $\forall j \in [1{:}N_2 - n_{h|k} - 1]$ represent the subset of a subsystem's states $x_{h+j|h}$ that correspond to sink states of any other subsystem, and let $\underline{y}^{cou}_{k|k}, \bar{y}^{cou}_{k|k}$ represent the corresponding elements of the RPI set bounds from Level 1. Here, *"cou"* stands for "coupled" to indicate that these states are coupled to states of other subsystems. For example, $x^{cou}_{h+j|h}$ of $\mathbf{G}_2$ in Fig. 5.4 would consist of the states associated with $v_{2,1}$, $v_{2,2}$, as they correspond to sinks $v^t_{1,1}$ and $v^t_{1,2}$ of $\mathbf{G}_1$. If $\mathbf{G}$ were used as the graph of the Level 1 controller, then $\underline{y}^{cou}_{k|k}, \bar{y}^{cou}_{k|k}$ of $\mathbf{G}_2$ would be the elements of $\underline{y}_{k|k}, \bar{y}_{k|k}$ corresponding to the states associated with v_4 and v_5 of $\mathbf{G}$. As aligns with Theorem 5.4 and Remark 5.5, if each subsystem ensures that $\underline{y}^{cou}_{k|k} \leq x^{cou}_{h+j|h} \leq \bar{y}^{cou}_{k|k}$ locally, then all subsystems can remain within their local subspace of the RPI set, and therefore the overall system can remain within the RPI set. Of importance to the result in this chapter is that *only* $x^{cou}_{h+j|h}$ of each subsystem is necessarily required to remain within the RPI set bounds between updates of the Level 1 controller.

Subsystem Modes:

Each subsystem model used for the Level 2 controllers is assumed to contain at least the same modes associated with its edge inputs as are included for the corresponding edges of the system model used for the Level 1 controller. The Level 2 controllers may also have access to additional modes. For example, at Level 1, the modes of operation included for a fluid pump might consist of only its maximum and minimum speeds, as may be sufficient for coordinating the slow timescale and long-term behavior of the overall system. However, under the subsystem-specific and faster timescale governance of the Level 2 controllers,

additional modes corresponding to multiple intermediate pump speeds can be included as well. This allows the Level 2 controllers to perform refinement of the high-level decisions above.

Let $e^{cou,sub}$ denote the subset of edges of a sub-graph by which it is coupled to any other sub-graph, and let $e^{cou,sys}$ be the corresponding edges of the system graph. For example, $e^{cou,sub}$ of $\mathbf{G}_1$ in Fig. 5.4 consists of $e_{1,3}$ and $e_{1,4}$, and the corresponding $e^{cou,sys}$ consists of e_3 and e_4. Let $u_\sigma^{cou,sub}$ denote the inputs associated with edges $e^{cou,sub}$ in a given Level 2 controller, as a function of its switching decision variable σ^{sub}. Similarly, let $u_\sigma^{cou,sys}$ denote the inputs associated with the corresponding edges in the Level 1 controller, as a function of its switching decision variable σ^{sys}. The Level 1 controller is responsible for determining the switching of inputs associated with edges $e^{cou,sys}$, coordinating the exchange of energy between subsystems. Therefore, the mode selection in each Level 2 controller is constrained as follows:

$$\sigma_{h+j|h}^{sub} \in \{\sigma^{sub} \mid u_{\sigma^{sub}}^{sub,cou} = u_{\sigma_{k|k}^{sys}}^{sys,cou}\},\ \forall j \in [0{:}N_2 - 1]. \tag{5.24}$$

This ensures that the Level 2 controllers are consistent in their knowledge of the edge inputs by which subsystems are coupled, without needing to be in direct communication with each other.

An implication of the above is that the subset of subsystem states included in $x_{h+j|h}^{cou}$, $\underline{y}_{k|k}^{cou}$, and $\bar{y}_{k|k}^{cou}$ can be mode-dependent. For example, in the demonstration in Sections 5.5 and 5.6, subsystems can be coupled or decoupled from each other by opening or closing solenoid valves. To indicate their dependence on σ, these signals will hereafter be denoted as $x_{h+j|h}^{cou,\sigma}$, $\underline{y}_{k|k}^{cou,\sigma}$, and $\bar{y}_{k|k}^{cou,\sigma}$.

Additional constraints may be required on the modes available to the Level 1 and Level 2 controllers as necessary to ensure satisfaction of Assumption 5.5. These constraints can be found a priori for each mode of the Level 1 controller by enumerating all combinations of modes of the Level 2 controllers satisfying (5.24) and removing those for which the overall system graph does not satisfy Assumption 5.5.

State Trajectories:

Three state trajectories are computed by each Level 2 controller, as exemplified on the lower axes of Fig. 5.5 for the case when $n_{h|k} = 0$ (i.e., the case when the Level 2 controllers have yet to be updated since the last update of the Level 1 controller). The middle trajectory, plotted in purple, represents a predicted state $\hat{x}_{h+j|h},\ \forall j \in [0{:}N_2]$. This is based on the predicted sink states $\hat{x}^t_{h+j|h},\ \forall j \in [0{:}N_2-1]$ and nominal source power flow decision variables $P^s_{h+j|h},\ \forall j \in [0{:}N_2-1]$. The remaining two trajectories, plotted in black, represent lower and upper bounds on the true state trajectory until the next update of the Level 1 controller, and are denoted as $\underline{z}_{h+j|h}, \bar{z}_{h+j|h},\ \forall j \in [0{:}N_2 - n_{h|k}]$. The lower bound $\underline{z}_{h+j|h}$ is calculated using the lower bounds on the sink states $\underline{x}^t_{h+j|j}$ and an associated source power flow signal $\underline{P}^{s,z}_{h+j|h}$, which is a decision variable of the controller. Similarly, the upper bound $\bar{z}_{h+j|h}$ is calculated using the upper bounds on the sink states $\bar{x}^t_{h+j|j}$ and an associated source power flow signal $\bar{P}^{s,z}_{h+j|h}$. By definition of a discrete-time cooperative system (i.e., the discrete-time equivalent of Definition 5.5), the true state trajectory will be bounded as $\underline{z}_{h+j|h} \leq x_{h+j|h} \leq \bar{z}_{h+j|h},\ \forall j \in [0{:}N_2 - n_{h|k}]$, provided that the sink states obey they their constraints and the applied source power flow falls between $\underline{P}^{s,z}_{h+j|h}$ and $\bar{P}^{s,z}_{h+j|h}$.

Leveraging the trajectory bounding property of cooperative systems is key to achieving less conservative operation under the Level 2 controllers than that planned by the Level 1 controller. The Level 2 controllers are given the flexibility to guide their states outside of the RPI set bounds calculated by the Level 1 controller, provided that they can ensure recovery to within the RPI set bounds by the next update of the Level 1 controller. This is depicted in Fig. 5.5, as both the predicted state trajectory and the maximal state trajectory leave the blue shaded area representing the RPI set, but return within it by step $k+1|k$, when the Level 1 controller is next due to update.

Figure 5.6 shows an example of a state trajectory computed by a Level 2 controller for the case when $n_{h|k} = 1$. The predicted state $\hat{x}_{h+j|h}$ is still calculated over a horizon of N_2 steps, however the state trajectory bounds are now one step shorter than in Fig. 5.5, still concluding on the step at which the Level 1 controller is next due to update.

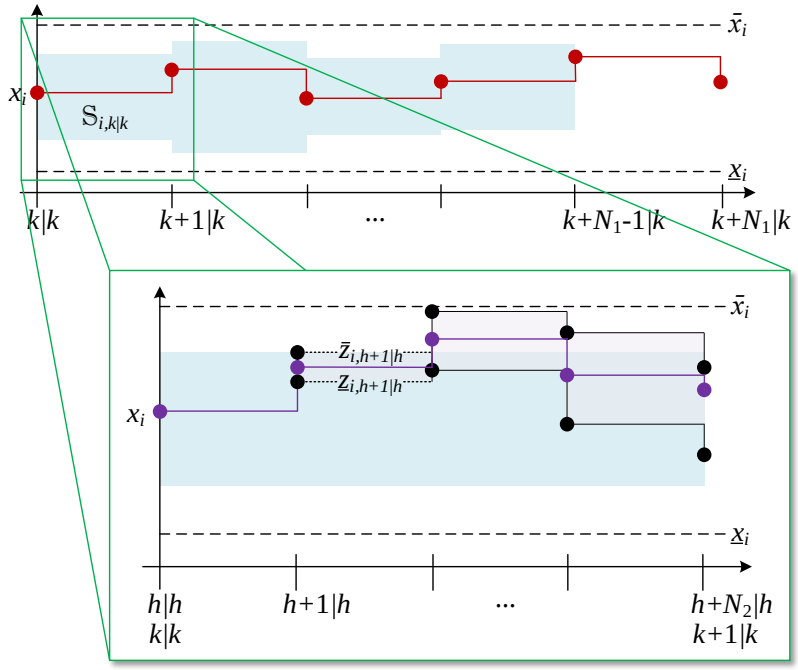

Figure 5.5: Sample predicted trajectory for a state x_i of a Level 2 controller for $n_{h|k} = 0$, showing the state constraints and equilibrium-based RPI set bounds.

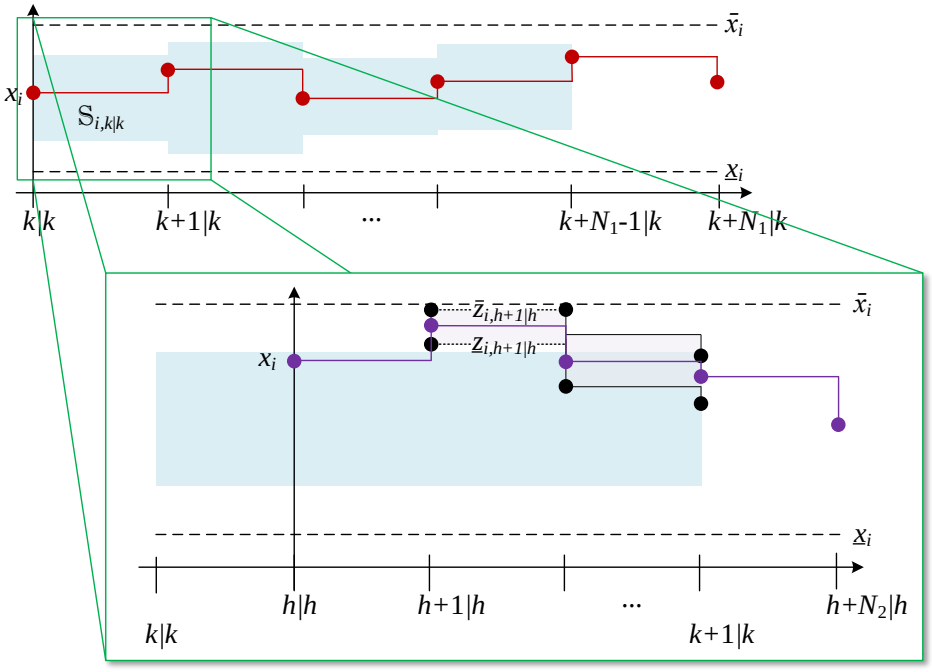

Figure 5.6: Sample predicted trajectory for a state x_i of a Level 2 controller for $n_{h|k} = 1$.

Each Level 2 MPC controller is formulated using the following MIQP:

$$J^*(x_h) = \min_{\boldsymbol{P}^s_h, \boldsymbol{\sigma}_h} \sum_{j=0}^{N_2-1} \ell_2\left(P^{s,des}_{h+j|h}, P^s_{h+j|h}, m_{h+j|h}\right), \tag{5.25a}$$

subject to, $\forall j \in [0{:}N_2-1]$,

$$\hat{x}_{h+j+1|h} = A^{d_2}_{\sigma_{h+j|h}} \hat{x}_{h+j|h} + B^{d_2}_{\sigma_{h+j|h}} P^s_{h+j|h} + E^{d_2}_{\sigma_{h+j|h}} \hat{x}^t_{h+j|h}, \tag{5.25b}$$

$$m_{i,h+j|h} = \begin{cases} 1 & \text{if } \sigma_{h+j|h} = i, \\ 0 & \text{else,} \end{cases} \tag{5.25c}$$

$$\sigma_{h+j|h} \in \{\sigma \mid u^{cou}_{\sigma} = u^{sys,cou}_{\sigma^{sys}_{k|k}} \text{ and Assumption 5.5 is satisfied}\}, \tag{5.25d}$$

$\forall j \in \left[0{:}N_2 - n_{h|k} - 1\right]$,

$$\underline{z}_{h+j+1|h} = A^{d_2}_{\sigma_{h+j|h}} \underline{z}_{h+j|h} + B^{d_2}_{\sigma_{h+j|h}} \underline{P}^{s,z}_{h+j|h} + E^{d_2}_{\sigma_{h+j|h}} \underline{x}^t_{h+j|h}, \tag{5.25e}$$

$$\bar{z}_{h+j+1|h} = A^{d_2}_{\sigma_{h+j|h}} \bar{z}_{h+j|h} + B^{d_2}_{\sigma_{h+j|h}} \bar{P}^{s,z}_{h+j|h} + E^{d_2}_{\sigma_{h+j|h}} \bar{x}^t_{h+j|h}, \tag{5.25f}$$

$$\underline{x} \le \underline{z}_{h+1|h} \le \bar{z}_{h+1|h} \le \bar{x}, \tag{5.25g}$$

$$\underline{P}^s \le \underline{P}^{s,z}_{h+j|h} \le P^s_{h+j|h} \le \bar{P}^{s,z}_{h+j|h} \le \bar{P}^s, \tag{5.25h}$$

$$\underline{y}^{cou,\sigma}_{k|k} \le \underline{z}^{cou,\sigma}_{h+1|h} \le \bar{z}^{cou,\sigma}_{h+1|h} \le \bar{y}^{cou,\sigma}_{k|k}, \tag{5.25i}$$

if $n_{h|k} \ge 1$, $\forall j \in \left[N_2 - n_{h|k}{:}N_2 - 1\right]$,

$$\underline{x} \le \hat{x}_{h+j+1|h} \le \bar{x}, \tag{5.25j}$$

$$\underline{P}^s \le P^s_{h+j|h} \le \bar{P}^s, \tag{5.25k}$$

$$\underline{y}_{k|k} \le \underline{z}_{N_2-n_{h|k}|h} \le \bar{z}_{N_2-n_{h|k}|h} \le \bar{y}_{k|k}, \tag{5.25l}$$

$$\hat{x}_{h|h} = \underline{z}_{h|h} = \bar{z}_{h|h} = x_h. \tag{5.25m}$$

Equation (5.25a) defines the objective function of the controller, for which the stage cost $\ell_2(\cdot)$ is a quadratic function. This will typically include a term penalizing errors in the tracking of $P^{s,des}$ by P^s, given by $\|P^{s,des}_{h+j|h} - P^s_{h+j|h}\|^2_{\Lambda_P}$. As in (5.23a), the third argument of $\ell_2(\cdot)$ allows costs associated with mode selection to be applied, where (5.25c) performs the same function as (5.23c) of the Level 1 controller formulation.

Equation (5.25b) defines the predicted subsystem dynamics at the update rate of the Level 2 controllers, subject to the nominal source power flow P^s and predicted sink state $\hat{x}^t$. Equation (5.25d) applies the mode selection constraints from (5.24), as well as any additional mode selection constraints required to ensure satisfaction of Assumption 5.5, as previously discussed. Equations (5.25e)-(5.25f) solve for the state trajectory bounds until the next update of the Level 1 controller. Equation (5.25g) requires that the trajectory bounds fall within the state constraints, while (5.25h) requires that the nominal source power flow falls between the source power flows used to compute the trajectory bounds, which in turn fall within the source power flow constraints. Equation (5.25i) requires that the trajectory bounds on any states of the subsystem corresponding to sink states of other subsystems in the current mode remain within their RPI set bounds computed by the Level 1 controller. This ensures that all other subsystem states can remain within their RPI set bounds.

Equations (5.25j) and (5.25k) constrain the predicted states and source power flows of steps following the next update of the Level 1 controller. These constraints are not explicitly necessary to the forthcoming proof, but are useful for bringing the Level 2 controllers to plan reasonable operation beyond the next update of the Level 1 controller.

Equation (5.25l) requires that the trajectory bounds fall within the RPI set constraints in the step corresponding to the next update of the Level 1 controller. This guarantees recursive feasibility of the Level 1 controller by ensuring that its initial state is within an RPI set. Lastly, (5.25m) sets the initial values of the predicted state and trajectory bounds equal to the current state of the subsystem.

Theorem 5.5: *Recursive Feasibility of Hierarchical Control Framework*
Consider the hierarchical control framework shown in Fig. 5.2. Between step h and step $h+1$ of the Level 2 controllers, the applied source power flows $P_{i,h}^{s,app}$ of the i^{th} subsystem are constant and bounded as $\underline{P}_{i,h}^{s,z} \leq P_{i,h}^{s,app} \leq \bar{P}_{i,h}^{s,z}$, and the sink state disturbances are constant and obey their assumed bounds, $\underline{x}_{i,h}^t \leq x_{i,h}^t \leq \bar{x}_{i,h}^t$. If the Level 1 controller has a feasible solution upon initialization, then all Level 2 controllers have a feasible solution and all controllers of the hierarchical framework are recursively feasible. Specifically, this implies that the discretized state trajectories always satisfy their their constraints.

Proof: Feasibility of the Level 1 controller upon initialization guarantees feasibility of all Level 2 controllers, as they can apply the same modes and source power flows as selected in the Level 1 controller and maintain all of their states within the RPI set bounds, in alignment with Theorem 5.4 and Remark 5.5. However, each Level 2 controller can also choose to apply different modes and/or govern its states outside the bounds of the RPI set if two two conditions are met. First, all states must remain within their constraints $\underline{x}, \bar{x}$ and recover to within the RPI set by the next update of the Level 1 controller. This is ensured by leveraging the properties of cooperative systems to directly calculate maximal and minimal state trajectories until the next update of the Level 1 controller in (5.25e)-(5.25f), and then imposing appropriate constraints on these trajectories in (5.25g) and (5.25l). The second condition is that any states that are sink states of other subsystems under the current mode always remain within their RPI set bounds, as imposed by (5.25i). This ensures that adjacent subsystems are able to maintain their states within their RPI set bounds if necessary.

The above ensures recursive feasibility of the Level 1 controller, as the RPI set and associated mode determined in the previous iteration can always be used for the current iteration. Feasibility of the Level 1 controller again guarantees feasibility of all Level 2 controllers, resulting in recursive feasibility of the hierarchical framework as a whole. □

5.5 Simulation Example

The efficacy of the proposed control approach is demonstrated with the following numerical example, which is extended to experimental application in the next section. Figure 5.7 shows a candidate fluid-thermal system architecture, consisting of components modeled in Chapter 2. Thermal energy is transferred to the system via heat loads applied to the four cold plate walls. In a thermal management system, these heat loads correspond to thermal energy generation due to inefficiencies in electrical equipment, such as the battery and converters of Chapter 3, or other components that generate thermal energy. Thermal energy can be removed from the fluid-thermal system by heat transfer to two sinks representing the ambient environment or other thermal storage and transport systems, such as vapor compression systems or air cycle machines. In this example, the temperatures of these sinks are treated

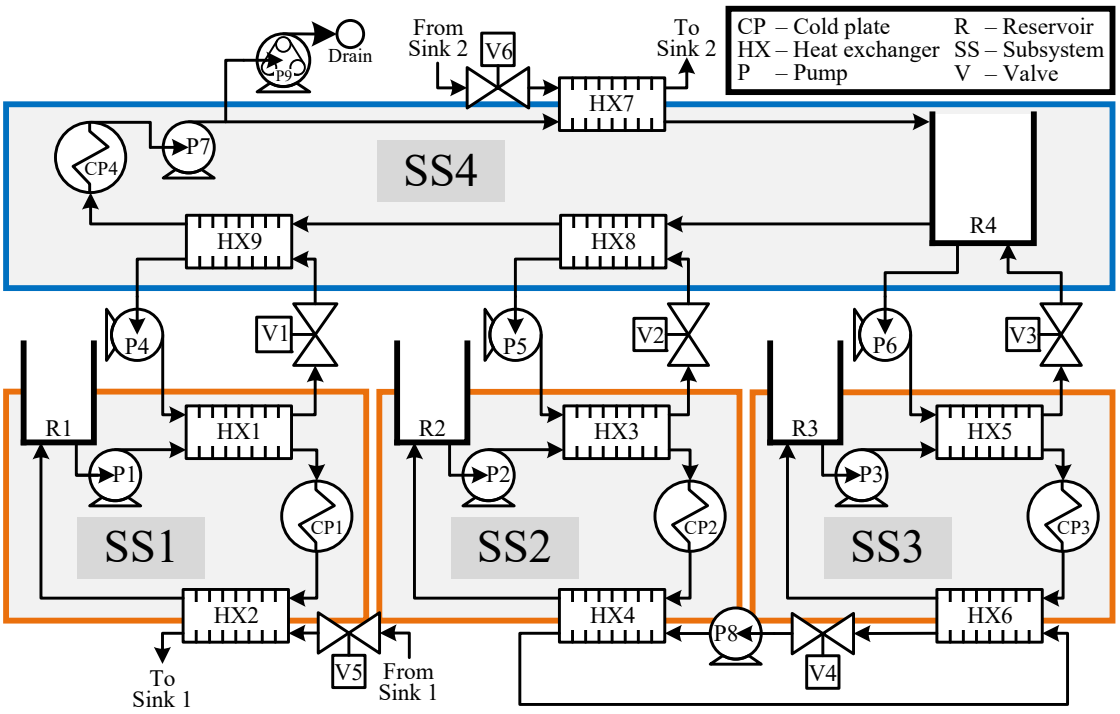

Figure 5.7: Schematic of the fluid-thermal system architecture used to demonstrate the proposed control approach.

as bounded exogenous disturbances for which nominal (but possibly inaccurate) preview is available.

Figure 5.8 shows the graph-based model corresponding to the schematic in Fig. 5.7. The model is parameterized in accordance with the specifications of the testbed components detailed in Section 2.4.1. The 39 states x of the model represent the temperatures of fluids, heat exchanger walls, and cold plate walls throughout the system. The three sink states x^t represent the fluid drain and the fluid temperatures of flow from the thermal sinks. The 47 power flows P capture the transfer of thermal energy among states and sinks of the system. Four source power flows P^s represent the heat loads applied to the walls of the four cold plates. The 12 inputs u are the unique fluid mass flow rates through the system. These are a function of the commanded inputs to the nine pumps (including the drain pump) and six valves. When the graph-based system model is represented in the switched linear form of (5.10), each mode corresponds to a particular set of these inputs. For the purposes of the demonstration in this chapter, the drain pump is not powered. While the remaining

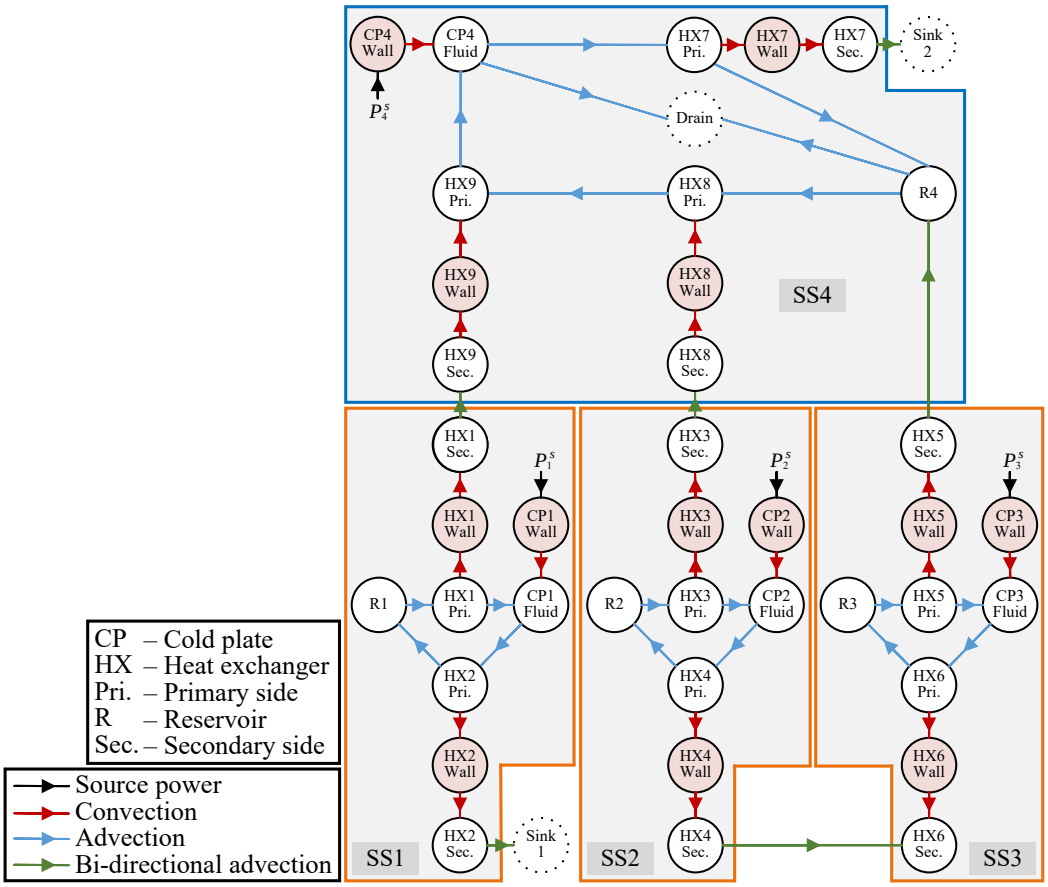

Figure 5.8: Graph-based for model of the fluid-thermal system in Fig. 5.7.

eight pumps of the system can have any number of modes corresponding to different speed commands, the valves can only be fully open or fully closed. Similar to the case study in Chapter 3, valves can be closed to decouple otherwise adjacent subsystems or bypass thermal sinks.

The goals of achieving high capability, efficiency, and safety that motivate the design of the hierarchical control framework in Section 5.4.3 apply directly to the context of thermal management in this example. Capability is represented as an objective of tracking references for the heat load applied to each cold plate. These references represent thermal energy generation associated with the desired operation of electrical or mechanical systems. Discrepancies between the desired and allowable heat loads are communicated to the systems generating the loads, which throttle down their operation as necessary to operate within their thermal budget. The scope of the current example focuses exclusively on the thermal

management aspect, controlling the fluid-thermal system to absorb as much of the desired thermal loading as possible over a mission while ensuring that at least minimum flight-critical loads can always be absorbed without violating temperature constraints. Satisfying these constraints can be essential to ensuring safe and reliable operation of both flight-critical and mission-critical equipment. Therefore, safety is captured by enforcing hard box constraints on the temperature states of the system. The goal of efficient operation is represented in the objective function by penalizing each mode of operation in proportion to the electrical energy consumption by the actuators of the thermal management system under that mode.

In this example, the stage costs $\ell_1(\cdot)$ and $\ell_2(\cdot)$ of the Level 1 and Level 2 control formulations, respectively, are defined as follows:

$$\ell_1\left(P^{s,des}_{k+j|k}, P^{s}_{k+j|k}, m_{k+j|k}\right) = \|P^{s,des}_{k+j|k} - P^{s}_{k+j|k}\|^2_{\Lambda_P} + \|m_{k+j|k}\|^2_{\Lambda_u}, \tag{5.26a}$$

$$\begin{aligned}\ell_2\left(P^{s,des}_{h+j|h}, P^{s}_{h+j|h}, m_{h+j|h}\right) = {} & \|P^{s,des}_{h+j|h} - P^{s}_{h+j|h}\|^2_{\Lambda_P} + \|m_{h+j|h}\|^2_{\Lambda_u} \\ & + \|P^{s}_{h+j|h} - P^{s}_{h+j-1|h}\|^2_{\Lambda_{\Delta P}} \\ & + \|m_{h+j|h} - m_{h+j-1|h}\|^2_{\Lambda_{\Delta u}}.\end{aligned} \tag{5.26b}$$

The first term on the right in (5.26a) represents costs associated with load shedding by penalizing deviations from the desired source power flows by the nominal source power flow decision variable P^s. The second term applies costs associated with mode selection, where (5.23c) defines $m_{k+j|k}$ as a vector of binary variables for each j, with unity in the row corresponding to the index of the active mode and zeros elsewhere. In addition to the two terms in (5.26a), (5.26b) includes additional terms to smooth the nominal source power flow trajectory and penalize mode switching to reduce actuator wear, where $P^s_{h-1|h} = P^s_{h-1|h-1}$ and $m_{h-1|h} = m_{h-1|h-1}$.

Figure 5.9 shows the specific hierarchical control framework used for this example. All four subsystems (SSs) are included in the Level 1 controller. Two controllers are used in Level 2. One of these governs SS1 and SS4, while the other governs SS2 and SS3. The reason for this grouping of subsystems in Level 2 will become clear in the experimental demonstration of the next section. The Level 1 controller updates every $\tau^{d_1} = 60$ seconds with a horizon

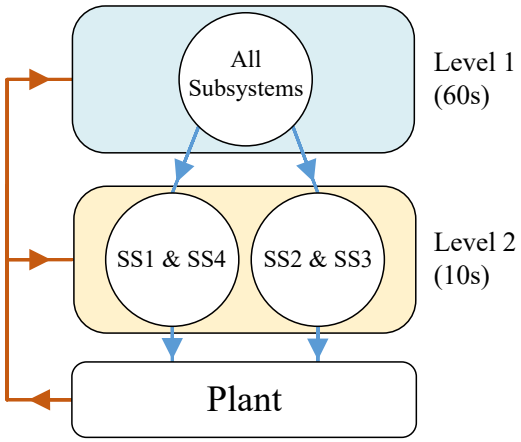

Figure 5.9: Hierarchical control framework for the simulation example.

of $N_1 = 5$ steps. This allows it to capture the slow dynamics of the system and optimize behavior over a five-minute prediction horizon. Each of the Level 2 controllers update every $\tau^{d_2} = 10$ seconds with a horizon of $N_2 = 6$ steps, allowing them to govern faster dynamics of the system.

Figure 5.10 shows the desired heat loads and sink states for the this example. The flight-critical loads are given by $\underline{P}^s = \begin{bmatrix} 0.75 & 0.3 & 0 & 0.5 \end{bmatrix}^T$ kW. The sheddable load profiles consist of low-frequency steps associated with the nominal operating state of heat-generating components, to which higher frequency variability is added, representing finer adjustments in operation made to achieve electrical and mechanical objectives, such as reference tracking and disturbance rejection. The profile for Heat Load 3 is designed to be characteristic of systems with high ratios of peak to average power, which are of particular interest in some vehicle energy system applications [20]. While the Level 2 controllers receive the profile of the sheddable loads over their prediction horizon, the Level 1 controller receives a nominal, downsampled representation of the sheddable loads over its prediction horizon, as shown by the dashed blue traces. This reflects the notion that more refined knowledge of the desired operation of the system may be available over the short time horizon of the Level 2 controllers than over the long time horizon of the Level 1 controller.

Excluding the sink state of the drain, which has no effect on the system in this example, the sink states are bounded as $\underline{x}^t = \begin{bmatrix} 15 & 15 \end{bmatrix}^T$°C, $\bar{x}^t = \begin{bmatrix} 35 & 35 \end{bmatrix}^T$°C. The Level 1 controller receives these constant bounds and the nominal trajectories shown by the dashed blue traces

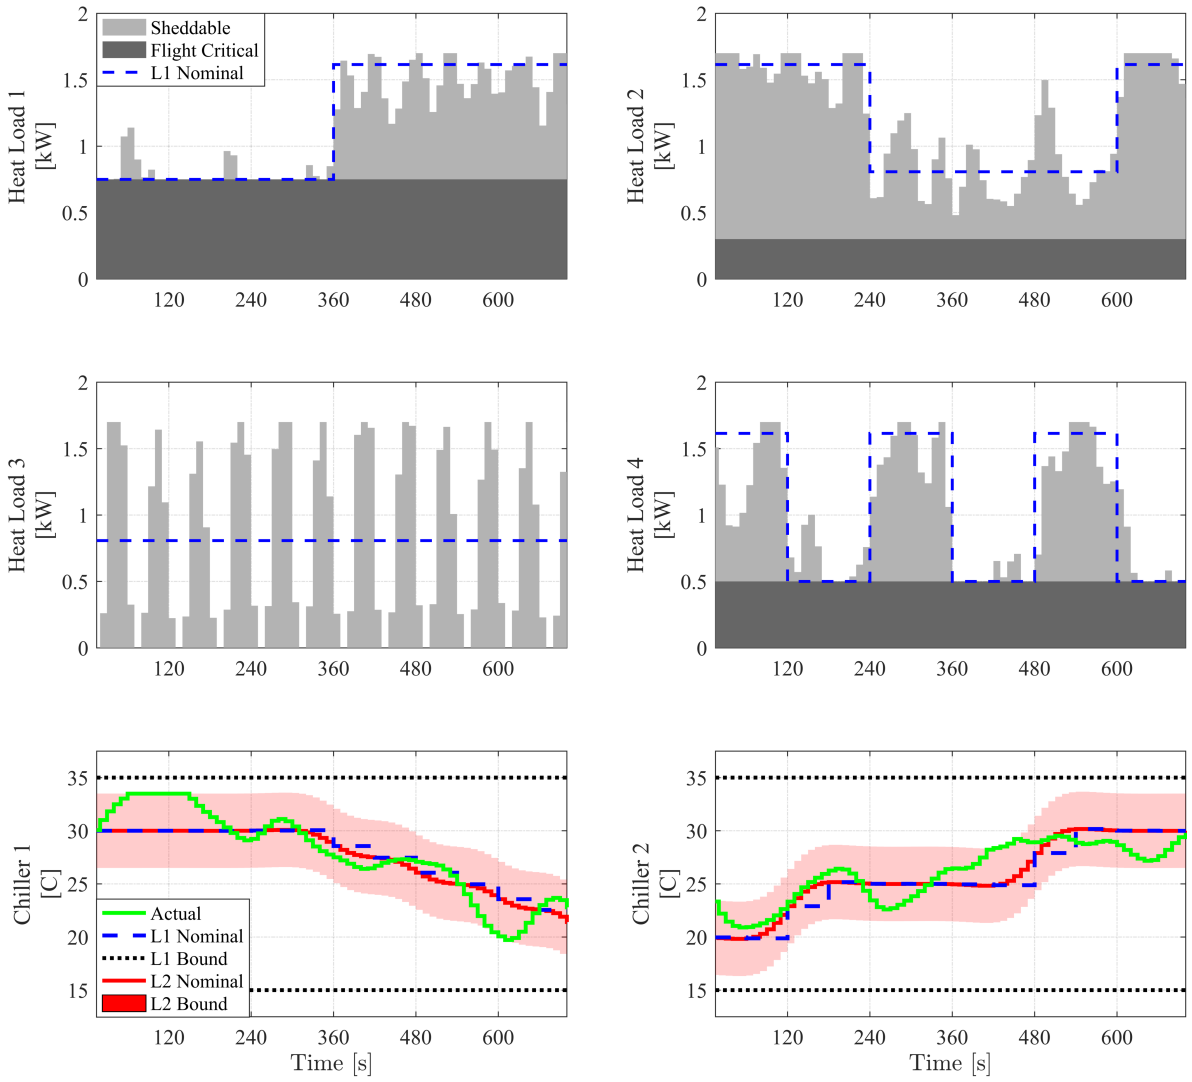

Figure 5.10: Desired source power flows and sink state disturbances of the simulation example.

in the bottom two subplots of Fig. 5.10. Over the short prediction horizon of the Level 2 controllers, less uncertainty is present in the sink states. These controllers receive the bounds defining the red envelopes in the bottom two subplots of Fig. 5.10, as well as nominal trajectories shown by the red traces. The actual sink states applied to the system are shown by the green traces. These deviate significantly from the nominal trajectories provided to each controller, but still satisfy the bounds assumed in Level 1 and Level 2. The hierarchical framework is robust in the sense that it guarantees state constraint satisfaction despite this presence of bounded uncertainties in the sink states.

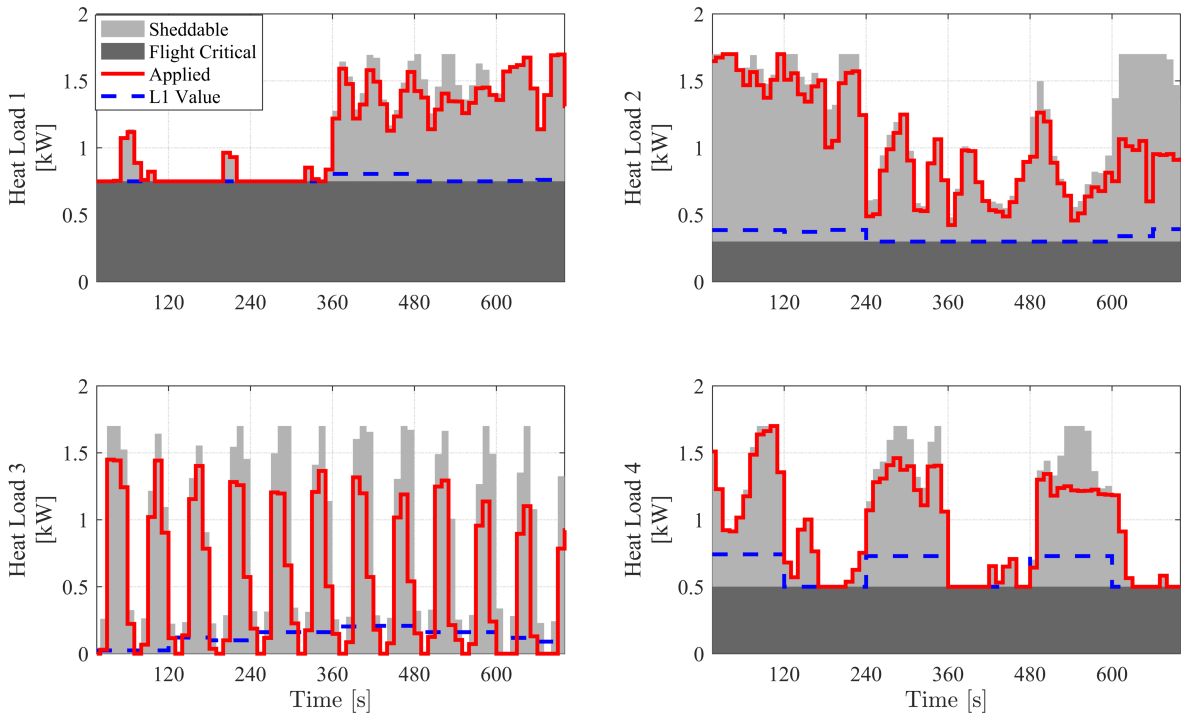

Figure 5.11: Desired and achieved source power flows of the simulation example.

The lower constraint on all temperature states is 15°C. The upper constraints are 40°C for fluids and 45°C for cold plate and heat exchanger walls, similar to those in the case study of Chapter 3.

While Theorem 5.5 guarantees recursive feasilbility and state constraint satisfaction subject to any applied power flows $P_h^{s,app}$ falling within the bounds $\underline{P}_h^{s,z} \leq P_h^{s,app} \leq \bar{P}_h^{s,z}$, in this example the nominal source power flow decision variables P_h^s are applied to the system to demonstrate the capability achieved when the controller has full authority over the source power flows. As shown in Figure 5.11, the hierarchical controller successfully achieves the flight critical loads and much of the sheddable loads, but does perform shedding as necessary to satisfy constraints and ensure recursive feasibility. The dashed blue traces in Figure 5.11 show the nominal source power flow decision variables of the Level 1 controller at each time step. In light of the much larger source power flows achieved by the Level 2 controllers, the equilibrium-based RPI sets calculated by the Level 1 controller clearly represent conservative bounds on the capability achievable by the system within its constraints. The Level 2

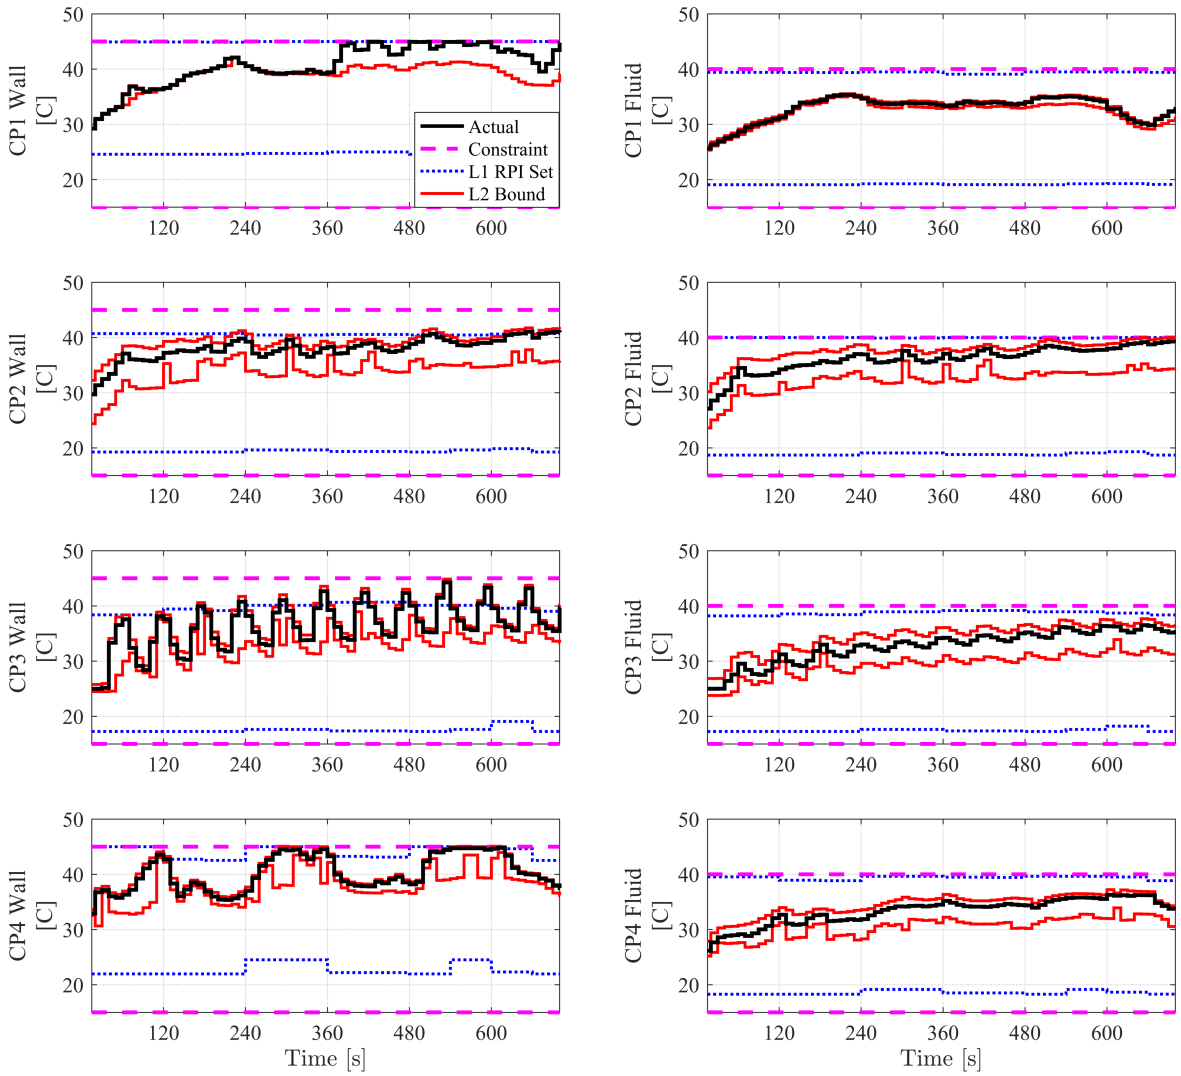

Figure 5.12: Selected states of the simulation example.

controllers significantly reduce this conservatism by bounding operation based on transient behavior rather than steady-state behavior.

Figure 5.12 shows the wall and outlet fluid temperatures of each cold plate, which are typically the highest temperatures of the system. As shown by the black traces, these always remain within their constraints. These states are also confirmed to be bounded by the trajectories $\underline{z}_h, \bar{z}_h$ determined by the Level 2 controllers, as shown by the red traces. Consistent with the goal of maximizing the capability of the system, the states often operate close to their upper trajectory bounds. This is particularly true of the CP1 wall, for which the trace for the state covers the trace for its upper bound. Figure 5.12 also shows the

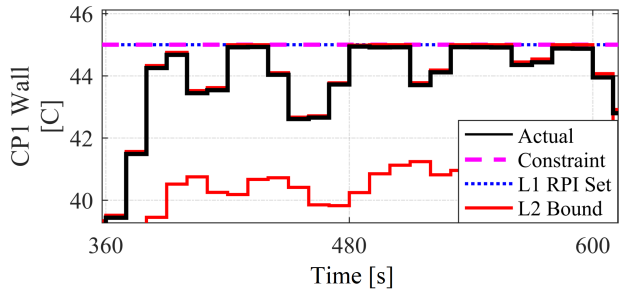

Figure 5.13: Closer view of the CP1 wall temperature between 360 and 600 s.

equilibrium-based RPI set bounds computed by the Level 1 controller for each state. Under the pulsed loading applied to the CP3 wall, the temperature of this wall is shown to exceed its RPI set bounds by up to 5°C at times. However, this state always recovers to within the RPI set bounds by the next update of the Level 1 controller (every 60 s), ensuring recursive feasibility. On several occasions, the system states or their trajectory bounds achieve their upper constraints $\bar{x}$. This is exemplified by Fig. 5.13, which shows a closer view of the CP1 wall temperature for a portion of the profile.

Figure 5.14 shows the mass flow rates applied to the system as a result of mode selection in the Level 1 and Level 2 controllers. When mass flow rates fall to exactly zero in the bottom three subplots, this indicates that a valve has been closed, decoupling two otherwise adjacent subsystems. In applications where less frequent switching is desirable to prevent actuator wear, the penalty on switching, given by $\Lambda_{\Delta u}$ in (5.26b), could be increased.

All MPC optimization problems of the simulation example were formulated with YALMIP [104] and solved with CPLEX [163] using a desktop computer with a 3.40 GHz Intel i7 processor and 16 GB of RAM. Computation times for all controllers were within their update intervals, as detailed in Table 5.1. This table also includes computation times for controllers of the experimental demonstration in the next section.

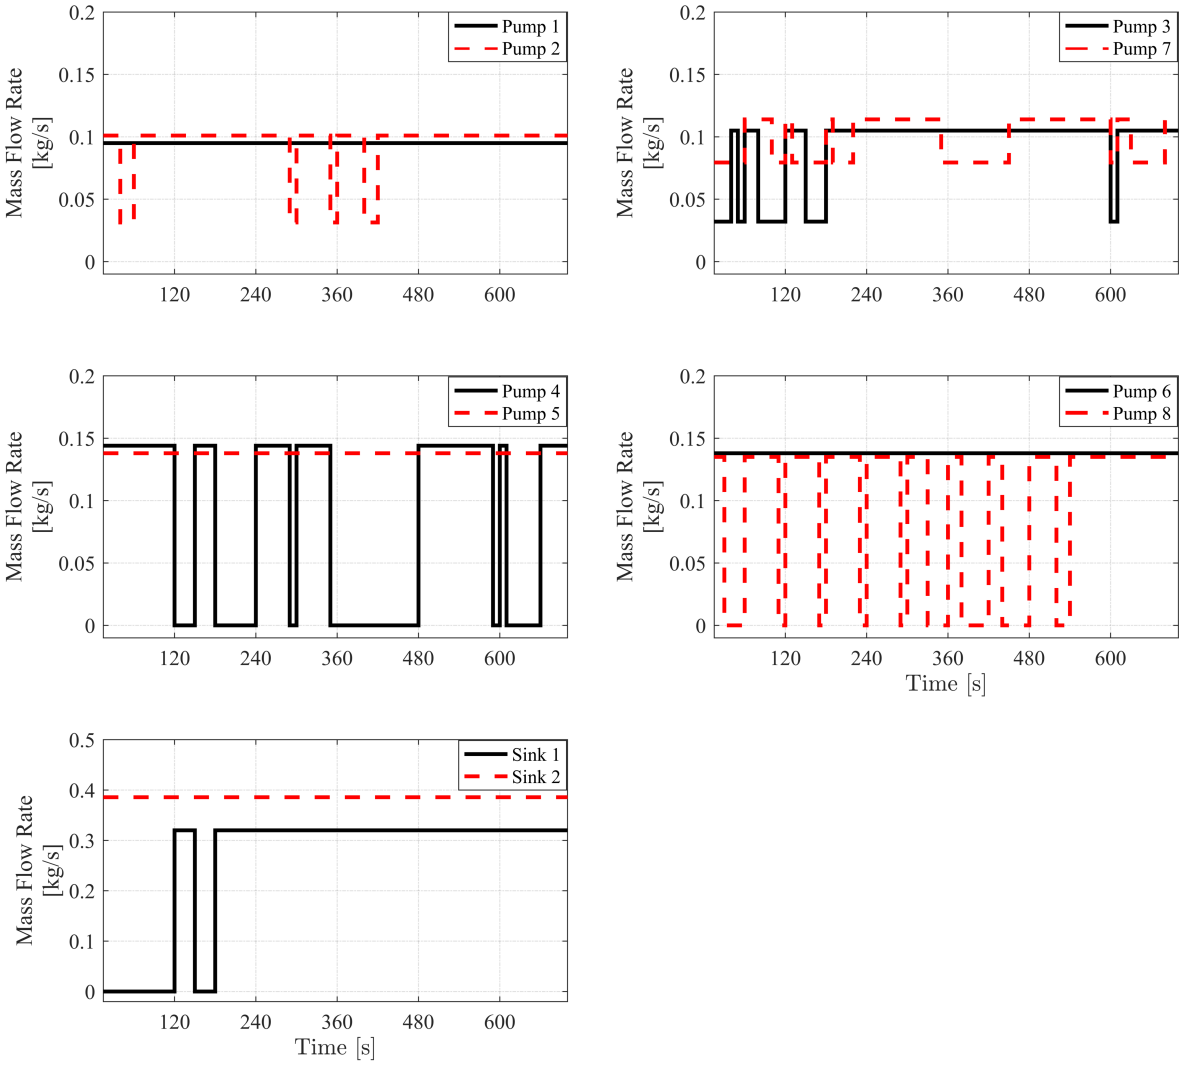

Figure 5.14: Mass flow rates of the simulation example.

Table 5.1: Number of modes, update intervals, and computation times of all controllers of the hierarchical MPC framework for the simulation example (sim) and experimental demonstration (exp).

Level	Controller	No. of modes	Update interval (s)	Avg. (s) (sim)	Avg. (s) (exp)	Max (s) (sim)	Max (s) (exp)
1	All Subsystems	12	60	49.0	43.0	57.2	57.2
2	SS1 & SS4	18	10	1.2	5.4	3.7	9.7
2	SS2 & SS3	16	10	1.8	5.8	4.1	9.7
3	SS1	N/A	2	N/A	0.14	N/A	0.78
3	SS2	N/A	2	N/A	0.13	N/A	0.66
3	SS3	N/A	2	N/A	0.14	N/A	0.75
3	SS4	N/A	2	N/A	0.13	N/A	0.59

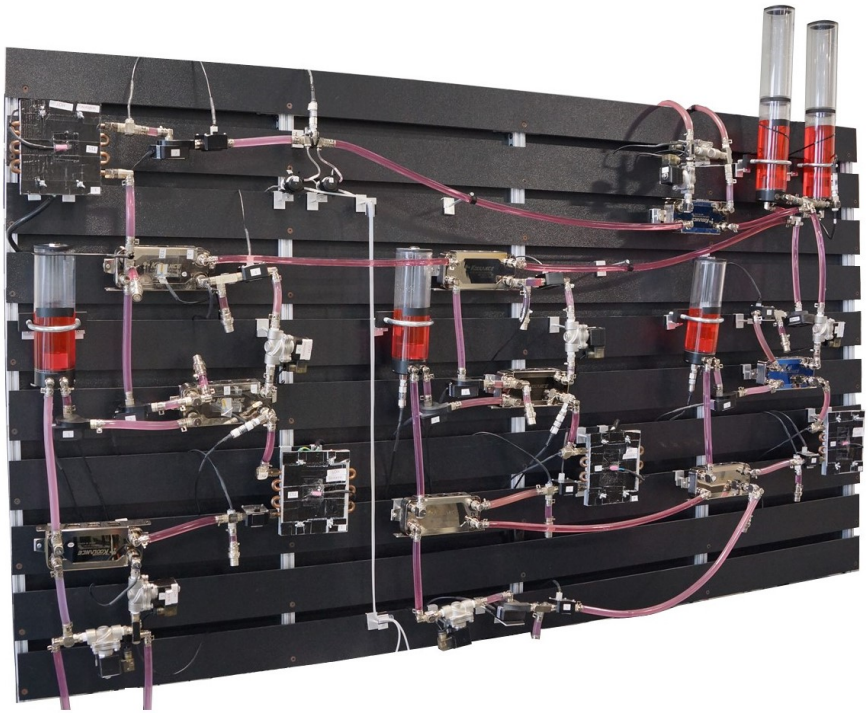

Figure 5.15: Fluid-thermal testbed configured to match the architecture in Fig. 5.7

5.6 Experimental Demonstration

To evaluate the ability of the control design in this chapter to translate from simulation to experimental application, the fluid-thermal testbed described in Section 2.4.1 was configured to match the architecture in Fig. 5.7 and used to further demonstrate the proposed hierarchical control approach. Fig. 5.15 shows the testbed in this configuration, while Fig. 5.16 shows the same image with fluid loops identified and component names labeled as defined in Fig. 5.7.

While the guarantee of hard state constraint satisfaction in Theorem 5.5 is a key contribution of this chapter, experimental application necessitates softening of these constraints to prevent model error, measurement and estimation error, and unknown disturbances from rendering controllers infeasible. Therefore, the control formulations in Section 5.4.3 are modified with slack variables on state constraints, similar to (3.12d) in the case study of Chapter 3. This is a standard practice in applied MPC as performed, for example, in [21]. The control formulations are also modified with known delays equal to their update intervals to account for computation times, similar to (3.12k)-(3.12l) in Chapter 3.

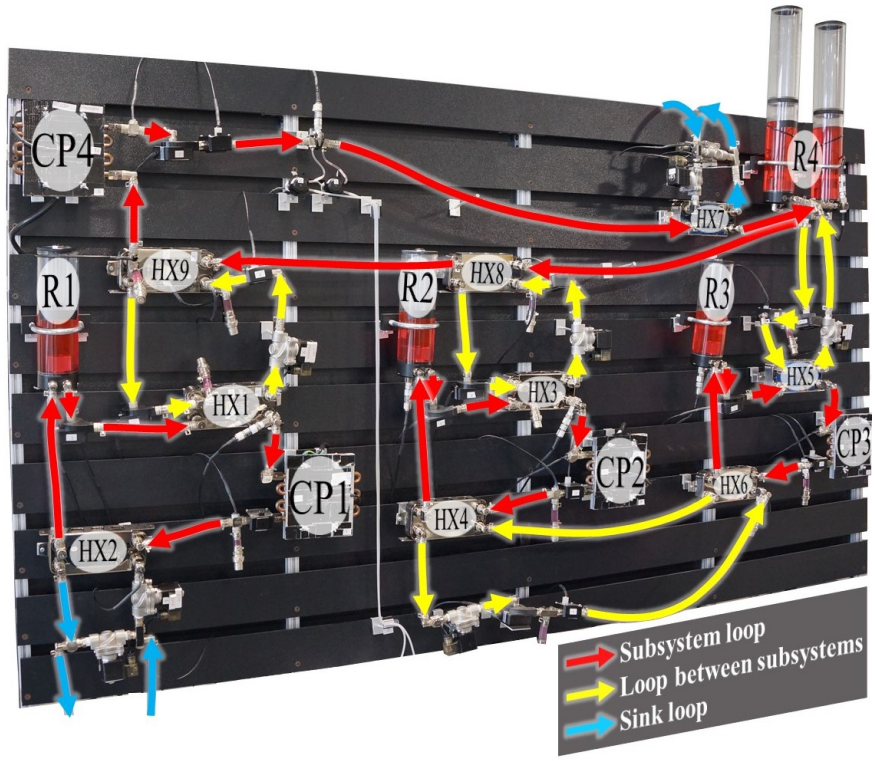

Figure 5.16: Reproduction of Fig. 5.15 with fluid loops identified and component names labeled as defined in Fig. 5.7.

While the hierarchical control framework studied in the theoretical results of this chapter consists of two levels, under the softening of constraints described above it is possible to include additional levels. To demonstrate this, the four-level hierarchical control framework in Fig. 5.17 is used in the following closed-loop experiments. In addition to the Level 1 and Level 2 controllers, a third level is added with controllers corresponding to each subsystem. These Level 3 controllers update every $\tau^{d_3} = 2$ seconds with a horizon of $N_3 = 5$ steps. These use the same general control formulation as in Level 2, with the exception that they are constrained to implement the modes selected in Level 2. The role of the Level 3 controllers in this case is to perform fast-timescale control of the applied source power flows, primarily compensating for model error in seeking to maintain the states within their softened constraints. Below Level 3, a tracking level similar to that used in Chapter 3 is included. This consists of SISO controllers that command pump speeds and valve positions to track the desired mass flow rates communicated from above.

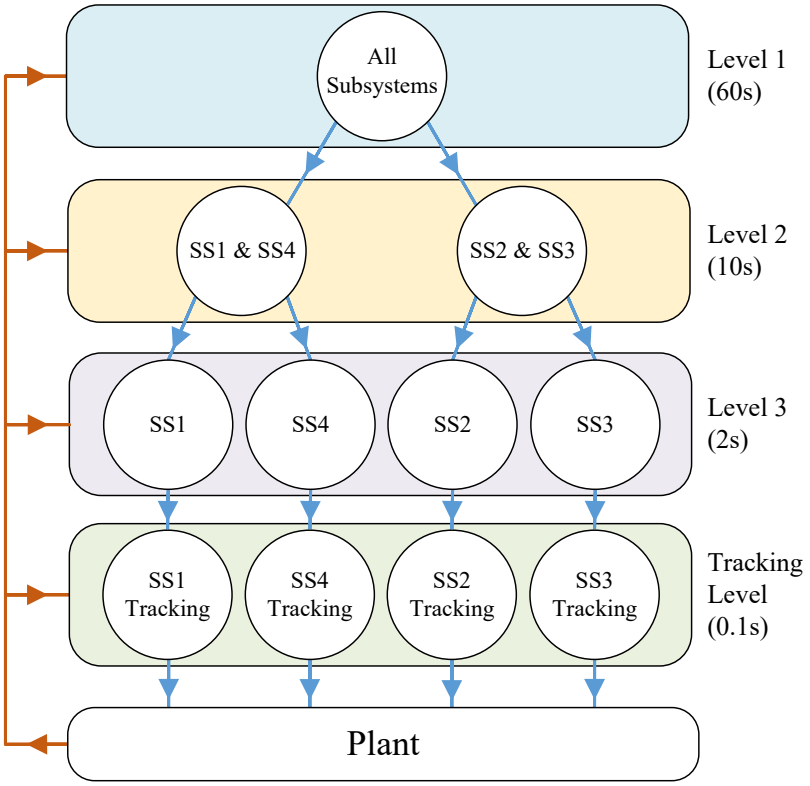

Figure 5.17: Hierarchical control framework for the experimental demonstration.

Fourteen of the 39 states of the system model are measured on the testbed. Estimates of the unmeasured states are made using an Extended Kalman Filter linearized about the current mode of operation at each update of the Level 3 controllers. The temperatures of both chiller sink states are also measured. Based on the current measurements and an empirically derived rate limit on the sink states, bounds on the sink state temperatures across the prediction horizon of the Level 3 controllers can be calculated. The midpoints of these bounds are used as the predicted sink states for the Level 3 controllers.

For the experimental demonstration, all MPC optimization problems were again formulated with YALMIP [104] and solved with CPLEX [163] using a desktop computer with the same specifications as in the simulation example. In this case, the MPC controllers of the hierarchy were solved in parallel to enable real time implementation. Computation times for all controllers were within their update intervals, as detailed in Table 5.1. However, the Level 2 controllers did generally take longer to solve in the experiment than in the simu-

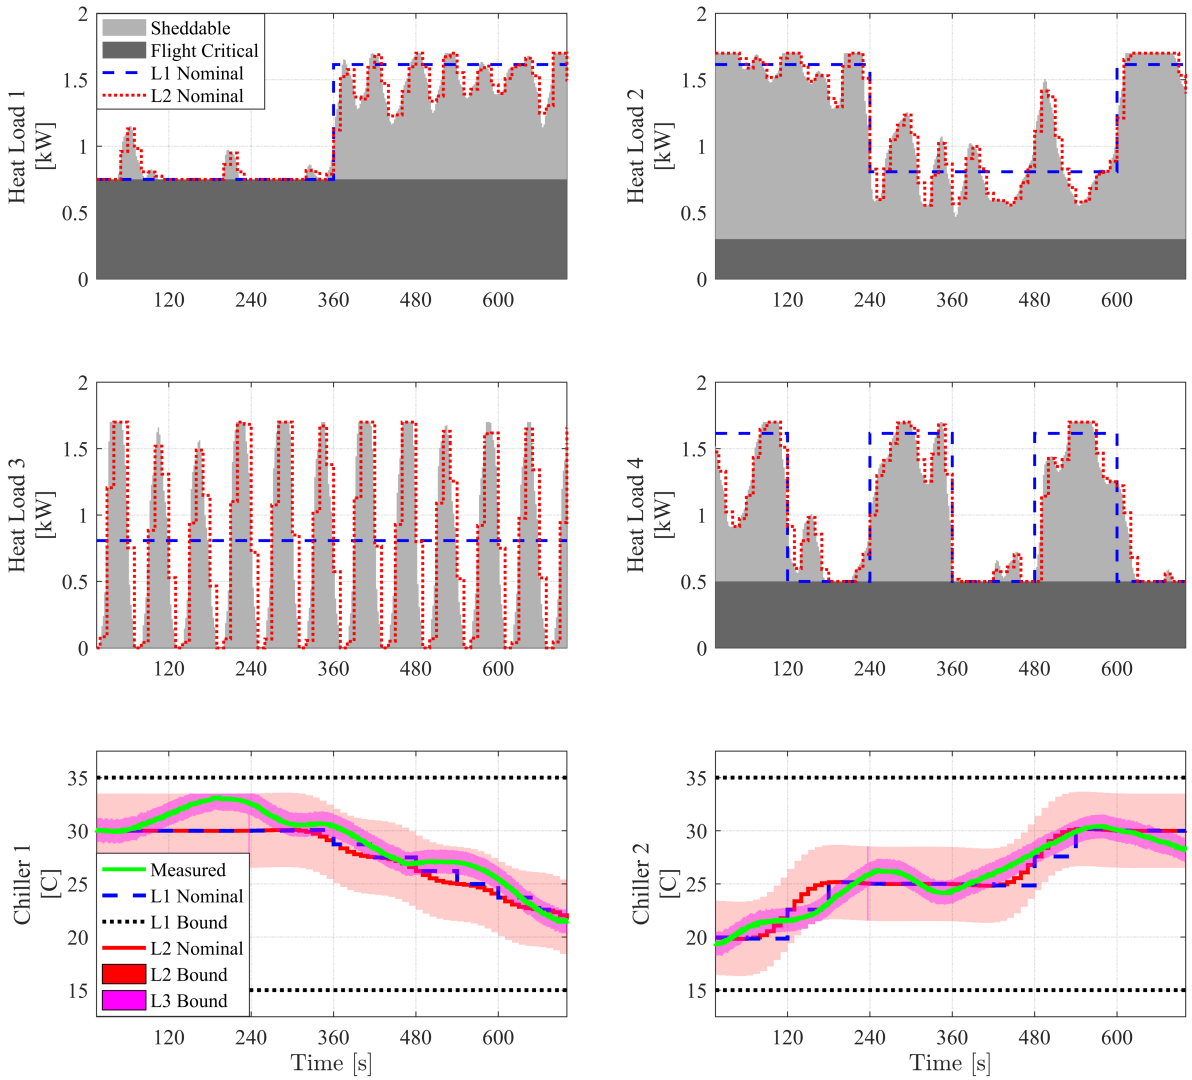

Figure 5.18: Desired source power flows and sink state disturbances of the experimental demonstration.

lation. This is likely due to a greater number of decision variables included in the control formulations for the experiment to soften state constraints.

Figure 5.18 shows the desired heat loads, measured sink states, and sink state bounds of the experimental demonstration. These are similar to those for the simulation example, with the exceptions that the sheddable load profiles evolve at the update rate of Level 3, and sink state bounds and predictions are included for Level 3.

Figure 5.19 shows the source power flows applied by the Level 3 controllers, as well as the corresponding nominal source power flow decision variables from the Level 1 and Level 2

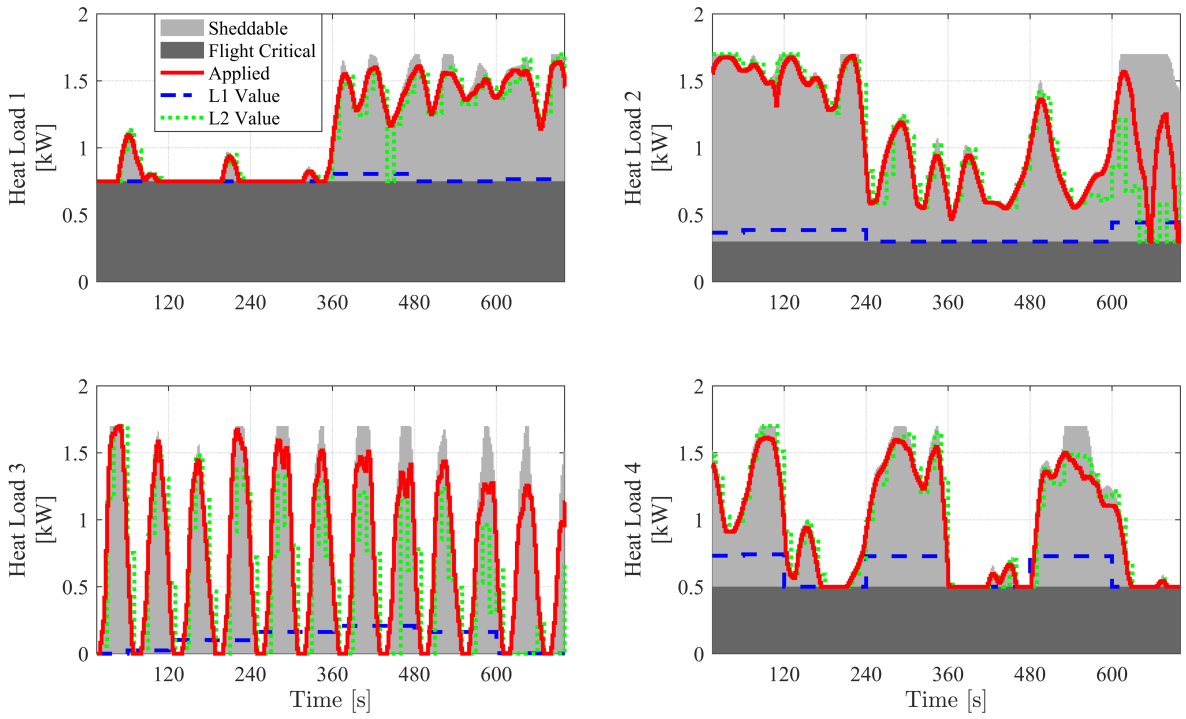

Figure 5.19: Desired and achieved source power flows of the experimental demonstration.

controllers. In general, similar capability is achieved as for the simulation example. Note that slightly different behavior is to be expected, as the sink states in the simulation differ from those measured in the experiment.

Figure 5.20 shows the measured wall and outlet fluid temperatures of each cold plate. Again, there is strong similarity to the simulation example. States such as the CP4 wall temperature are operated close to their constraints without violations. However, small temperature constraint violations do occur in the CP2 and CP3 fluid states. This is likely due to measurement and model errors, as well as reference tracking error in the Tracking Level responsible for achieving the fluid mass flow rates expected by the levels above. However, the state constraint violations in this case are smaller than in Chapter 3, even though the control formulations in Chapter 3 included artificial constraint tightening of 2°C. The state constraint violations are also smaller than the hierarchical MPC implementation for the same testbed in [64]. This suggests that accounting for sink state uncertainty throughout the hierarchical controller can improve its ability to respect state constraints in experimental application.

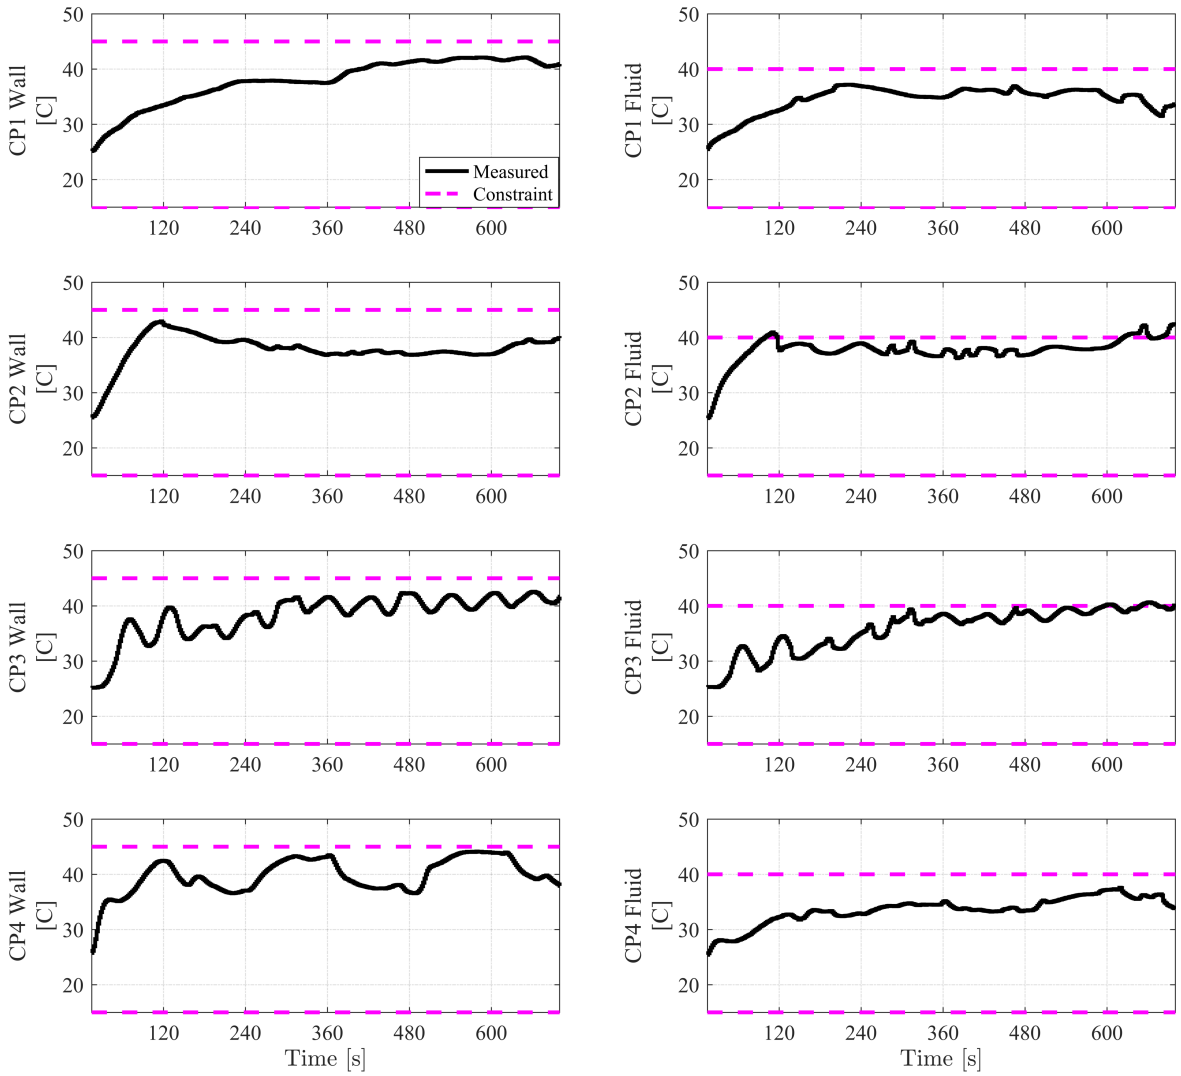

Figure 5.20: Selected state measurements of the experimental demonstration.

Figure 5.21 shows the desired mass flow rates sent to the Tracking Level a result of mode selection throughout the hierarchical framework. Again, the results are similar to those for the simulation example. If desired, the penalty on switching could again be increased to cause the Level 2 controllers to switch less frequently.

5.7 Chapter Summary

This chapter presents a two-level switched hierarchical control framework that guarantees satisfaction of state and input constraints for a class of graph-based dynamic systems. Each

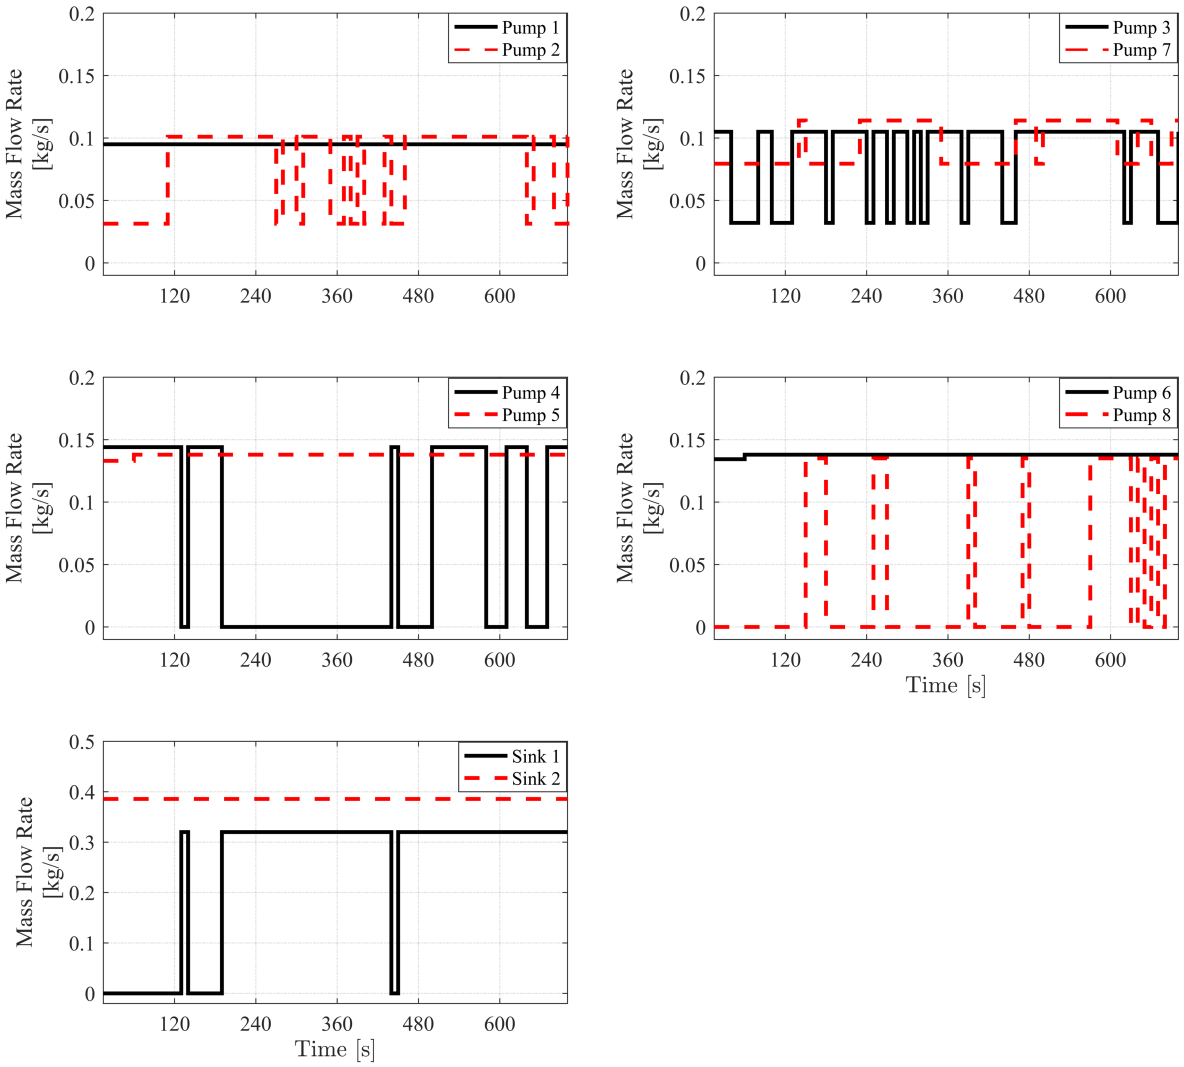

Figure 5.21: Mass flow rate references sent to the Tracking Level of the experimental demonstration.

mode in this class is a cooperative system, and can be proven open-loop stable under appropriate conditions on the connectivity of the graph and form of the equations governing power flow along its edges. The two-level hierarchical control framework is designed to achieve high capability, efficiency, and safety of switched energy systems, subject to bounded uncertainty in exogenous disturbances associated with the ability to dissipate energy to sinks. At the top level of the hierarchy, a slowly-updating controller plans operation over a long time horizon and performs online computation of RPI sets, leveraging the existence of equilibrium-based box-invariant sets for cooperative systems. At the second level of the hierarchy, multiple controllers corresponding to specific subsystems operate at a faster timescale. These con-

trollers have the flexibility to drive states beyond the RPI sets received from the level above, provided that they can guarantee recovery to within the sets by the next update of the level above. This guarantee is made by leveraging trajectory bounding properties of cooperative systems. Under the proposed MPC formulations, recursive feasibility and satisfaction of hard constraints on states is guaranteed. This is achieved without requiring the controllers at the second level to directly exchange information with each other.

The applicability of this approach is demonstrated in both simulation and experimental application using a fluid-thermal system. The thermal management objectives are to achieve high capability of the system to absorb heat loads, while also operating with high efficiency and maintaining all temperature states within specified constraints. The simulation example shows that the theoretical guarantee of recursive feasibility and constraint satisfaction holds. The experimental demonstration extends this theoretical result to real-world application by softening state constraints and extending the hierarchical framework to include additional levels. The experimental results agree closely with the simulation results. While small state constraint violations do occur, they are smaller than under previous control designs applied to the same testbed (e.g., that in Chapter 3 and [64]), despite the fact that no artificial tightening of constraints is performed in this case, as has been done previously.

Chapter 6

Conclusion

6.1 Summary of Research Contributions

As a result of the increasing electrification and complexity of vehicle energy systems, legacy integration and control approaches have become limiting factors of performance and cannot accommodate the requirements of next-generation systems. Addressing this requires control frameworks that coordinate dynamics across multiple physical domains and timescales of these systems, and govern both continuous and switched behavior. This dissertation presents such a framework, with the goal of enabling new paradigms in the capability, efficiency, and safety of vehicle energy systems.

Chapter 2 motivates, derives, and experimentally validates a dynamic graph-based modeling approach for power flow systems. Rooted in conservation of mass and energy, this approach can be applied to capture interactions within and between systems characterized by multiple physical domains and dynamic timescales. Both continuous and switched behaviors of energy systems can also be captured. The governing equations of graph-based models make explicit the underlying structure of coupling within the system under study. This structural knowledge is leveraged throughout the dissertation for analysis and control design.

Chapter 3 presents and experimentally demonstrates a switched hierarchical control framework for vehicle energy management. The candidate system for this demonstration consists of a hardware-in-the-loop experimental testbed that embodies key dynamic characteristics of aircraft electro-thermal systems. As compared to a baseline controller representative of conventional control approaches, the hierarchical controller achieves improved performance in figures of merit quantifying the capability, safety, and efficiency of the closed-loop system.

Chapter 4 presents a decentralized approach to certifying closed-loop passivity of a class of switched, nonlinear graph-based dynamic models. A common storage function is applied to guarantee open-loop passivity of individual switched subsystems. Passivity is then shown to be preserved under the interconnection of multiple graph-based subsystems. Decentralized Model Predictive Controllers paired with each subsystem are formulated with a local passivity-preserving constraint to ensure closed-loop stability. A simulation example demonstrates the efficacy of the proposed approach on a fluid tank system controlled by a hierarchical framework.

Chapter 5 presents a two-level switched hierarchical MPC framework that guarantees satisfaction of state and input constraints for a class of switched graph-based dynamic systems. The proposed framework is robust in the sense that constraint satisfaction and recursive feasibility are guaranteed subject to bounded uncertainty in the exogenous disturbances associated with the ability of the system to dissipate energy to sinks. The applicability of this approach is demonstrated in both simulation and experimental application using a fluid-thermal system. The simulation study confirms that the proposed hierarchical framework enforces hard constraints on states, and shows how the goals of achieving high capability, efficiency, and safety are reflected in the control design. The experimental demonstration bridges from theoretical results to real-world application by softening state constraints and extending the hierarchical framework to include additional levels.

6.2 Future Work

This dissertation represents an initial contribution towards the objectives of achieving high-performing, stable, and robust vehicle energy management. Future work will build on this foundation with advances in both theory and application, as well as the development of supporting design and estimation techniques.

6.2.1 Theory

1. **Passivity-based MPC under discretization:** Analysis of the passivity-based MPC formulation in Chapter 4 is performed using a continuous-time representation of the system dynamics. However, this formulation must be discretized for implementation. Future work should address the implications of this discretization to analysis via explicit consideration of the sampled-data nature of MPC implementation, similar to [137, 138].

2. **Cooperativity-based hierarchical control:** While the cooperativity-based hierarchical control framework in Chapter 5 consists of two levels, its extension to an arbitrary number of levels is a topic of ongoing study. This will greatly extend its applicability to complex systems with dynamic timescales ranging across multiple orders of magnitude. The guarantee of closed-loop state constraint satisfaction in Chapter 5 is made with respect to the discretized dynamics at the rate of the fastest controller of the hierarchy. In many applications, this update rate can be chosen fast enough such that the error from the continuous evolution of the states is negligible. However, the guarantee could be extended to encompass the continuous dynamics as well. In addition, future work will seek to increase the domain of the controller and reduce conservatism by removing the requirement that the system is within an RPI set at each update of the Level 1 controller, instead requiring only that an RPI set can be reached by the end of this controller's horizon. Lastly, while Chapter 5 focuses on state constraint satisfaction rather than stability about an equilibrium, future work could seek to design the stage and terminal costs of the objective function to achieve stability as well. As in [164], this modification would be enabled by the fact that the current guarantee of state constraint satisfaction is achieved without imposing restrictions on objective functions, leaving them available to be used in achieving additional properties. An alternate path to guaranteeing stability could be to impose a dwell time or average dwell time on switching [26].

6.2.2 Application

1. **Hierarchical control of additional physical systems and vehicle classes:** While the hardware-in-the-loop case study in Chapter 3 includes a physical fluid-thermal system, the electrical system is simulated. Therefore, future work should extend the electro-thermal demonstration to feature physical electrical systems as well. Lower levels can be included in the hierarchical control framework to govern faster timescale electrical operation, such as voltage regulation of power electronic devices. Future work should also include experimental demonstration of hierarchical control on testbeds with mechanical components and propulsion systems. While aircraft energy systems are the application of focus in this dissertation, the proposed methods are appropriate for any complex vehicle energy system. Demonstration to other vehicle classes is a topic of ongoing and future research. For example, application of hierarchical control to an automotive electric vehicle energy system has been performed in simulation in [66].

2. **Experimental application of passivity-based MPC:** To bridge the contribution of Chapter 3 from theory to practice, the proposed passivity-based MPC approach should be demonstrated in experimental application on an appropriate testbed, such as the fluid-thermal testbed used throughout this dissertation.

6.2.3 Supporting Tools

1. **Computational considerations:** While the experimental demonstrations in this dissertation exhibit the ability of the proposed control approaches to be implemented in real time, controllers were solved on desktop computers that may have significantly greater computational power than the embedded hardware on board some vehicle energy systems. Therefore, future research should explore tradeoffs between computational cost and performance, as a function of update rates and prediction horizons throughout a hierarchical control framework.

2. **Controller parameterization:** Hierarchical control frameworks can have many more design parameters than centralized controllers. These include the number of levels and

partitioning of the system model within each level, the update rates and prediction horizons of each level, and the design and weightings of objective functions. The development of tools that can aid in selection of these controller parameters will significantly decrease the time required to tune controllers and may improve the performance of the final design. For thermal systems governed within hierarchical frameworks, objective function designs based on analysis of exergy destruction [98] could be particularly promising.

3. **Model reduction:** It may not be necessary to include the full-order system model in the upper levels of a hierarchical control framework. Instead, model reduction can be performed to develop a lower-order representation of the system that captures the most salient behavior at the timescale of a given level and requires much less computational power to optimize than the full-order model. However, care must be taken to ensure that this reduction does not introduce significant model mismatch between levels of the hierarchy that could degrade performance or result in destabilization. Future work will develop tools for performing reduction of the graph-based modeling approach utilized throughout this dissertation, and will leverage these reduced models for hierarchical control.

4. **Hierarchical estimation:** Just as centralized control may not be feasible for many complex vehicle energy systems, communication bandwidth limitations may present barriers to centralized state estimation. As such, recent work has begun to study the pairing of hierarchical control frameworks with hierarchical estimation frameworks [165]. Future work will continue to develop these estimation frameworks, in particular by extending them to achieve accurate estimation under switched behavior of vehicle energy systems.

5. **Fault detection and response:** While this dissertation explores the performance benefits of hierarchical control frameworks when the systems they control operate as intended, it remains to develop tools for detecting and accounting for system faults, sudden changes to the desired vehicle task, and unanticipated environmental condi-

tions. Future research can address this need by augmenting the low-level, fast-updating controllers of a hierarchical framework to detect undesirable operation and trigger appropriate corrective actions throughout the system.

6. **Design optimization:** Recent work has explored the utilization of the graph-based modeling approach for design optimization, where graph-theoretic techniques aid in the programmatic generation and evaluation of large numbers of candidate energy system architectures within a class [33, 93]. Ongoing and future work will perform architecture design and component sizing in tandem with control design, optimizing the system dynamics to best be leveraged in advanced control.

References

[1] M. A. Williams, "A Framework for the Control of Electro-Thermal Aircraft Power Systems," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2017.

[2] W. D. Gerstler and R. S. Bunker, "Aircraft Engine Thermal Management: The Impact of Aviation Electric Power Demands," *ASME Global Gas Turbine News*, 2008.

[3] P. Cazzola, M. Gorner, L. Munuera, R. Schuitmaker, and E. Maroney, "Global EV Outlook 2017: Two Million and Counting," International Energy Agency, Tech. Rep., 2017.

[4] A. R. Gnadt, R. L. Speth, J. S. Sabnis, and S. R. Barrett, "Technical and environmental assessment of all-electric 180-passenger commercial aircraft," *Progress in Aerospace Sciences*, 2018.

[5] F. Baronti, G. Fantechi, R. Roncella, R. Saletti, G. Pede, and F. Vellucci, "Design of the battery management system of LiFePO4 batteries for electric off-road vehicles," in *IEEE International Symposium on Industrial Electronics*, 2013, pp. 1–6.

[6] J. S. Thongam, M. Tarbouchi, a. F. Okou, D. Bouchard, and R. Beguenane, "All-electric ships: A review of the present state of the art," in *Eighth International Conference and Exhibition on Ecological Vehicles and Renewable Energies*, 2013, pp. 1–8.

[7] R. A. Roberts and D. D. Decker, "Energy Optimization of an Aircraft Focused on Component Sizing and Control Architecture Interactions," in *11th International Energy Conversion Engineering Conference*, 2013.

[8] A. Boglietti, A. Cavagnino, A. Tenconi, and S. Vaschetto, "The safety critical electric machines and drives in the more electric aircraft: A survey," in *Industrial Electronics Conference*, 2009, pp. 2587–2594.

[9] K. J. Karimi, "Future Aircraft Power Systems- Integration Challenges," 2007.

[10] M. Dooley, C. Lui, and R. W. Newman, "Efficient Propulsion, Power, and Thermal Management Integration," in *49th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, 2013, pp. 1–8.

[11] J. Doty, K. Yerkes, L. Byrd, J. Murthy, A. Alleyne, M. Wolff, S. Heister, and T. S. Fisher, "Dynamic Thermal Management for Aerospace Technology: Review and Outlook," *Journal of Thermophysics and Heat Transfer*, vol. 31, no. 1, pp. 86–98, 2017.

[12] M. Bodie, G. Russell, K. Mccarthy, E. Lucus, J. Zumberge, and M. Wolff, "Thermal Analysis of an Integrated Aircraft Model," in *48th AIAA Aerospace Sciences Meeting*, 2010, pp. 1–12.

[13] C. Bailey, "Thermal management technologies for electronic packaging: Current capabilities and future challenges for modeling tools," *10th Electronics Packaging Technology Conference*, pp. 527–532, 2008.

[14] S. V. Garimella, A. S. Fleischer, J. Y. Murthy, A. Keshavarzi, R. Prasher, C. Patel, S. H. Bhavnani, R. Venkatasubramanian, R. Mahajan, Y. Joshi, B. Sammakia, B. a. Myers, L. Chorosinski, M. Baelmans, P. Sathyamurthy, and P. E. Raad, "Thermal Challenges in Next-Generation Electronic Systems," *IEEE Transactions on Components and Packaging Technologies*, vol. 31, no. 4, pp. 801–815, 2008.

[15] O. Veneri, Ed., *Technologies and Applications for Smart Charging of Electric and Plug-in Hybrid Vehicles*. Springer, 2017.

[16] M. Boyd, J. Mcnichols, M. Wolff, M. Corbett, and P. Lamm, "Effects of Transient Power Extraction on an Integrated Hardware-in-the-Loop Aircraft/Propulsion/Power System," in *Power Systems Conference*, 2008.

[17] J. Lo, "Effect of Temperature on Lithium-Iron Phosphate Battery Performance and Plug-in Hybrid Electric Vehicle Range," *M.S. Thesis, University of Waterloo*, 2013.

[18] T. Yuksel, S. Litster, V. Viswanathan, and J. J. Michalek, "Plug-in hybrid electric vehicle LiFePO4battery life implications of thermal management, driving conditions, and regional climate," *Journal of Power Sources*, vol. 338, pp. 49–64, 2017.

[19] D. B. Doman, "Fuel Flow Control for Extending Aircraft Thermal Endurance," *Journal of Thermophysics and Heat Transfer*, vol. 32, no. 1, pp. 35–50, 2018.

[20] T. Mahefkey, K. Yerkes, B. Donovan, and M. L. Ramalingam, "Thermal Management Challenges For Future Military Aircraft Power Systems," in *Power Systems Conference*, 2004.

[21] D. J. Burns, C. Danielson, J. Zhou, and S. D. Cairano, "Reconfigurable Model Predictive Control for Multievaporator Vapor Compression Systems," *IEEE Transactions on Automatic Control*, vol. 26, no. 3, pp. 984–1000, 2018.

[22] J. Bendtsen, K. Trangbaek, and J. Stoustrup, "Hierarchical Model Predictive Control for Plug-and-Play Resource Distribution," in *Distributed Decision Making and Control*. London: Springer, 2012, pp. 339–358.

[23] M. Maasoumy, P. Nuzzo, F. Iandola, M. Kamgarpour, A. Sangiovanni-Vincentelli, and C. J. Tomlin, "Optimal Load Management System for Aircraft Electric Power Distribution," in *IEEE Conference on Decision and Control*, 2013, pp. 2939–2945.

[24] A. Barzegar, R. Su, C. Wen, L. Rajabpour, Y. Zhang, A. Gupta, C. Gajanayake, and M. Y. Lee, "Intelligent power allocation and load management of more electric aircraft," *International Conference on Power Electronics and Drive Systems*, pp. 533–538, 2015.

[25] M. Williams, "A Hierarchical Control Strategy For Aircraft Thermal Systems," *M.S. Thesis, University of Illinois at Urbana-Champaign*, 2014.

[26] D. Liberzon, *Switching in Systems and Control*. Boston: Birkhauser, 2003.

[27] T. O'Connell, G. Russell, K. Mccarthy, E. Lucus, J. Zumberge, and M. Wolff, "Energy Management of an Aircraft Electrical System," in *46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, 2010.

[28] R. A. Roberts and S. M. Eastbourn, "Vehicle Level Tip-to-Tail Modeling of an Aircraft," *International Journal of Thermodynamics*, vol. 17, no. 2, pp. 107–115, 2014.

[29] K. L. Yerkes, J. D. Scofield, D. L. Courson, and H. Jiang, "Steady-Periodic Acceleration Effects on the Performance of a Loop Heat Pipe," *Journal of Thermophysics and Heat Transfer*, vol. 28, no. 3, pp. 440–454, 2014.

[30] S. K. Thomas and K. L. Yerkes, "Quasi-Steady-State Performance of a Heat Pipe Subjected to Transient Acceleration Loadings," *Journal of Thermophysics and Heat Transfer*, vol. 11, no. 2, pp. 306–308, 1997.

[31] M. Bodie and M. Wolff, "Robust Optimization of an Aircraft Power Thermal Management System," in *8th Annual International Energy Conversion Engineering Conference*, 2010.

[32] D. Rajaram, Y. Cai, I. Chakraborty, and D. N. Mavris, "Integrated Sizing and Optimization of Aircraft and Subsystem Architectures in Early Design," *Journal of Aircraft*, vol. 55, no. 5, 2018.

[33] Satya R. T. Peddada, Daniel R. Herber, Herschel C. Pangborn, Andrew G. Alleyne, and James T. Allison, "Optimal Flow Control and Single Split Architecture Exploration for Fluid-Based Thermal Management," *Journal of Mechanical Design*, 2019.

[34] C. Lui and M. Dooley, "Electric Thermal Management Architectures," *SAE Technical Paper 2013-01-2164*, 2013.

[35] D. B. Doman, "Rapid Mission Planning for Aircraft Thermal Management," in *AIAA Guidance, Navigation, and Control Conference*, 2015, pp. 1–19.

[36] B. Rasmussen, "Dynamic Modeling and Advanced Control of Air Conditioning and Refrigeration Systems," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2005.

[37] M. Kania, J. Koeln, and A. Alleyne, "A Dynamic Modeling Toolbox for Air Vehicle Vapor Cycle Systems," in *Power Systems Conference*, 2012.

[38] K. Mccarthy, E. Walters, A. Heltzel, P. C. Krause, and J. Dalton, "Dynamic Thermal Management System Modeling of a More Electric Aircraft," in *Power Systems Conference*, 2008.

[39] M. Williams, S. Sridharan, S. Banerjee, C. Mak, C. Pauga, P. Krein, A. Alleyne, A. Jacobi, and S. D. Urso, "PowerFlow: A Toolbox for Modeling and Simulation of Aircraft Systems," *SAE Technical Paper 2015-01-2417*, 2015.

[40] R. Scattolini, "Architectures for Distributed and Hierarchical Model Predictive Control - A Review," *Journal of Process Control*, vol. 19, pp. 723–731, 2009.

[41] A. R. Behbahani, A. Von Moll, R. Zeller, and J. Ordo, "Aircraft Integration Challenges and Opportunities for Distributed Intelligent Control, Power, Thermal Management, and Diagnostic and Prognostic Systems," in *Aerospace Systems and Technology Conference*, 2014.

[42] T. Michalak, S. Emo, and J. Ervin, "Control Strategy for Aircraft Vapor Compression System Operation," *International Journal of Refrigeration*, vol. 48, pp. 10–18, 2014.

[43] H. Pangborn, J. E. Hey, T. O. Deppen, A. G. Alleyne, and T. S. Fisher, "Hardware-in-the-Loop Validation of Advanced Fuel Thermal Management Control," in *46th AIAA Thermophysics Conference*, 2016, pp. 1–8.

[44] M. R. Amini, J. Sun, and I. Kolmanovsky, "Two-Layer Model Predictive Battery Thermal and Energy Management Optimization for Connected and Automated Electric Vehicles," in *IEEE Conference on Decision and Control*, 2018.

[45] T. O. Deppen, A. G. Alleyne, K. A. Stelson, and J. J. Meyer, "Optimal Energy Use in a Light Weight Hydraulic Hybrid Passenger Vehicle," *Journal of Dynamic Systems, Measurement, and Control*, vol. 134, pp. 1–11, 2012.

[46] H. Borhan, A. Vahidi, A. M. Phillips, M. L. Kuang, I. V. Kolmanovsky, and S. Di Cairano, "MPC-Based Energy Management of a Power-Split Hybrid Electric Vehicle," *IEEE Transactions on Control Systems Technology*, vol. 20, no. 3, pp. 593–603, 2012.

[47] N. Jain and A. Alleyne, "Exergy-based optimal control of a vapor compression system," *Energy Conversion and Management*, vol. 92, pp. 353–365, 2015.

[48] H. Pangborn, J. E. Hey, T. O. Deppen, A. G. Alleyne, and T. S. Fisher, "Hardware-in-the-Loop Validation of Advanced Fuel Thermal Management Control," *Journal of Thermophysics and Heat Transfer*, pp. 1–9, 2017.

[49] T. O. Deppen, J. E. Hey, A. G. Alleyne, and T. S. Fisher, "A Model Predictive Framework for Thermal Management of Aircraft," in *ASME Dynamic Systems and Control Conference*, 2015, pp. 1–9.

[50] J. Seok, I. Kolmanovsky, and A. Girard, "Integrated/coordinated control of aircraft gas turbine engine and electrical power system: Towards large electrical load handling," in *IEEE Conference on Decision and Control*, 2016, pp. 3183–3189.

[51] ——, "Coordinated Model Predictive Control of Aircraft Gas Turbine Engine and Power System," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 10, pp. 1–18, 2017.

[52] J. Seok, D. M. Reed, I. V. Kolmanovsky, and A. R. Girard, "Coordinated Model Predictive Control of Aircraft Gas Turbine Engine with Simplified Electrical System Model," in *American Control Conference*, 2018, pp. 1460–1466.

[53] J. Hou, J. Sun, and H. Hofmann, "Integrated Control of Power Generation, Electric Motor and Hybrid Energy Storage for All-Electric Ships," in *American Control Conference*, 2016, pp. 6797–6802.

[54] V. Chandan, "Decentralized Thermal Control of Building Systems," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2013.

[55] R. Negenborn, A. Sahin, Z. Lukszo, B. De Schutter, and M. Morarri, "A Non-Iterative Cascaded Predictive Control Approach for Control of Irrigation Canals," in *IEEE International Conference on Systems, Man and Cybernetics*, 2009, pp. 3552–3557.

[56] C. Ocampo-Martinez, D. Barcelli, V. Puig, and A. Bemporad, "Hierarchical and Decentralised Model Predictive Control of Drinking Water Networks: Application to Barcelona Case Study," *IET Control Theory and Applications*, vol. 6, no. 1, pp. 62–71, 2012.

[57] M. J. Tippett and J. Bao, "Distributed Model Predictive Control Based on Dissipativity," *AIChE Journal*, vol. 59, no. 3, pp. 787–804, 2012.

[58] S. Rahnama, J. D. Bendtsen, J. Stoustrup, S. Member, and H. Rasmussen, "Robust Aggregator Design for Industrial Thermal Energy Storages in Smart Grid," *IEEE Transactions on Smart Grid*, vol. 8, no. 2, pp. 902–916, 2017.

[59] Y. Zheng, S. Li, and R. Tan, "Distributed Model Predictive Control for On-Connected Microgrid Power Management," *IEEE Transactions on Control Systems Technology*, vol. 26, no. 3, pp. 1028–1039, 2018.

[60] X. Zhang, J. Bao, R. Wang, C. Zheng, and M. Skyllas-Kazacos, "Dissipativity based distributed economic model predictive control for residential microgrids with renewable energy generation and battery energy storage," *Renewable Energy*, vol. 100, pp. 18–34, 2017.

[61] W. Dunham, B. Hencey, A. Girard, and I. Kolmanovsky, "Distributed MPC via ADMM for Coordination and Control of More Electric Aircraft Subsystems," in *ASME Dynamic Systems and Control Conference*, 2017.

[62] J. M. Maestre and R. R. Negenborn, Eds., *Distributed Model Predictive Control Made Easy*. Springer, 2014, vol. 69.

[63] J. P. Koeln, "Hierarchical Power Management in Vehicle Systems," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2016.

[64] H. C. Pangborn, J. P. Koeln, M. A. Williams, and A. G. Alleyne, "Experimental Validation of Graph-Based Hierarchical Control for Thermal Management," *Journal of Dynamic Systems, Measurement, and Control*, vol. 140, pp. 1–16, 2018.

[65] M. A. Williams, J. P. Koeln, and A. G. Alleyne, "Hierarchical Control of Multi-Domain Power Flow in Mobile Systems - Part II: Aircraft Application," in *ASME Dynamic Systems and Control Conference*, 2015, pp. 1–10.

[66] D. J. Docimo, H. C. Pangborn, and A. G. Alleyne, "Hierarchical Control for Electro-Thermal Power Management of an Electric Vehicle Powertrain," in *ASME Dynamic Systems and Control Conference*, 2018, pp. 1–10.

[67] A. Puntel, S. Emo, T. E. Michalak, J. Ervin, L. Byrd, V. Tsao, and T. Reitz, "Refrigerant Charge Management and Control for Next-Generation Aircraft Vapor Compression Systems," *SAE Technical Paper 2013-01-2241*, 2013.

[68] T. O. Deppen, "Optimal energy use in mobile applications with storage," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2013.

[69] J. P. Koeln, M. A. Williams, H. C. Pangborn, and A. G. Alleyne, "Experimental Validation of Graph-Based Modeling for Thermal Fluid Power Flow Systems," in *ASME Dynamic Systems and Control Conference*, 2016, pp. 1–10.

[70] S. K. Srivastava, D. A. Cartes, F. Maturana, F. Ferrese, M. Pekala, M. Zink, R. Meeker, D. Carnahan, R. Staron, D. Scheidt, and K. Huang, "A Control System Test Bed for Demonstration of Distributed Computational Intelligence Applied to Reconfiguring Heterogeneous Systems," *IEEE Instrumentation & Measurement Magazine*, vol. 11, no. 1, 2008.

[71] S. Yang, M. Chagas, and J. Ordonez, "Modeling, cross-validation, and optimization of a shipboard integrated energy system cooling network," *Applied Thermal Engineering*, vol. 145, pp. 516–527, 2018.

[72] H. Pangborn, A. G. Alleyne, and N. Wu, "A Comparison Between Finite Volume and Switched Moving Boundary Approaches for Dynamic Vapor Compression System Modeling," *International Journal of Refrigeration*, vol. 53, pp. 101–114, 5 2015.

[73] J. T. Wen and S. Mishra, Eds., *Advances in Industrial Control Intelligent Building Control Systems*. Springer, 2018.

[74] M. Cantoni, E. Weyer, Y. Li, S. K. Ooi, I. Mareels, and M. Ryan, "Control of Large-Scale Irrigation Networks," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 75–91, 2007.

[75] P. D. Christofides, R. Scattolini, D. Muñoz de la Peña, and J. Liu, "Distributed Model Predictive Control: A Tutorial Review and Future Research Directions," *Computers and Chemical Engineering*, vol. 51, pp. 21–41, 2013.

[76] R. Zamora and A. K. Srivastava, "Controls for Microgrids with Storage: Review, Challenges, and Research Needs," *Renewable and Sustainable Energy Reviews*, vol. 14, no. 7, pp. 2009–2018, 2010.

[77] M. A. Williams, J. P. Koeln, H. C. Pangborn, and A. G. Alleyne, "Dynamical Graph Models of Aircraft Electrical, Thermal, and Turbomachinery Components," *Journal of Dynamic Systems, Measurement, and Control*, vol. 140, no. 4, 2018.

[78] D. C. Karnopp, D. L. Margolis, and R. C. Rosenberg, *System Dynamics: Modeling, Simulation, and Control of Mechatronic Systems*. Hoboken, NJ, USA: John Wiley & Sons, 2012.

[79] T. Ersal, H. K. Fathy, and J. L. Stein, "Structural simplification of modular bond-graph models based on junction inactivity," *Simulation Modelling Practice and Theory*, vol. 17, no. 1, pp. 175–196, 2009.

[80] S. S. Jogwar, S. Rangarajan, and P. Daoutidis, "Multi-time scale dynamics in energy-integrated networks: A graph theoretic analysis," in *18th IFAC World Congress*, 2011, pp. 6085–6090.

[81] S. Heo, S. Rangarajan, P. Daoutidis, and S. S. Jogwar, "Graph Reduction of Complex Energy-Integrated Networks: Process Systems Applications," *AIChE Journal*, vol. 60, no. 3, pp. 995–1012, 2014.

[82] H. A. Preisig, "A Graph-Theory-Based Approach to the Analysis of Large-Scale Plants," *Computers and Chemical Engineering*, vol. 33, no. 3, pp. 598–604, 2009.

[83] K. L. Moore, T. L. Vincent, F. Lashhab, and C. Liu, "Dynamic Consensus Networks with Application to the Analysis of Building Thermal Processes," in *18th IFAC World Congress*, 2011, pp. 3078–3083.

[84] S. Mukherjee, S. Mishra, and J. T. Wen, "Building temperature control: A passivity-based approach," in *IEEE Conference on Decision and Control*, 2012, pp. 6902–6907.

[85] F. Blanchini, E. Franco, G. Giordano, V. Mardanlou, and P. L. Montessoro, "Compartmental flow control: Decentralization, robustness and optimality," *Automatica*, vol. 64, pp. 18–28, 2016.

[86] H. Behjati, A. Davoudi, and F. Lewis, "Modular DC-DC converters on graphs: Cooperative control," *IEEE Transactions on Power Electronics*, vol. 29, no. 12, pp. 6725–6741, 2014.

[87] D. B. West, *Introduction to Graph Theory*, 2nd ed. London: Pearson, 2001.

[88] P. J. Tannous, S. R. Peddada, J. T. Allison, T. Foulkes, R. C. Pilawa-Podgurski, and A. G. Alleyne, "Dynamic Temperature Estimation of Power Electronics Systems," in *American Control Conference*, 2017, pp. 3463–3469.

[89] C. T. Aksland, T. W. Bixel, L. C. Raymond, M. A. Rottmayer, and A. G. Alleyne, "Graph-Based Modeling and Validation of Electrical and Mechanical Systems for a Hybrid Unmanned Aerial Vehicle," in *American Control Conference*, 2019.

[90] V. Gnielinski, "New Equations for Heat and Mass-transfer in Turbulent Pipe and Channel Flow," *International Chemical Engineering*, vol. 16, no. 1, pp. 8–16, 1976.

[91] G. A. Longo and A. Gasparella, "Refrigerant R134a vaporisation heat transfer and pressure drop inside a small brazed plate heat exchanger," *International Journal of Refrigeration*, vol. 30, no. 5, pp. 821–830, 2007.

[92] D. J. Docimo and A. G. Alleyne, "Electro-Thermal Graph-Based Modeling for Hierarchical Control with Application to an Electric Vehicle," in *IEEE Conference on Control Technology and Applications*, 2018.

[93] S. Peddada, D. Herber, H. Pangborn, A. A.G., and A. J.T., "Optimal Flow Control and Single Split Architecture Exploration for Fluid-Based Thermal Management," in *ASME 2018 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference*, 2018, pp. 1–11.

[94] J. Riordan, "The enumeration of trees by height and diameter," *IBM Journal of Research and Development*, vol. 4, no. 5, pp. 473–478, 1960.

[95] R. P. Stanley, *Enumerative Combinatorics, Volume 1*, 2nd ed. Cambridge University Press, 2011.

[96] C. T. Aksland, J. P. Koeln, and A. G. Alleyne, "A Graph-Based Approach For Dynamic Compressor Modeling in Vapor Compression Systems," in *Dynamic Systems and Control Conference*, 2017, pp. 1–8.

[97] W. Dunham, B. Hencey, I. Kolmanovsky, and A. Girard, "Predictive propulsion and power control for large transient power loads in a More Electric Aircraft," in *American Control Conference*, 2017, pp. 4055–4061.

[98] N. Jain and B. M. Hencey, "Increasing Fuel Thermal Management System Capability via Objective Function Design," in *American Control Conference*, 2016.

[99] D. Schlabe and J. Lienig, "Energy Management of Aircraft Electrical Systems - State of the Art and Further Directions," in *Electrical Systems for Aircraft, Railway and Ship Propulsion*, 2012.

[100] N. Ozay, U. Topcu, and R. M. Murray, "Distributed power allocation for vehicle management systems," in *IEEE Conference on Decision and Control and European Control Conference*, 2011, pp. 4841–4848.

[101] J. P. Koeln, H. C. Pangborn, M. A. Williams, M. L. Kawamura, and A. G. Alleyne, "Hierarchical Control of Aircraft Electro-Thermal Systems," *IEEE Transactions on Control Systems Technology*, 2019.

[102] P. Falcone, F. Borrelli, H. E. Tseng, J. Asgari, and D. Hrovat, "Linear time-varying model predictive control and its application to active steering systems: Stability analysis and experimental validation," *International Journal of Robust and Nonlinear Control*, vol. 18, pp. 862–875, 2008.

[103] R. Cagienard, P. Grieder, E. C. Kerrigan, and M. Morari, "Move blocking strategies in receding horizon control," *Journal of Process Control*, vol. 17, no. 6, pp. 563–570, 2007.

[104] J. Lofberg, "YALMIP: A Toolbox for Modeling and Optimization in MATLAB," in *IEEE International Symposium on Computer Aided Control Systems Design*, 2004, pp. 284–289.

[105] Gurobi Optimization Inc., "Gurobi Optimizer Reference Manual," 2016.

[106] J. C. Willems, "Dissipative Dynamical Systems Part I: General Theory," *Archive for Rational Mechanics and Analysis*, vol. 45, no. 5, pp. 321–351, 1972.

[107] D. Hill and P. Moylan, "The Stability of Nonlinear Dissipative Systems," *IEEE Transactions on Automatic Control*, vol. 21, no. 5, pp. 708–711, 1976.

[108] M. Arcak, C. Meissen, and A. Packard, *Networks of Dissipative Systems: Compositional Certification of Stability, Performance, and Safety.* Springer, 2016.

[109] R. Sepulchre, M. Jankovic, and P. Kokotovic, *Constructive Nonlinear Control.* London: Springer, 1997.

[110] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River: Prentice Hall, 2002.

[111] A. v. d. Schaft, *L2-Gain and Passivity Techniques in Nonlinear Control.* Springer, 2000.

[112] C. C. Okaeme, S. Mishra, and J. T. Wen, "A Comfort Zone Set-Based Approach for Coupled Temperature and Humidity Control in Buildings," in *IEEE Conference of Automation Science and Engineering*, 2016.

[113] J. Bao and P. L. Lee, *Process Control: The Passive Systems Approach.* Springer, 2010.

[114] A. Ulbig, "Passivity-based Nonlinear Model Predictive Control," Ph.D. dissertation, University of Stuttgart, 2007.

[115] J. Liu, B. Krogh, and B. Ydstie, "Passivity-based robust control for power systems subject to wind power variability," in *American Control Conference*, 2011, pp. 4149–4154.

[116] M. Arcak and E. D. Sontag, "A passivity-based stability criterion for a class of interconnected systems and applications to biochemical reaction networks," *Mathematical Biosciences and Engineering*, 2008.

[117] M. Spong, J. Holm, and D. Lee, "Passivity-Based Control of Bipedal Locomotion," *IEEE Robotics & Automation Magazine*, vol. 14, no. 2, pp. 30–40, 2007.

[118] P. Falugi, "Model Predictive Control: A Passive Scheme," in *19th IFAC World Congress*, 2014, pp. 1017–1022.

[119] C. Løvaas, M. M. Seron, and G. C. Goodwin, "A dissipativity approach to robustness in constrained model predictive control," *IEEE Conference on Decision and Control*, pp. 1180–1185, 2007.

[120] T. Raff, C. Ebenbauer, and F. Allgower, "Nonlinear Model Predictive Control: A Passivity-Based Approach," in *Assessment and Future Directions of Nonlinear Model Predictive Control*, 2007, pp. 151–162.

[121] S. Sredojev and R. Eaton, "Model Predictive Controller for a Class of Nonlinear Dissipative Systems," in *American Control Conference*, 2014, pp. 5570–5573.

[122] H. Yu, F. Zhu, M. Xia, and P. J. Antsaklis, "Robust Stabilizing Output Feedback Nonlinear Model Predictive Control by Using Passivity and Dissipativity," in *European Control Conference*, 2013, pp. 2050–2055.

[123] P. Varutti, B. Kern, and R. Findeisen, "Dissipativity-Based Distributed Nonlinear Predictive Control for Cascaded Systems," in *IFAC Symposium on Advanced Control of Chemical Processes*, 2012, pp. 439–444.

[124] H. Yu and P. J. Antsaklis, "A Passivity Measure of Systems in Cascade Based on Passivity Indices," in *Conference on Decision and Control*, 2010, pp. 2186–2191.

[125] J. P. Koeln and A. G. Alleyne, "Stability of Decentralized Model Predictive Control of Graph-Based Power Flow Systems via Passivity," *Automatica*, vol. 82, pp. 29–34, 2017.

[126] W. Haddad and V. Chellaboina, "Dissipativity Theory and Stability of Feedback Interconnections for Hybrid Dynamical Systems," in *American Control Conference*, 2000, pp. 2688–2694.

[127] N. Chopra and M. W. Spong, "Passivity-Based Control of Multi-Agent Systems," in *Advances in Robot Control: From Everyday Physics to Human-Like Movements*. Springer, 2006, pp. 107–134.

[128] W. Chen and M. Saif, "Passivity and Passivity Based Controller Design of a Class of Switched Control Systems," in *Triennial World Congress of IFAC*, 2005, pp. 676–681.

[129] M. Zefran, F. Bullo, and M. Stein, "A Notion of Passivity for Hybrid Systems," in *IEEE Conference on Decision and Control*, 2001, pp. 768–773.

[130] J. Zhao and D. J. Hill, "Dissipativity Theory for Switched Systems," *IEEE Transactions on Automatic Control*, vol. 53, no. 4, pp. 941–953, 2008.

[131] J. Zhao, "A Notion of Passivity for Switched Systems with State-Dependent Switching," *Journal of Control Theory and Applications*, vol. 1, pp. 70–75, 2006.

[132] M. J. McCourt and P. J. Antsaklis, "Stability of networked passive switched systems," in *IEEE Conference on Decision and Control*, 2010, pp. 1263–1268.

[133] J. Zhao and D. J. Hill, "Completeness, Passivity and Stability of Switched Systems," in *27th Chinese Control Conference*, 2008, pp. 618–622.

[134] H. C. Pangborn, J. P. Koeln, and A. G. Alleyne, "Passivity and Decentralized MPC of Switched Graph-Based Power Flow Systems," in *American Control Conference*, 2018, pp. 6658–6665.

[135] S. Su and Z. Lin, "Distributed Consensus Control of Multi-Agent Systems with Higher Order Agent Dynamics and Dynamically Changing Directed Interaction Topologies," *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 515–519, 2016.

[136] A. Bemporad, G. Bianchini, and F. Brogi, "Passivity Analysis and Passification of Discrete-Time Hybrid Systems," *IEEE Transactions on Automatic Control*, vol. 53, no. 4, pp. 1004–1009, 2008.

[137] L. Magni and R. Scattolini, "Model Predictive Control of Continuous-Time Nonlinear Systems with Piecewise Constant Control," *IEEE Transactions on Automatic Control*, vol. 49, no. 6, pp. 900–906, 2004.

[138] D. Nešić and L. Grüne, "A Receding Horizon Control Approach to Sampled-Data Implementation of Continuous-Time Controllers," *Systems and Control Letters*, vol. 55, no. 8, pp. 660–672, 2006.

[139] J. P. Koeln, M. A. Williams, and A. G. Alleyne, "Hierarchical Control of Multi-Domain Power Flow in Mobile Systems - Part I: Framework Development and Demonstration," in *ASME Dynamic Systems and Control Conference*, 2015.

[140] A. Wächter, "An Interior Point Algorithm for Large-Scale Nonlinear Optimization with Applications in Process Engineering," Ph.D. dissertation, Carnegie Mellon University, 2002.

[141] T. Tran and J. Bao, "Supervisory Stability Assurance Layer for Hierarchical Plant-wide Process Control," in *American Control Conference*, 2010, pp. 4409–4414.

[142] J. D. Engerer and T. S. Fisher, "Flash boiling from carbon foams for high-heat-flux transient cooling," *Applied Physics Letters*, vol. 109, no. 2, 2016.

[143] H. C. Pangborn and A. G. Alleyne, "Cooperativity and Hierarchical MPC of State-Constrained Switched Power Flow Systems," in *American Control Conference*, 2019.

[144] H. L. Smith, *Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems*. American Mathematical Society, 1996.

[145] D. Angeli and E. D. Sontag, "Monotone Control Systems," *IEEE Transactions on Automatic Control*, vol. 48, no. 10, pp. 1684–1698, 2003.

[146] E. D. Sontag, "Monotone and near-monotone biochemical networks," *Systems and Synthetic Biology*, vol. 1, no. 2, pp. 59–87, 2007.

[147] S. Coogan, M. Arcak, and A. A. Kurzhanskiy, "Mixed monotonicity of partial first-in-first-out traffic flow models," in *IEEE Conference on Decision and Control*, 2016, pp. 7611–7616.

[148] P. J. Meyer, H. Nazarpour, A. Girard, and E. Witrant, "Experimental implementation of UFAD regulation based on robust controlled invariance," in *European Control Conference*, 2014, pp. 1468–1473.

[149] P. J. Meyer, A. Girard, and E. Witrant, "Safety control with performance guarantees of cooperative systems using compositional abstractions," in *5th IFAC Conference on Analysis and Design of Hybrid Systems*, 2015.

[150] F. Blanchini and S. Miani, *Set-Theoretic Methods in Control*, 2nd ed. Birkhauser, 2008.

[151] A. Rantzer, "Separable Lyapunov functions for monotone systems," in *IEEE Conference on Decision and Control*, 2013.

[152] H. Ito, B. S. Ruffer, and A. Rantzer, "Max- and sum-separable Lyapunov functions for monotone systems and their level sets," in *IEEE Conference on Decision and Control*, 2014, pp. 2371–2377.

[153] G. Dirr, H. Ito, A. Rantzer, and B. S. Rüffer, "Separable Lyapunov functions for monotone systems: Constructions and limitations," *Discrete and Continuous Dynamical Systems - Series B*, vol. 20, no. 8, pp. 2497–2526, 2015.

[154] S. Coogan, "Separability of Lyapunov functions for contractive monotone systems," in *IEEE Conference on Decision and Control*, 2016, pp. 2184–2189.

[155] A. Rantzer, "Distributed Control of Positive Systems," *European Journal of Control*, vol. 24, pp. 72–80, 2015.

[156] J. P. Koeln and A. G. Alleyne, "Robust hierarchical model predictive control of graph-based power flow systems," *Automatica*, vol. 96, pp. 127–133, 2018.

[157] E. D. Sontag, "Nonlinear Regulation: The Piecewise Linear Approach," *IEEE Transactions on Automatic Control*, vol. 26, no. 2, pp. 346–358, 1981.

[158] M. Lazar, "Model Predictive Control of Hybrid Systems: Stability and Robustness," Ph.D. dissertation, Eindhoven University of Technology, 2006.

[159] F. Blanchini, P. Colaneri, and M. E. Valcher, *Switched Positive Linear Systems*. Now Publishers, 2015.

[160] A. Abate, A. Tiwari, and S. Sastry, "Box invariance in biologically-inspired dynamical systems," *Automatica*, vol. 45, no. 7, pp. 1601–1610, 2009.

[161] P. J. Meyer, A. Girard, and E. Witrant, "Robust controlled invariance for monotone systems: Application to ventilation regulation in buildings," *Automatica*, vol. 70, pp. 14–20, 2016.

[162] P. Trodden, "A One-Step Approach to Computing a Polytopic Robust Positively Invariant Set," *IEEE Transactions on Automatic Control*, vol. 61, no. 12, pp. 4100–4105, 2016.

[163] "IBM ILOG CPLEX Optimization Studio." [Online]. Available: https://www.ibm.com/products/ilog-cplex-optimization-studio

[164] C. Danielson, L. Bridgeman, and S. Di Cairano, "Constraint satisfaction for switched linear systems with restricted dwell-time," in *American Control Conference*, 2017, pp. 3682–3687.

[165] P. J. Tannous, D. J. Docimo, H. C. Pangborn, and A. G. Alleyne, "Hierarchical Estimation for Complex Multi-Domain Dynamical Systems," in *American Control Conference*, 2019.