

Herschel C. Pangborn¹

Mechanical Science and Engineering
Department,
University of Illinois at Urbana-Champaign,
Urbana, IL 61801
e-mail: pangbor2@illinois.edu

Justin P. Koeln

Department of Mechanical Engineering,
University of Texas at Dallas,
Richardson, TX 75080
e-mail: justin.koeln@utdallas.edu

Matthew A. Williams

Northrop Grumman Corporation,
Falls Church, VA 22042
e-mail: matt.a.williams@ngc.com

Andrew G. Alleyne

Mechanical Science and Engineering
Department,
University of Illinois at Urbana-Champaign,
Urbana, IL 61801
e-mail: alleyne@illinois.edu

Experimental Validation of Graph-Based Hierarchical Control for Thermal Management

This paper proposes and experimentally validates a hierarchical control framework for fluid flow systems performing thermal management in mobile energy platforms. A graph-based modeling approach derived from the conservation of mass and energy inherently captures coupling within and between physical domains. Hydrodynamic and thermodynamic graph-based models are experimentally validated on a thermal-fluid testbed. A scalable hierarchical control framework using the graph-based models with model predictive control (MPC) is proposed to manage the multidomain and multi-timescale dynamics of thermal management systems. The proposed hierarchical control framework is compared to decentralized and centralized benchmark controllers and found to maintain temperature bounds better while using less electrical energy for actuation.

[DOI: 10.1115/1.4040211]

1 Introduction

Modern aircraft and large ground vehicles are complex machines consisting of interconnected subsystems that interact in multiple energy domains with a variety of dynamic timescales. A longstanding trend of these systems is increasing electrical power requirements, necessarily accompanied by increased thermal loading [1,2]. This additional burden on thermal management systems is compounded by the continual desire to reduce the size and weight of mobile platforms, and further exacerbated in aviation by a decreased ability to remove thermal energy due to the use of composite skin materials with high thermal resistance and reduction in ram air heat exchanger cross sections [3]. Many approaches have been explored to meet this challenge, with a focus on employing transient analyses to improve upon traditional steady-state analyses [4]. These approaches include integrated aircraft modeling [3], characterization of component performance under transient operation [5,6], design optimization [7], improved topologies for thermal management [8], enhanced mission planning [9], and advanced control approaches [10–13].

Due to the inherent complexity and multidomain nature of thermal management systems, modular modeling frameworks have been developed to support system and control design. A graph-based approach to modeling not only provides this modularity and scalability but can also readily facilitate model-based control design, as shown in Ref. [14] for building thermal systems, [15] for process systems, and [12] and [16] for mobile thermal energy management systems.

To validate both modeling and control approaches for energy systems, experimental testbeds have been developed across a range of application areas. Examples include the vapor compression refrigeration testbeds of Refs. [17] and [18], the hydraulic hybrid vehicle testbed of Ref. [19], the aircraft fuel thermal

management system (FTMS) testbed of [12], and the shipboard chilled water distribution system testbed of Ref. [20].

This paper proposes a hierarchical framework to address the challenge of controlling the multiple timescales and energy domains present in fluid-based thermal management systems. Hierarchical control approaches have been widely investigated in the literature, including [21–25]. While much of this work focuses on deriving analytical guarantees of stability and/or robustness under appropriate assumptions, practical demonstration of modeling and hierarchical control frameworks in application to physical systems is not readily found in the literature. Therefore, the purpose of this paper is to provide such a demonstration. The graph-based model validation of Ref. [26] is first extended to a more complex testbed configuration. Then, this paper builds upon the initial simulation study of Ref. [13] by experimentally demonstrating the hierarchical control approach and comparing this to several baseline control approaches. In comparison to these benchmarks, the hierarchical controller is found to maintain temperature bounds better while using less energy for actuation.

The remainder of this paper is organized as follows. Section 2 presents the graph-based modeling framework. Model validation on an experimental testbed is shown in Sec. 3. Section 4 presents the proposed hierarchical control framework, and Sec. 5 discusses the baseline decentralized and centralized control approaches. Section 6 experimentally demonstrates and compares the three control approaches. Conclusions are provided in Sec. 7.

Notation: A vector x of elements x_i is denoted by $x = [x_i]$, while $M = [m_{ij}]$ denotes a matrix M of elements m_{ij} . Lower case superscripts are used throughout this paper in the naming of variables, while upper case superscripts indicate mathematical functions, such as a transpose.

2 Graph-Based Dynamic Modeling

2.1 Generic Graph Formulation. The graph-based models of this paper are derived by applying conservation equations to a component or system, inherently capturing the storage and transport of mass or energy. When modeling a system as a graph, capacitive elements that store energy, or mass, are represented as

¹Corresponding author.

Contributed by the Dynamic Systems Division of ASME for publication in the JOURNAL OF DYNAMIC SYSTEMS, MEASUREMENT, AND CONTROL. Manuscript received September 11, 2017; final manuscript received May 1, 2018; published online June 4, 2018. Assoc. Editor: Mahdi Shahbakhhti.

vertices and the paths for the transport of energy, or mass, between storage elements are represented as edges. A key feature of graph-based models that makes them attractive for model-based hierarchical control is that a graph-based model of a complete system can easily be partitioned into separate models for individual subsystems by clustering the graph into subgraphs based on an analysis of the edges and vertices. Edges that are cut because of this partitioning represent coupling terms between the subsystems, for which controllers of the hierarchy can account by exchanging information. This ease of partitioning reduces the labor required to generate the many models used throughout a model-based hierarchical control framework.

Let the oriented graph $G = (v, e)$ represent the structure of storage and exchange of a conserved quantity throughout a system S . Graph G is said to be of order N_v with vertices $v = [v_i], i \in [1, N_v]$, and of size N_e with edges $e = [e_j], j \in [1, N_e]$. As shown in the notional graph example of Fig. 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 $e_i^{\text{head}} = \{e_j | v_j^{\text{head}} = v_i\}$, while the set of edges directed out of vertex v_i is given by $e_i^{\text{tail}} = \{e_j | v_j^{\text{tail}} = v_i\}$ [27].

Each vertex v_i of graph G has an associated dynamic state in system S , denoted as x_i , representing an amount of stored energy or mass. Similarly, each edge e_j has an associated value y_j describing the transfer rate of energy (i.e., power flow) or mass between adjacent vertices. The orientation of each edge indicates the direction of positive flow, from v_j^{tail} to v_j^{head} . Therefore, the dynamics of each vertex satisfy the conservation equation

$$C_i \dot{x}_i = \sum_{\{j | e_j \in e_i^{\text{head}}\}} y_j - \sum_{\{j | e_j \in e_i^{\text{tail}}\}} y_j \quad (1)$$

where C_i is the storage capacitance of the i th vertex.

The transfer rate y_j along each edge is a function of the states of the vertices to which it is incident and may also be a function of an input signal u_j . The transfer rate along each edge is, therefore, given by

$$y_j = f_j(x_j^{\text{tail}}, x_j^{\text{head}}, u_j) \quad (2)$$

Figure 1 includes examples of Eqs. (1) and (2) as applied to several vertices and edges.

In addition to capturing the exchange of energy or mass within the graph, the modeling framework must account for exchange with entities external to the graph. Sources to graph G are modeled by source edges $e^{\text{in}} = [e_j^{\text{in}}], j \in [1, N_{\text{in}}]$ with associated power flows $y^{\text{in}} = [y_j^{\text{in}}]$, which are treated as disturbances to the system that may come from neighboring systems or the environment. Therefore, edges belonging to e^{in} are not counted among the edges e of graph G , and transfer rates in y^{in} are not counted among the internal transfer rates y of system S .




Fig. 1 Notional graph example to demonstrate key features of the modeling framework

Sinks of graph G are modeled by sink vertices $v^{\text{out}} = [v_j^{\text{out}}], j \in [1, N_{\text{out}}]$ with associated states $x^{\text{out}} = [x_j^{\text{out}}]$. The sink vertices are counted among the vertices v of graph G but the sink states x^{out} are not included in the state vector x of system S . Instead, the sink states x^{out} are treated as disturbances to the system associated with neighboring systems or the environment.

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

$$m_{ij} = \begin{cases} +1 & v_i \text{ is the tail of } e_j \\ -1 & v_i \text{ is the head of } e_j \\ 0 & \text{else} \end{cases} \quad (3)$$

M can then be partitioned as

$$M = \begin{bmatrix} \bar{M} \\ \underline{M} \end{bmatrix} \text{ with } \bar{M} \in \mathbb{R}^{(N_v - N_{\text{out}}) \times N_e} \quad (4)$$

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

Similarly, the connection of external sources to the system is given by $D = [d_{ij}] \in \mathbb{R}^{(N_v - N_{\text{out}}) \times N_{\text{in}}}$ where

$$d_{ij} = \begin{cases} 1 & v_i \text{ is the head of } e_j^{\text{in}} \\ 0 & \text{else} \end{cases} \quad (5)$$

Following from the conservation equation for each vertex (1) and the above definitions of \bar{M} and D , the dynamics of all states in system S are given by

$$S: C \dot{x} = -\bar{M}y + Dy^{\text{in}} \quad (6)$$

where $C = \text{diag}([C_i])$ is the diagonal matrix of capacitances.

Following from Eq. (2), the vector of all power flows y in S is given by

$$y = F(x, x^{\text{out}}, u) = \left[f_j(x_j^{\text{tail}}, x_j^{\text{head}}, u_j) \right] \quad (7)$$

2.2 Domain-Specific Modeling. An exposition of how conservation of mass and energy can be applied to derive lumped parameter graph-based models for a variety of fluid-thermal components is found in Ref. [16]. These components include a fluid reservoir, a flow split/junction, a pump, a pipe, a cold plate (CP) heat exchanger, and a liquid-to-liquid brazed-plate heat exchanger. Figure 2 shows the mass conservation and thermal energy conservation graphs for each component. For graphs in the hydraulic domain, vertices represent dynamic states of pressure, while edges represent the rate of mass transfer between vertices. For graphs in the thermal domain, vertices represent dynamic states of temperature, while edges represent thermal power flow between vertices due to convection in heat exchangers (HXs) or fluid transport. Dashed lines, indicating sources or sinks of each component, consist of variables determined by neighboring components or disturbances.

2.3 Hydraulic Graph Modeling. For notation purposes, a superscript m is used to denote capacitances, functions, and inputs associated with hydraulic graphs. The reader is referred to Ref. [16] 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 is the fluid volume in the component and both the fluid density ρ and the bulk modulus E are assumed to be constant. For reservoirs, $C_i^m = A_{c,i} / g$, where A_c is the reservoir cross-sectional area and g is the gravitational constant.

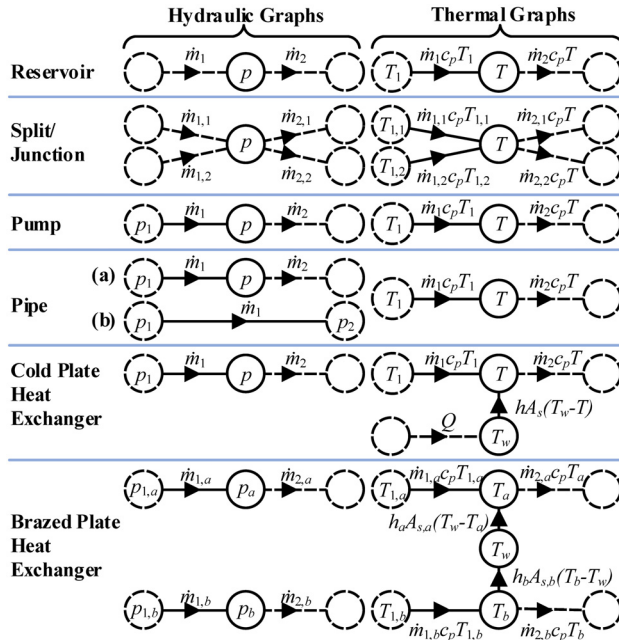


Fig. 2 Hydraulic and thermal graphs for individual fluid-thermal component models

For all hydraulic edges except those of a pump, the mass flow rate $\dot{m}_j = f_j^m(p_j^{\text{tail}}, p_j^{\text{head}}, u_j^m)$ is specifically given by

$$\dot{m}_j = \rho A_{c,j} \sqrt{\frac{2(p_j^{\text{tail}} - p_j^{\text{head}} + \rho g \Delta h_j)}{\rho \left(f_j \frac{L_j}{D_j} + K_{L,j} \right)}} \quad (8)$$

where L , D , and A_c are the fluid flow length, diameter, and cross-sectional area, respectively, Δh is the height difference between the inlet and outlet flow, f is the friction factor, and K_L is the minor loss coefficient. For pumps, the mass flow rate is given by

$$\dot{m}_j = \rho A_{c,j} \sqrt{2g \left(H_j - \frac{p_j^{\text{head}} - p_j^{\text{tail}}}{\rho g} \right)} \quad (9)$$

where the pump head H is determined using an empirical map as a linear function of pump speed ω and pressure differential across the pump

$$H = \alpha_1 + \alpha_2 (p_j^{\text{head}} - p_j^{\text{tail}}) + \alpha_3 \omega \quad (10)$$

where α_1 , α_2 , and α_3 are constants.

When hydraulic graphs of multiple components are interconnected to represent a system, the hydrodynamics can be represented in the form of Eqs. (6) and (7). The fluid system configuration used for demonstration in this paper consists of closed loops such that fluid mass does not enter or exit the system. Thus in the notation of the hydraulic graph variables, Eq. (6) reduces to

$$C^m \dot{p} = -\bar{M}^m \dot{m} \quad (11)$$

And Eq. (7) is given by

$$\dot{m} = F^m(p, u^m) = \left[f_j^m(p_j^{\text{tail}}, p_j^{\text{head}}, u_j^m) \right] \quad (12)$$

2.4 Thermal Graph Modeling. For notation purposes, a superscript t is used to denote capacitances and functions associated with thermal graphs. The reader is referred to Ref. [16] for a detailed derivation of the model equations that follow.

For all vertices associated with a fluid temperature, the thermal capacitance is given by $C_i^t = \rho V_i c_p$, where the specific heat capacitance of the fluid c_p is assumed to be constant. For all vertices associated with heat exchanger wall temperatures, $C_i^t = M_{w,i} c_{p,w}$, where M_w is the mass of the wall and $c_{p,w}$ is the specific heat capacitance of the wall material.

The thermal power flow $P_j = f_j^t(T_j^{\text{tail}}, T_j^{\text{head}}, \dot{m}_j^t)$ is specifically given by

$$P_j = \dot{m}_j^t c_p T_j^{\text{tail}} \quad (13)$$

for thermal power flow due to fluid flow, and by

$$P_j = h_j A_{s,j} (T_j^{\text{tail}} - T_j^{\text{head}}) \quad (14)$$

for convective thermal power flow in HXs, where A_s is the convective surface area and h is the heat transfer coefficient, approximated in this paper as a function of mass flow rate and temperature as $h \approx \beta_1 + \beta_2 \dot{m} T$, where β_1 and β_2 are constants.

When thermal graphs of multiple components are interconnected to represent a system, the thermodynamics can be represented in the form of Eqs. (6) and (7). In the notation of the thermal graph variables, Eq. (6) is given by

$$C^t \dot{T} = -\bar{M}^t P + D^t P^{\text{in}} \quad (15)$$

and Eq. (7) is given by

$$P = F^t(T, T^{\text{out}}, \dot{m}^t) = \left[f_j^t(T_j^{\text{tail}}, T_j^{\text{head}}, \dot{m}_j^t) \right] \quad (16)$$

2.5 Multigraph System Representation. Table 1 summarizes the quantities associated with each element of the graph-based framework, including in the generic sense of Sec. 2.1 and in the specific hydraulic and thermal domains of Secs. 2.3 and 2.4, respectively. The combined hydraulic and thermal system dynamics may be represented as two coupled graphs, as shown in Fig. 3. The hydraulic graph, denoted as G^m in Fig. 3, is governed by mass conservation laws, while the thermal graph, denoted as G^t , is governed by energy conservation laws. It is assumed in this paper that the coupling between these two graphs is limited to the unidirectional influence of hydrodynamics on thermodynamics.

Table 1 and Fig. 3 show that the mass flow rates calculated along the edges of G^m become inputs to the edges of G^t . However, the edges in G^m may not align one-to-one with the edges in G^t , in particular when a single mass flow rate affects multiple edges of the thermal graph. It is also possible that some mass flow rate inputs to the thermal system originate from its surroundings and are not modeled within the hydraulic graph. For example, this could include flow rates on the secondary side of HXs by which thermal energy is transferred to and from neighboring systems. These external flow rates are denoted by $\dot{m}^{\text{ext}} = [\dot{m}_i^{\text{ext}}]$, $i \in [1, N_{\text{ext}}]$ and treated as disturbances to the thermal model. The

Table 1 Summary of quantities in generic, hydraulic, and thermal graph-based models

Generic graph, G	Hydraulic graph, G^m	Thermal graph, G^t
Conserved quantity	Mass	Thermal energy
Vertex storage state, x	Pressure, p	Temperature, T
Edge transfer rate, y	Mass flow rate, \dot{m}	Power flow, P
Edge input, u	Actuator effort, u^m	Mass flow rate, \dot{m}^t




Fig. 3 Interconnection between a hydraulic graph (middle) and a thermal graph (top), with pump dynamics (bottom) affecting the hydraulic graph and hydrodynamics affecting the thermal graph

mass flow rates within G^m and \dot{m}^{ext} can then be mapped to the input mass flow rates of G^t by

$$\dot{m}^t = Z \begin{bmatrix} \dot{m} \\ \dot{m}^{\text{ext}} \end{bmatrix} \quad (17)$$

where $Z \in \mathbb{R}^{N_e^t \times (N_e^m + N_{\text{ext}})}$.

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 affecting the hydraulic graph, each u_i^m is paired with a single-input single-output (SISO) system S_i^p as shown in Fig. 3. Each S_i^p models the state of the i th pump u_i^m as a function of its commanded value u_i^p .

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

$$u_i^m(s) = \frac{e^{-\zeta_i^p s}}{\tau_i^p s + 1} u_i^p(s) \quad (18)$$

2.6 Hydraulic Graph Linearization. In general, the graph-based models have a nonlinear form but satisfy the generic relationships of Eqs. (1) and (2) for each vertex and edge. For control design, in particular, 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.

From Eqs. (8)–(10), the nonlinear hydraulic mass flow rate equations for all components follow the general form:

$$\dot{m}_j = c_{1,j} \sqrt{c_{2,j} + c_{3,j} (p_j^{\text{tail}} - p_j^{\text{head}})} + c_{4,j} u_j^m \quad (19)$$

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

$$\Delta \dot{m}_j = a_j^m (\Delta p_j^{\text{tail}} - \Delta p_j^{\text{head}}) + b_j^m \Delta u_j^m \quad (20)$$

where the Δ denote deviations from the linearization point and a_j, b_j are constant coefficients.

By substitution into Eq. (11), the linear equations for a complete hydraulic system model are given by

$$\dot{p} = A^m \Delta p + B^m \Delta u^m \quad (21)$$

where

$$A^m = -(C^m)^{-1} \bar{M}^m \text{diag}([a_j^m]) (\bar{M}^m)^T \quad (22)$$

$$B^m = -(C^m)^{-1} \bar{M}^m \bar{b}^m$$

\bar{M}^m represents the columns of \bar{M}^m corresponding to edges associated with pumps, and \bar{b}^m is the vector of input coefficients for edges associated with pumps (i.e., edges j for which $c_{4,j} \neq 0$ in Eq. (19)).

The output equation of the linearized hydraulic model relating pressures and pump efforts to mass flow rates is given by

$$\Delta \dot{m} = C^m \Delta p + D^m \Delta u^m \quad (23)$$

where

$$C^m = \text{diag}([a_i^m]) (\bar{M}^m)^T \quad (24)$$

$$D^m = [d_{jk}^m] \in \mathbb{R}^{N_e^m \times N_p^m}$$

and

$$d_{jk}^m = \begin{cases} b_j^m & \text{if edge } j \text{ is associated with pump } k \\ 0 & \text{else} \end{cases} \quad (25)$$

2.7 Thermal Graph Linearization. From the power flow equations detailed in Eqs. (13) and (14) and assumed expression for heat transfer coefficient, the nonlinear power flow equations for all components follow the general form

$$P_j = c_{1,j} T_j^{\text{tail}} + c_{2,j} T_j^{\text{head}} + c_{3,j} T_j^{\text{tail}} \dot{m}_j^t + c_{4,j} (T_j^{\text{tail}} - T_j^{\text{head}}) T_j^{\text{head}} \dot{m}_j^t \quad (26)$$

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

$$\Delta P_j = a_{1,j}^t \Delta T_j^{\text{tail}} + a_{2,j}^t \Delta T_j^{\text{head}} + b_j^t \Delta \dot{m}_j^t \quad (27)$$

By substitution into Eq. (15), the linear equations for a complete thermal system model are given by

$$\dot{T} = A^t \Delta T + B_1^t \Delta T^{\text{out}} + B_2^t \Delta \dot{m}^t + B_3^t \Delta P^{\text{in}} \quad (28)$$

where

$$A^t = -(C^t)^{-1} \bar{M}^t (\bar{M}^t)^T$$

$$B_1^t = -(C^t)^{-1} \bar{M}^t (\underline{M}^t)^T \quad (29)$$

$$B_2^t = -(C^t)^{-1} \text{diag}([b_j^t])$$

$$B_3^t = -(C^t)^{-1} D^t$$

and $\underline{M}^t = [m_{i,j}^t] \in \mathbb{R}^{(N_v^t + N_{\text{out}}^t) \times N_e^t}$ is a weighted incidence matrix for the thermal graph with

$$m_{i,j}^t = \begin{cases} a_{1,j}^t & \text{if } v_i \text{ is the tail of } e_j \\ a_{2,j}^t & \text{if } v_i \text{ is the head of } e_j \\ 0 & \text{else} \end{cases} \quad (30)$$

3 Modeling Example and Validation

3.1 Example Configuration Description. Appendix A describes the experimental testbed used for validation in this paper. The testbed is pictured in Fig. 4 in an example system configuration used for demonstration throughout this paper. The corresponding schematic is shown in Fig. 5. This configuration is notionally representative of a simplified aircraft FTMS in which heat from actuators, generators, engine oil, and other transient loads is absorbed, stored in liquid fuel, and rejected through transfer to neighboring systems [28].

The sample configuration has eight pumps arranged in four sets of two in series. In this paper, the *a* and *b* pumps of each set receive the same commands. Therefore, for notational convenience, the two pumps in each set are referred collectively. For example, pumps 1a and 1b are collectively termed as “pump 1.”

The secondary loop (identified as the left half of the system in Fig. 5) absorbs thermal loads from the heaters mounted to cold plate 1 (CP 1), through which fluid is driven by pump 1. This loop has dedicated thermal storage available in reservoir 2, and the ability to exchange thermal energy with the main loop across brazed-plate heat exchanger 1 (HX 1) with fluid driven by pump 2.

The primary loop (identified as the right half of the system in Fig. 5) includes two parallel fluid flow paths out of reservoir 2.

The path driven by pump 3 passes through heat exchanger 1, exchanging thermal energy with the secondary loop. The path driven by pump 4 passes through CP 2 and CP 3 absorbing thermal energy produced by their heaters. The two flow paths then junction and pass through HX 2, by which thermal energy is transferred out of the system to the thermal sink.

3.2 Graph-Based Representation of Example Configuration. The hydrodynamics of the example testbed configuration in Fig. 5 are represented by the system graph shown in Fig. 6, formed by the interconnection of the individual hydraulic component graphs from Fig. 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 7 shows the thermal graph for the example testbed configuration, formed by the interconnection of the individual thermal component graphs from Fig. 2. The edges exiting the three leftmost dashed vertices indicate heat transfer from the resistive heaters to the cold plates, treated as disturbances to the system. The two edges on the right side of the graph connected to dashed vertices denote source and sink power flows from and to the chiller, which is treated in this work as an infinite source/sink of thermal energy. Thus in Eq. (15)




Fig. 4 Example thermal-fluid testbed configuration for experimental validation

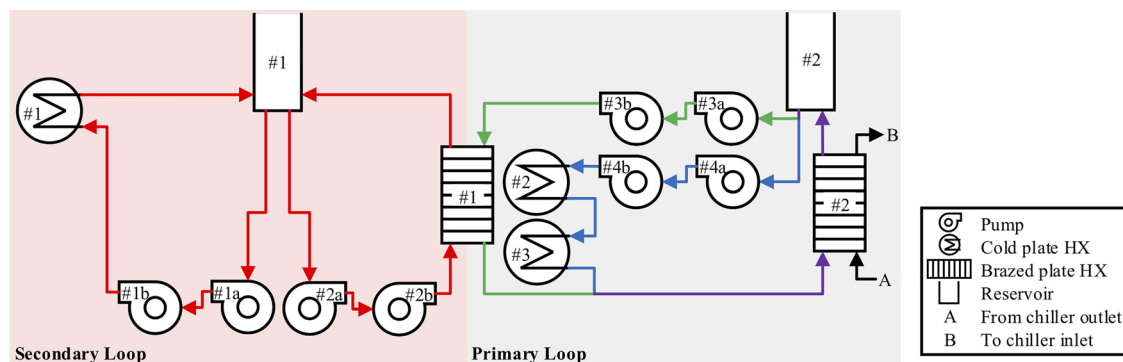


Fig. 5 Schematic of example testbed configuration

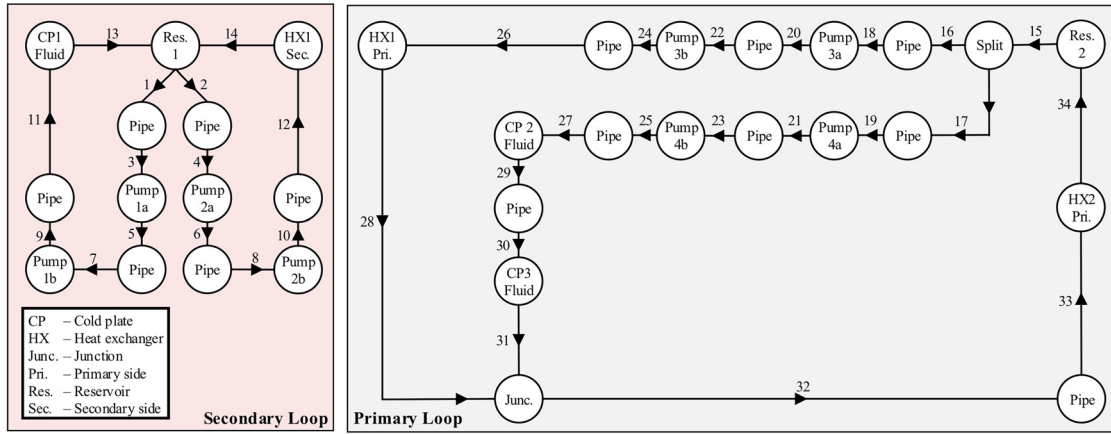


Fig. 6 Hydraulic graph for example testbed configuration

$$P^{in} = [Q_1 \quad Q_2 \quad Q_3 \quad \dot{m}_{ext} c_p T_c]^T \quad (31)$$

where each Q is the heat load to the corresponding cold plate heat exchanger, \dot{m}_{ext} is the mass flow rate of chilled fluid through the right side of HX 2, and T_c is the temperature of the fluid exiting the chiller and entering the right side of HX 2. 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.

3.3 Experimental Validation of Example Configuration.

Figure 8 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. 8 show the envelope between the maximum and minimum values measured at each time among the five trials.

The heat loads plotted in Fig. 8 are translated into a reference current for each bank of resistive heaters using an empirical map between the applied electrical current and the achieved thermal 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. Figure 8 shows that error from this set point of about 0.5 °C on

average is present due to measurement and tracking error within the chiller unit's internal controller.

Figures 9(a) and 9(b) show a selection of hydraulic and thermal signals, respectively, that result from applying the inputs and disturbances of Fig. 8. All experimental traces plotted show the envelope between the maximum and minimum values measured at each time among five experimental trials. To make the width of these envelopes more clear, a closer view of several signals is provided in Fig. 10, which demonstrates that the testbed exhibits a high degree of repeatability.

Figure 9(a) 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 pumps 2b and 3b mass flow rate, the nonlinear model is more accurate. This is due to the error incurred by linearization of the terms under the square root in Eqs. (8) and (9).

Figure 9(b) similarly demonstrates a high degree of accuracy in the nonlinear thermal graph-based model. The discrepancies that occur can be attributed to unmodeled friction/losses and errors in heat transfer coefficients. However, in the interval from 750 to 850 s significant error occurs for the linear model in signals pertaining to cold plates of the primary loop. This is largely due to the combination of a low speed in pump 3 and a high speed in pump 4, which falls far from the linearization conditions. Linearization of the inherently bilinear fluid-thermal power flow equation $P = \dot{m} c_p T$ 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

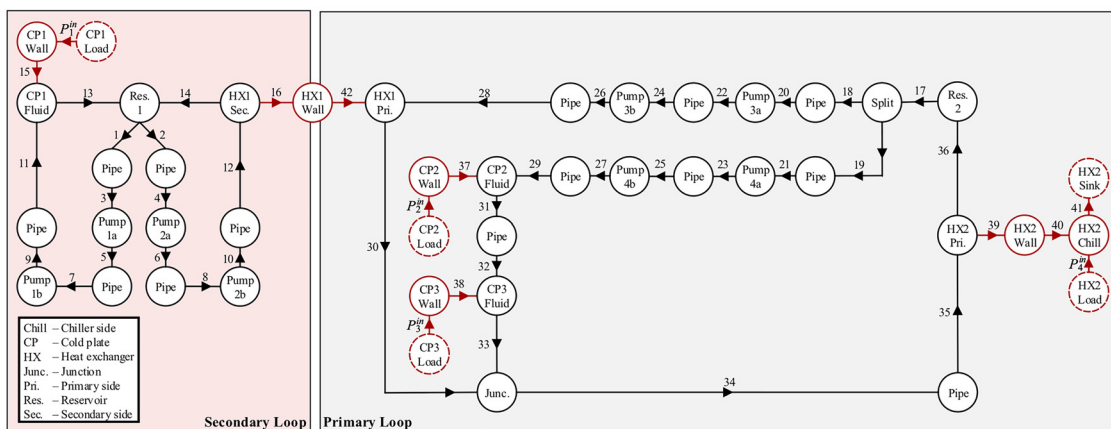


Fig. 7 Thermal graph for example testbed configuration

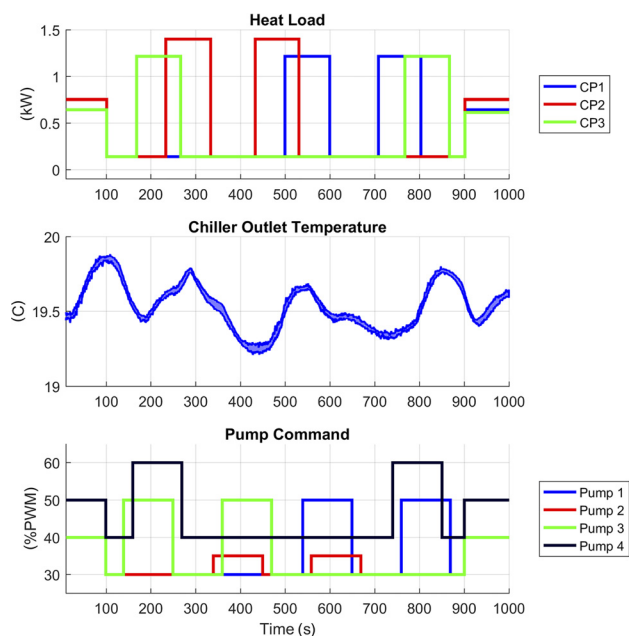


Fig. 8 Inputs and disturbances used for model validation

the correct 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 these models could be improved at the cost of increased complexity, the accuracy of the

proposed graph-based models is sufficient for their intended use in closed-loop control.

4 Graph-Based Hierarchical Control

4.1 Overview of Hierarchical Control Framework. Figure 11 diagrams the hierarchical control framework developed for the example testbed configuration. This framework consists of four layers, as enclosed by the dashed box. The update period of each control layer is included in Fig. 11 below its name. Lower layers have faster update rates because they are responsible for governing dynamics of faster timescales, progressing from slow thermal dynamics, to faster thermal dynamics, to hydrodynamics, and finally to pump dynamics in the pump control layer.

The thermal system control layer at the top of the hierarchy is responsible for coordinating the overall long-term thermal behavior of the system, leveraging preview of expected thermal disturbances for the mission at hand. Using model predictive control (MPC) with an update interval of 80 s, this controller is designed to primarily manage slow thermal dynamics, such as reservoir fluid temperatures, cold plate wall temperatures, and heat exchanger wall temperatures.

While the thermal system control layer plans for the coordination of the full thermal system over a long-time horizon, its slow update rate prohibits governing faster system dynamics or quickly compensating for model error or differences between preview of expected future disturbances and the true disturbances that affect the system. However, simply increasing the update rate of the thermal system control layer while maintaining the same time horizon may not be computationally tractable for complex systems, as this would increase the number of prediction steps to be solved by the optimization program while decreasing the time between consecutive updates in which the program must be

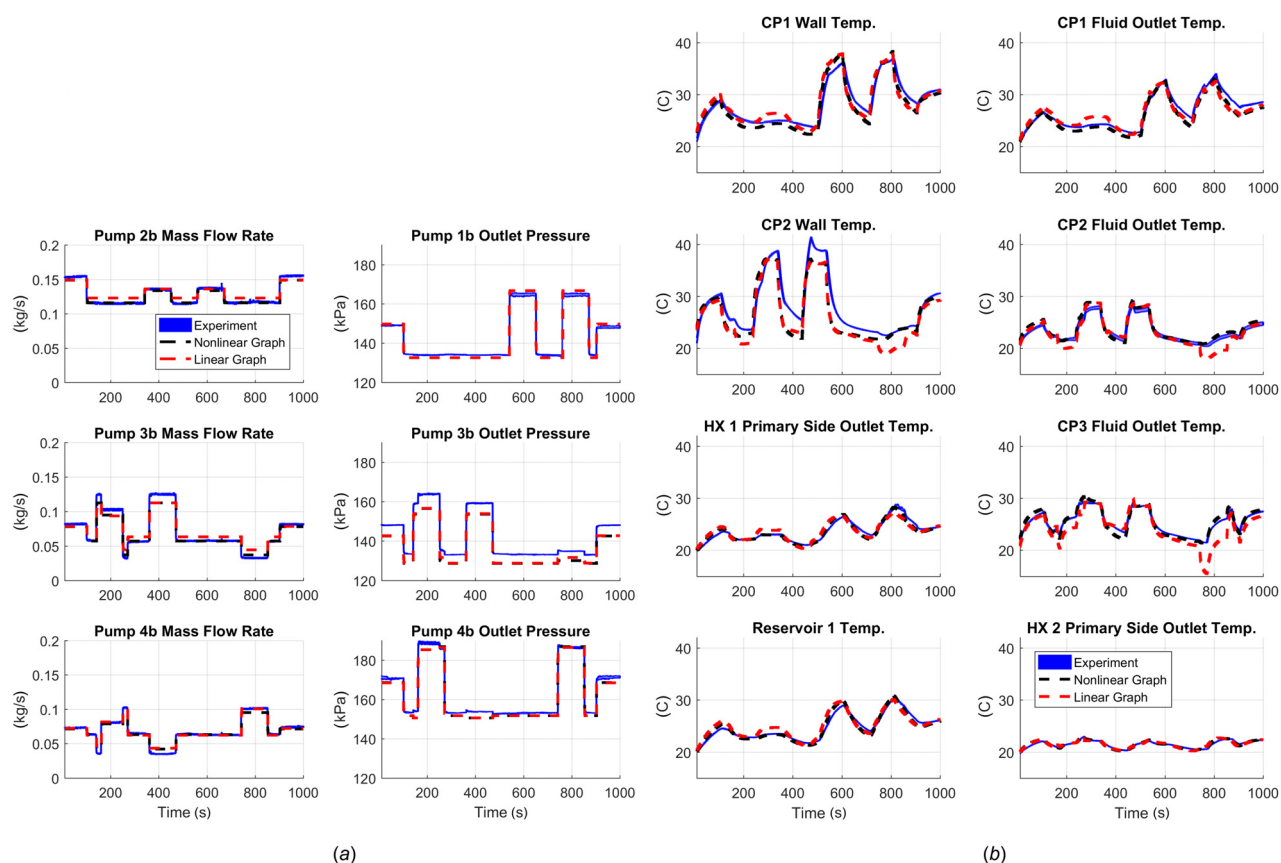


Fig. 9 Selected signals for validation of experimental data with nonlinear and linear graph-based models: (a) hydraulic signals and (b) thermal signals

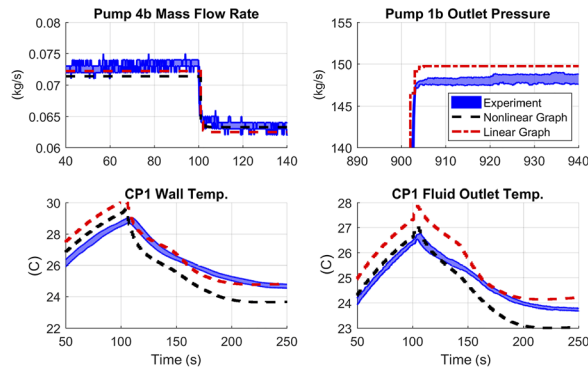


Fig. 10 Closer view of several signals from Fig. 9. All experimental traces show the envelope between the maximum and minimum values measured at each time among five experimental trials.

solved. This motivates the introduction of the faster updating thermal subsystem control layer.

The thermal subsystem control layer consists of two MPC controllers governing the thermal dynamics of the primary and secondary loops. The graph of the full thermal system in Fig. 7 is partitioned into separate graphs for each loop. Using the same general control formulation as for the thermal system controller, graph-based thermal controllers of a 2s update period are designed for each of the thermal subsystems.

References to be tracked for select temperatures and thermal power flows are communicated from the thermal system control layer to the thermal subsystem control layer. This leverages the

system-level coordination and long-time horizon in the thermal system controller, while the faster update rate of the thermal subsystem layer compensates for model and preview error to achieve thermal objectives. Because long-horizon planning is performed by the thermal system controller, the thermal subsystem controllers can be constructed with a relatively small number of steps in their prediction horizons.

The third layer in the hierarchy is the hydraulic subsystem control layer, which like the thermal subsystem control layer is partitioned into subsystems corresponding to each fluid loop. References for mass flow rates are sent from each controller in the thermal subsystem control layer to the corresponding controller in the hydraulic subsystem control layer. In this paper, we assume that only pressure measurements are available as hydraulic feedback, as pressure sensors may be preferred over mass flow rate sensors in implementation for their reduced cost, increased accuracy, and faster response time.

The controllers in the hydraulic subsystem control layer determine references for the pump states. In the pump control layer, a set of decoupled SISO controllers track the desired pump states by commanding the pump inputs, compensating for dynamics and time delays within each pump.

4.2 Thermal System Control Layer Formulation. The thermal system control layer leverages available preview of upcoming thermal loads in calculating references for temperatures and power flows over a prediction horizon. The primary objective of thermal management is to regulate temperature states of the system $T_i, i \in [1, N_v^t]$ such that $\underline{T}_i \leq T_i \leq \bar{T}_i$ where \underline{T}_i and \bar{T}_i are lower and upper bounds, respectively, on the i th temperature. However, a secondary objective to maintaining temperature

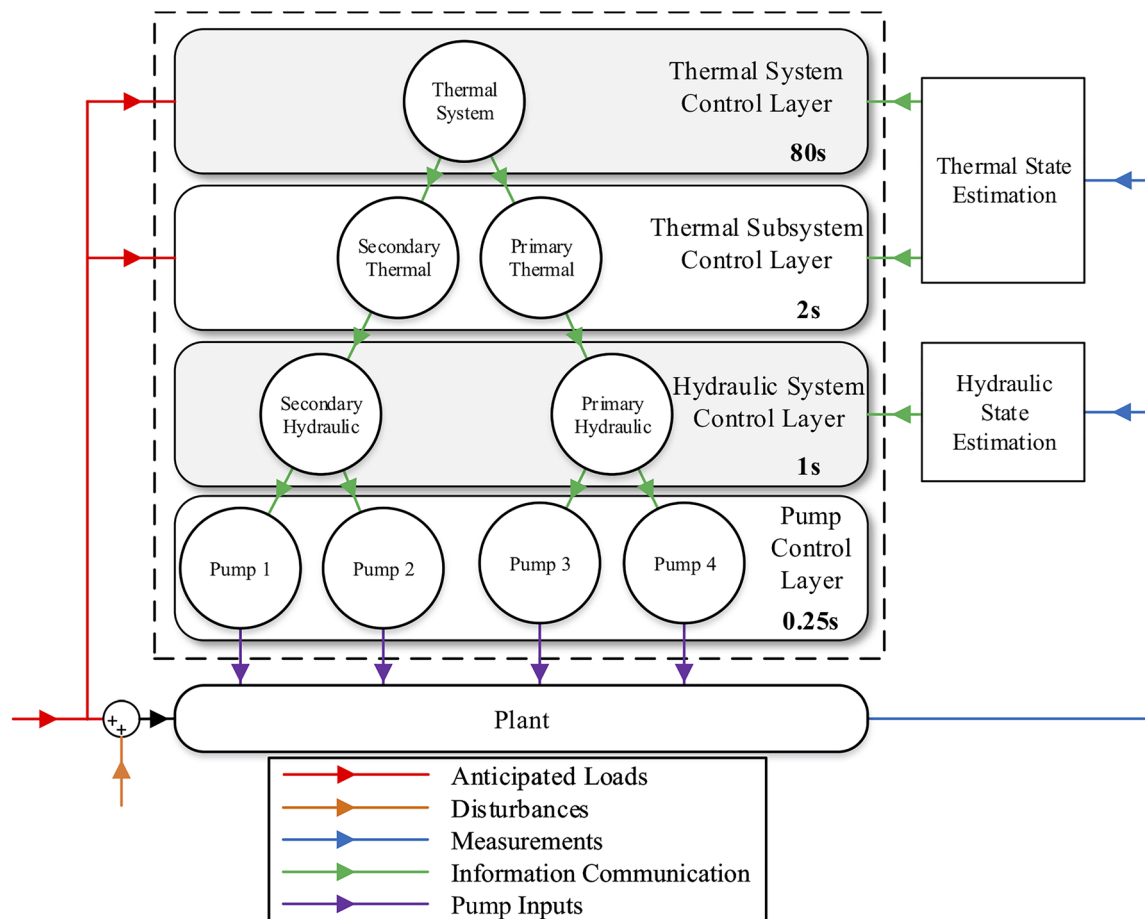


Fig. 11 Four layer hierarchical control framework for the testbed configuration

constraints is keeping the mass flow rates required to achieve the references small when possible, reducing the pump effort required for the system. The MPC in the thermal control layer solves the following constrained nonlinear program:

$$\begin{aligned} \min_{\dot{m}} \sum_{k=0}^{N_h^t-1} & \left(\lambda_s^t \|s^t(k)\|_2^2 + \lambda_u^t \|\dot{m}(k)\|_2^2 \right. \\ & + \lambda_v^t \|T^{\text{track}}(k+1) - T_{\text{preview}}^{\text{ref}}(k+1)\|_2^2 \\ & + \lambda_e^t \|P^{\text{track}}(k) - P_{\text{preview}}^{\text{ref}}(k)\|_2^2 \Big) \\ & + \sum_{k=0}^{N_h^t-2} \left(\lambda_d^t \|\dot{m}(k+1) - \dot{m}(k)\|_2^2 \right) \end{aligned} \quad (32)$$

subject to

$$T(0) = T_{\text{est}} \quad (33a)$$

$$\underline{T}_i - s_i^t(k) \leq T_i(k+1) \leq \bar{T}_i + s_i^t(k), i \in [1, N_v^t] \quad (33b)$$

$$s_i^t(k) \geq 0, i \in [1, N_v^t] \quad (33c)$$

$$\begin{aligned} \Delta T(k+1) = & A_{d,1}' \Delta T(k) + B_{d,1}' \Delta T_{\text{preview}}^{\text{out}}(k) \\ & + B_{d,2}' \Delta \dot{m}^t(k) + B_{d,3}' \Delta P_{\text{preview}}^{\text{in}}(k) \end{aligned} \quad (33d)$$

$$P^{\text{track}}(k) = F^{t,\text{track}} \left(T(k), T_{\text{preview}}^{\text{out}}(k), \dot{m}^t(k) \right) \quad (33e)$$

$$\dot{m}^t(k) = Z \begin{bmatrix} \dot{m}(k) \\ \dot{m}_{\text{preview}}^{\text{ext}}(k) \end{bmatrix} \quad (33f)$$

$$H \dot{m}(k) \leq z \quad (33g)$$

for $k \in [0, N_h^t - 1]$ and

$$\dot{m}(0) = \dot{m}_{\text{last}}(1) \quad (34)$$

In Eqs. (32)–(33), N_h^t is the number of steps in the prediction horizon of the thermal system control layer MPC. The subscript “preview” is included in some variables to indicate that preview of expected values for these signals is assumed to be available over the time horizon of the controller as part of the system’s planned mission. Equation (32) defines the cost function as minimizing the thermal slack variable $s^t = [s_i^t]$, $i \in [1, N_v^t]$ with weighting λ_s^t , the mass flow rates under control \dot{m} with weighting λ_u^t , the tracking error of a subset of system temperatures T^{track} from references T^{ref} with weighting λ_v^t , the tracking error of a subset of thermal power flows P^{track} from references P^{ref} with weighting λ_e^t , and the sum of the differences between mass flow rates in consecutive steps with weighting λ_d^t . The last of these objectives is used to ensure that control decisions remain smooth in time across the prediction horizon.

Equation (33a) sets the temperatures of the first step in the horizon equal to the current estimated values T_{est} . Equations (33b) and (33c) define s^t as measuring the extent by which temperature constraints are violated and impose that $s_i^t(k) \geq 0 \forall i, k$. Equation (33d) imposes a discrete form of the linear thermal graph-based model from Eq. (28). Equation (33e) imposes the nonlinear thermal power flow equation from Eq. (16) to compute those power flows with tracking objectives P^{track} . While the linearized thermal power flow equation of Eq. (27) could be used instead, the algebraic nature of this equation means that large error can occur at any steps in the horizon when values of \dot{m}^t are far from their linearization points, and so when computational resources are sufficient to employ a nonlinear solver, the nonlinear power flow model is preferred.

Equations (33f) and (33g) constrain the mass flow rates of the system. From Eqs. (17), (33f) defines the relationship between the mass flow rates that affect the thermal system \dot{m}^t , the mass flow rates controlled by the hydraulic system \dot{m} , and the mass flow rates that are external disturbances to the system \dot{m}^{ext} .

Equation (33g) defines the envelope of allowable mass flow rates. The presence of flow splits and junctions in the fluid loops means that it is possible for fluid to flow in the reverse direction from the orientation indicated by the arrows in Fig. 5. As reverse flow would typically be undesirable in an aircraft FTMS, constraints must be included in the controllers to avoid this behavior. In addition, constraints are required to define the envelope of achievable mass flow rates in each loop. The approach used to determine these constraints is detailed in Appendix B. For the thermal system model, $\dot{m} = \begin{bmatrix} \dot{m}_{\text{pri}} \\ \dot{m}_{\text{sec}} \end{bmatrix}$, $H = \begin{bmatrix} H_{\text{pri}} & 0 \\ 0 & H_{\text{sec}} \end{bmatrix}$, and

$z = \begin{bmatrix} z_{\text{pri}} \\ z_{\text{sec}} \end{bmatrix}$ from Eqs. (B1) and (B2) of Appendix B.

Finally, Eq. (34) dictates that the mass flow rates chosen for the first step in the prediction horizon $\dot{m}(0)$ be equal to those for the corresponding time chosen in the previous iteration of the controller $\dot{m}_{\text{last}}(1)$. This is performed under the assumption that the controller is afforded one step in duration to calculate a solution, after which the controller output decisions corresponding to the $k = 1$ step are applied.

4.3 Thermal Subsystem Control Layer Formulation. To create a model for each subsystem controller, the system graph in Fig. 7 is partitioned into subsystem graphs corresponding to the primary and secondary loops. The vertex corresponding to the wall temperature of HX 1, which couples the thermal dynamics of the two loops, is represented as a sink vertex in each subsystem graph.

Select temperatures and thermal power flows determined in the thermal system control layer are passed as references to the thermal subsystem control layer. Temperature references are downsampled to the update rate of the latter using a first-order hold, while power flow references are downsampled using a zero-order hold because they are calculated with an algebraic equation and can change quickly due to the dependence on mass flow rates.

The MPC formulation for each of the thermal subsystem controllers is identical to that in Eqs. (32)–(34) subject to use of signals and model parameters corresponding to the thermal subsystem graph-based models diagrammed in Fig. 7. For the primary thermal subsystem $\dot{m} = \dot{m}_{\text{pri}}$, $H = H_{\text{pri}}$, and $z = z_{\text{pri}}$, and for the secondary thermal subsystem $\dot{m} = \dot{m}_{\text{sec}}$, $H = H_{\text{sec}}$, and $z = z_{\text{sec}}$.

4.4 Hydraulic Control Layer Formulation. Mass flow rates determined in the thermal subsystem control layer are passed as references to the hydraulic subsystem control layer. These references are downsampled to the update rate of the hydraulic subsystem controllers using a zero-order hold. Each of the MPC controllers in the hydraulic control layer solves a constrained quadratic program using a discrete form of the linear hydraulic graph-based model from Eqs. (21) and (23). The formulation for the hydraulic controllers is found in Ref. [13].

4.5 Pump Control Layer Formulation. The pump control layer consists of MPC controllers for each pump. These track pump state references from the hydraulic subsystem control layer u^m by commanding the pump inputs u^p , compensating for pump dynamics and delays. These references are downsampled to the update rate of the pump controllers using a zero-order hold. Each of the MPC controllers in the pump control layer solves a constrained quadratic program using a discrete form of the pump model from Eq. (18). The formulation for the pump controllers is found in Ref. [13].

5 Benchmark Controllers

5.1 Decentralized Benchmark Control. Figure 12 diagrams the decentralized benchmark controller for the example system configuration of this paper. Each set of pumps is paired with a PI loop that seeks to bring a temperature measurement from the system to track a constant reference. Table 2 lists the signal used as feedback for each PI loop.

The choice of the feedback signal for each loop reflects the primary purpose of the overall system to manage the temperature of the cold plate walls. PI 1 tracks a reference temperature for the CP 1 wall. PI 2 and PI 3 track reference temperatures for reservoir 1 and the HX 1 primary side outlet, respectively, governing the exchange of thermal energy between the secondary and primary loops by controlling fluid flow on either side of HX 1. Because pump 4 moves fluid through both CP 2 and CP 3, using the wall temperature of just one of these cold plates as the feedback signal to PI 4 could result in an inability to properly manage the temperature of the other cold plate. Therefore, the CP 3 outlet fluid temperature, which is affected by the wall temperature of both CP 2 and CP 3, is used as the feedback signal to control pump 4.

5.2 Centralized Benchmark Control. The centralized benchmark controller is identical to the thermal system control layer formulation of Sec. 4.2. The mass flow rates determined in this controller are translated into input commands to each pump using a static mapping, as depicted in Fig. 13.

6 Closed-Loop Experiments

While heat loads in the experimental testbed are generated by resistive heaters, in application these would be generated by high-power electrical equipment, such as electrical actuators and batteries, which may have tightly constrained operating temperatures. In this paper, temperature limits are embedded into thermal MPC formulations as soft constraints \underline{T} and \bar{T} , parametrized in accordance with Table 3.

To evaluate the ability of the controllers to manage tight constraints, only 5 °C separate the minimum and maximum temperature constraints for CP 1. Managing the temperature of CP 1 not only involves proper control of the pumps in the secondary

subsystem but also coordination with the primary subsystem to transfer thermal energy across heat exchanger 1, providing cooling to the secondary subsystem.

The heat load to each cold plate, the temperature of the chiller outlet fluid, and the mass flow rate of the chiller outlet fluid all serve as disturbances to the system. The chiller is given a constant temperature set point of 20 °C and mass flow rate of 0.35 kg/s. These are assumed to be known to the thermal MPC controllers as preview information across their prediction horizon.

Figure 14 shows the cold plate heat loads for a 4800 s mission, which serves as a case study to evaluate closed-loop performance. Each heat load nominally performs several large step changes on the order of 1 kW, as shown in the top subplot of Fig. 14. However, the applied loads also include additive disturbances consisting of higher frequency uniformly distributed noise stepped every 40 s of up to 200 W in each direction, as shown in the bottom subplot of Fig. 14. The heat loads are saturated to enforce a 1.7 kW maximum load to each cold plate. The nominal loading profile represents reconfiguration of the system as mission phases change over hundreds or thousands of seconds, which in application could coincide with changes in duty cycles of electrical equipment generating heat, charge and discharge of batteries, etc. As such, the nominal loads are assumed to be known to the thermal MPC controllers as mission preview information across their prediction horizon. However, the higher frequency noise added to the nominal heat loads is not included in this preview information and instead serves as an unknown thermal disturbance.

Linear Kalman filters based on a discrete form of the linear graph-based models estimate the unmeasured states. These serve as the thermal state estimation and hydraulic state estimation blocks pictured in Figs. 11 and 13.

The MPC-based controllers are formulated using the YALMIP toolbox [29]. Constrained quadratic programs are solved with the Gurobi optimization suite [30], while nonlinear programs are solved with the IPOPT software package [31].

6.1 Decentralized Benchmark Control. Table 2 lists the specific reference temperatures and gains for each PI controller loop, corresponding to the input/output pairings shown in Fig. 12. These were designed manually by iterating over closed-loop simulations

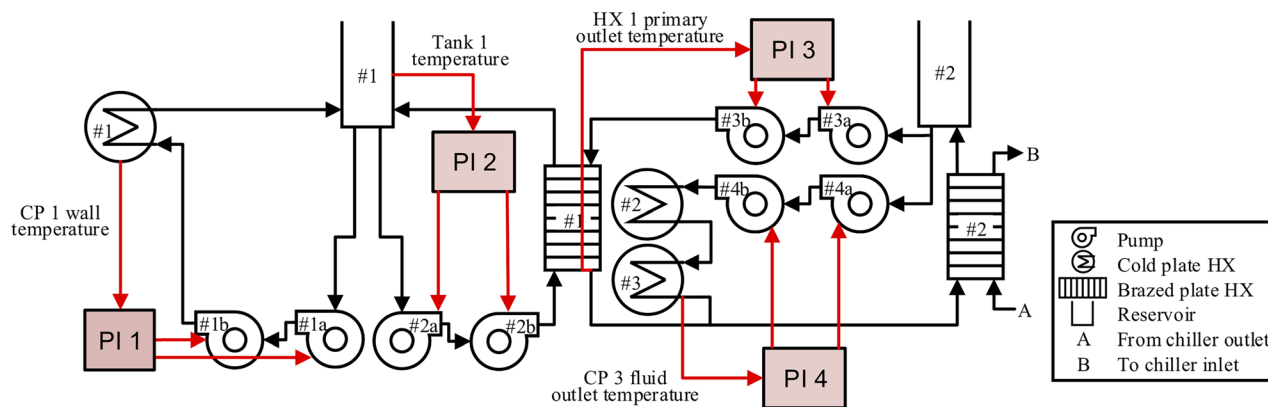


Fig. 12 Decentralized benchmark controller framework

Table 2 Feedback signals for PI loops of decentralized benchmark controller

Controller loop	Pump	Feedback signal	Reference temp. (°C)	Proportional gain	Integral gain
PI 1	Pump 1	CP 1 wall temperature	42.5	3	0.05
PI 2	Pump 2	Reservoir 1 temperature	35	5	0.05
PI 3	Pump 3	HX 1 primary outlet temperature	35	5	0.05
PI 4	Pump 4	CP 3 fluid outlet temperature	40	1	0.05

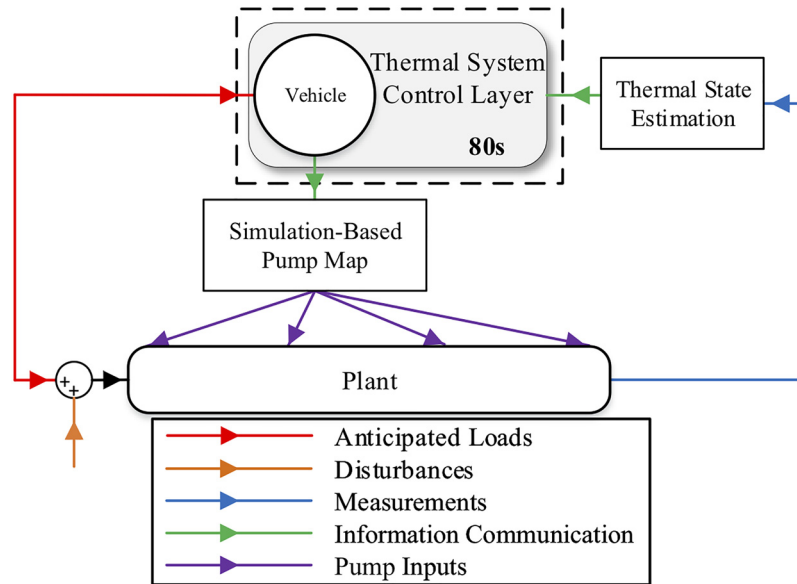


Fig. 13 Centralized benchmark controller framework

with the goal of bringing the system to achieve thermal objectives. The reference temperature for the CP 1 wall of 42.5 °C falls exactly in the middle of its constraints. The CP 3 fluid outlet temperature reference is 10 °C below the CP 2 wall temperature upper bound, and 5 °C below the CP 3 wall temperature upper bound.

Figure 15(a) shows the pump commands chosen by the decentralized PI controller in experimental application and Fig. 15(b) shows the resulting cold plate wall temperatures. While the temperatures of CP 1 and CP 3 are held closely within their constraints, the temperature of CP 2 is not, with multiple durations of violations more than 5 °C at peak. This emphasizes a drawback of employing a decentralized single-input single-output approach to control complex systems. One actuator may be responsible for managing many elements but the control approach does not incorporate a comprehensive awareness of the condition of all those elements. This is often the case in thermal management systems, where a single fluid line may be responsible for cooling multiple thermal loads [28].

6.2 Centralized Benchmark Control. The centralized controller is updated every 80 s. The number of steps in the prediction horizon is $N_h^{t,sys} = 20$, therefore the time horizon is 1600 s. For this controller, T^{track} consists of only the temperature of the CP 1 wall, which tracks a constant reference of 42.5 C, and there are no thermal power flow tracking objectives. The following weightings are used in the objective function of Eq. (32): $\lambda_s^{t,sys} = 10^4$, $\lambda_u^{t,sys} = 10^{-4}$, $\lambda_v^{t,sys} = 10^5$, $\lambda_d^{t,sys} = 10^{-3}$.

The contribution to the total objective function cost of each term of Eq. (32) is further scaled by the dimension and magnitude of the signals in each term. Therefore, the weightings above indicate that the highest priority is placed via $\lambda_s^{t,sys}$ on minimizing violations of the temperature bounds. The weighting on minimizing mass flow rates $\lambda_u^{t,sys}$ is comparatively small, indicating that this

should only be done when possible without violating the temperature constraints.

Figure 16(a) shows the pump commands chosen by the centralized controller in experimental application and Fig. 16(b) shows the resulting simulated cold plate wall temperatures. Large constraint violations occur in cold plate temperatures of the primary loop, which exceed constraints by up to 10 °C. This is due to model error in the slow updating controller. During periods of constraint violation, the controller's model predicts that the CP 2 and CP 3 wall temperatures will decrease to within constraints over the next update period, during which the mass flow rates applied are those chosen for the second step in the horizon of the previous iteration of the controller due to the one-step delay built into the control design to accommodate computation time. This limits the ability of the controller to leverage measurement feedback in compensating for model and preview errors.

6.3 Hierarchical Control. The thermal system control layer of the hierarchical framework in Fig. 11 is identical to the centralized benchmark controller of the previous section. The two controllers in the thermal subsystem control layer are updated every 2 s, with $N_h^{t,sub} = 2$ steps in the prediction horizon. A long-time horizon is not critical for the thermal subsystem controllers because their primary objective is to track references from above

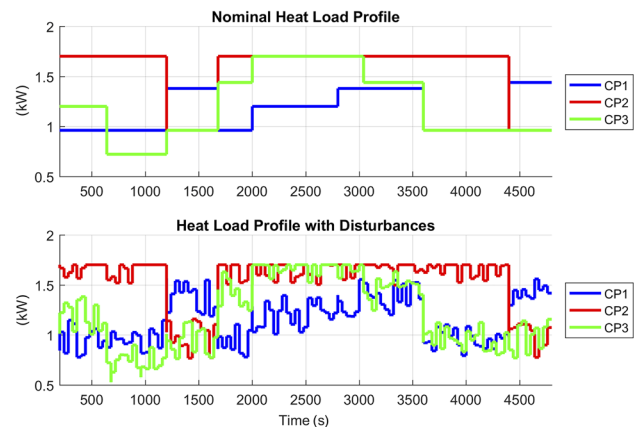


Fig. 14 Heat load disturbance profile for closed-loop experiments

Table 3 Temperature constraints for closed-loop experiments

Temperature	\underline{T} (°C)	\bar{T} (°C)
CP 1 Wall	40	45
CP 2 Wall	15	45
CP 3 Wall	15	50
All others	15	50

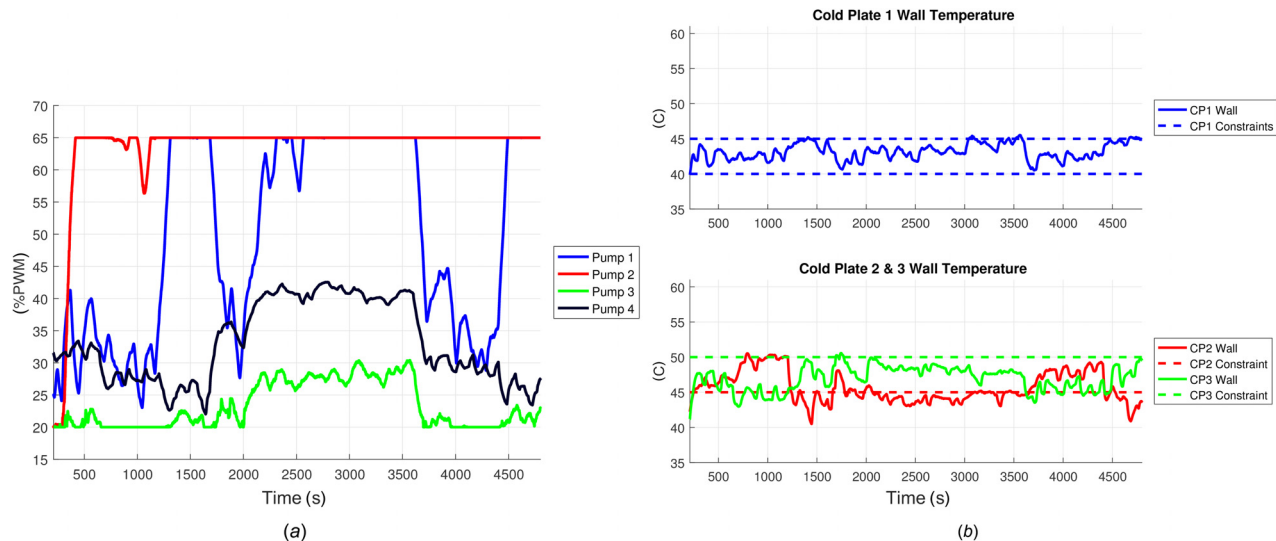


Fig. 15 Decentralized PI controller in experiment: (a) pump commands and (b) cold plate temperatures

in the hierarchy that have already been optimized over a long-time horizon.

For each thermal subsystem controller, T^{track} consists of references from the thermal system control layer for the temperature of all CP walls, HX walls, and reservoirs present in the subsystem. These represent the slowest dynamics, whose behaviors evolve over the timescale of the thermal system control layer. Because the HX 1 wall is treated as a sink temperature in both subsystem thermal graphs, the temperature for this vertex predicted by the thermal system control layer is included in the sink state preview $T^{\text{out}}_{\text{preview}}$ to each subsystem controller. P^{track} for each thermal subsystem controller consists of references for the thermal power flow along edges that couple the subsystems to each other and to the chiller. In Fig. 7, these are labeled in the primary loop as edges 39 and 42, and in the secondary loop as edge 16. These power flow references ensure that the coordination among subsystems planned in the thermal system control layer is achieved by the thermal subsystem control layer with no requirement for direct communication between subsystem controllers.

The following weightings are used in the objective function of Eq. (32) for each thermal subsystem controller: $\lambda_s^{\text{sub}} = 10^7$, $\lambda_u^{\text{sub}} = 10^{-6}$, $\lambda_v^{\text{sub}} = 10^5$, $\lambda_e^{\text{sub}} = 10^3$, $\lambda_d^{\text{sub,pri}} = 10^8$, $\lambda_d^{\text{sub,sec}} = 3 \times 10^9$.

These weightings indicate that the highest priority is on tracking references from the layer above in the hierarchical framework.

The two controllers in the hydraulic subsystem control layer are updated every 1 s, with $N_h^{\text{sub}} = 3$ steps. For all pressures, $p_i = 10 \text{ kPa} \forall i$, and $\bar{p}_i = 200 \text{ kPa} \forall i$.

The four controllers in the pump control layer are updated every 0.25 s with $N_p^p = 10$ steps. For all pumps, $u_i^p = 20\% \forall i$, and $\bar{u}_i^p = 65\% \forall i$.

Figure 17(a) shows the pump commands chosen by the hierarchical controller in the experiment. Figure 17(b) shows the resulting cold plate wall temperatures, which exhibit significantly improved constraint management as compared to the centralized benchmark controller. This improvement is largely due to the inclusion of the thermal subsystem control layer, which leverages a faster update rate to compensate for thermal model error and reject much of the unknown noise added to the nominal heat load profile in tracking references from above in the hierarchy.

6.4 Controller Comparison Summary. Figure 18(a) shows the total temperature constraint violations under each controller, computed by integrating the magnitude of constraint violations over time for each cold plate. The integral of violations in CP 1 is

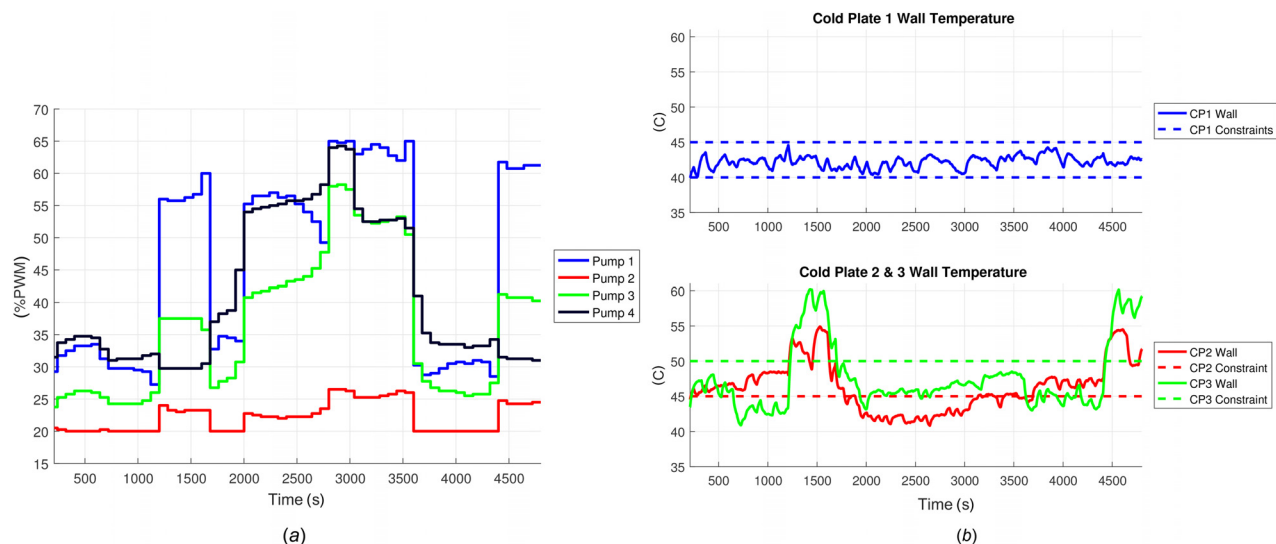


Fig. 16 Centralized MPC controller in experiment: (a) pump commands and (b) cold plate temperatures

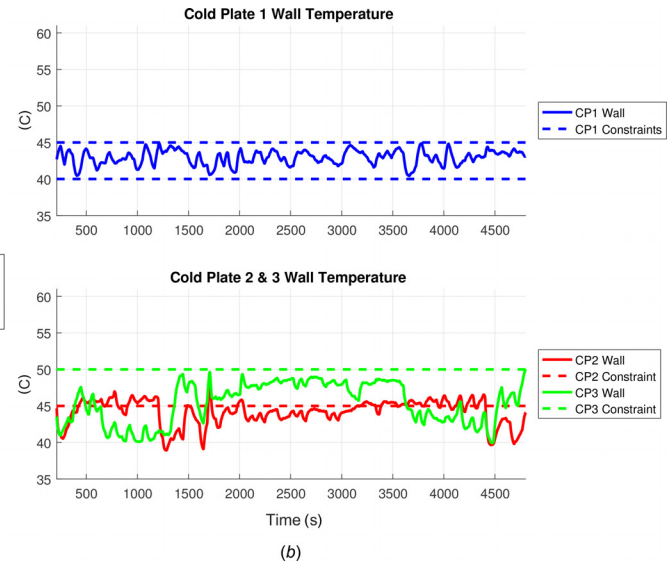



Fig. 17 Hierarchical MPC controller in experiment: (a) pump commands and (b) cold plate temperatures

small enough for all controllers to not be visible in this figure. Figure 18(b) shows the peak temperature violation across the mission in each cold plate under each controller.

From Figs. 18(a) and 18(b), it is clear that the hierarchical framework outperforms both of the baseline controllers, having 7.7% of the total violations of the centralized baseline controller, 21% of the total violations of the decentralized PI baseline controller, and greatly reduced peak violations in CP 2 and CP 3.

Figure 19 shows the total energy consumed by the pumps under each controller, calculated as a function of the electrical current measured for each pump during the experiments. Although the objective of minimizing mass flow rates is given a relatively small weight in the MPC formulations, the centralized and hierarchical approaches still require less total pump energy than the decentralized PI approach. Most importantly, Fig. 19 shows that differences in thermal performance between the three control approaches are not due to increasing the overall actuator effort of the system but instead due to a better allocation of actuation.

7 Conclusions

This paper applies a graph-based modeling approach and hierarchical control framework for thermal management. Graph-based dynamic models are derived from conservation of mass and thermal energy, where vertices represent storage elements and edges

capture the transport of mass and energy. Experimental validation with a testbed fluid-thermal system demonstrates the high accuracy of the modeling approach. The graph-based framework facilitates model-based control design and is especially well suited to hierarchical control designs, where the control structure should reflect the structure inherent in the graph framework. A scalable hierarchical framework is proposed to manage the multidomain and multi-timescale dynamics present in the system. The proposed hierarchical controller is experimentally demonstrated on a testbed system and compared to decentralized and centralized benchmark controllers, where it is found to perform significantly better in managing thermal objectives by compensating for model error and rejecting unknown disturbances.

In cases where a centralized controller can be solved using a complete system model at a fast update rate and long-time horizon, implementation of hierarchical control is likely not warranted. However, beyond a certain level of complexity executing such a centralized controller becomes intractable given the limited computational resources on board vehicle systems. Therefore, hierarchical control represents an enabling technology for achieving high performance in highly complex thermal management systems, where centralized control is not sufficiently scalable and decentralized control can result in poor performance or necessitate overconservative designs due to an inability to manage coupling between subsystems.

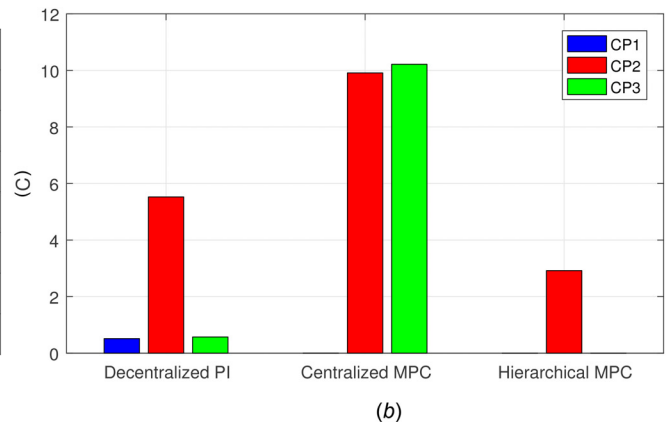
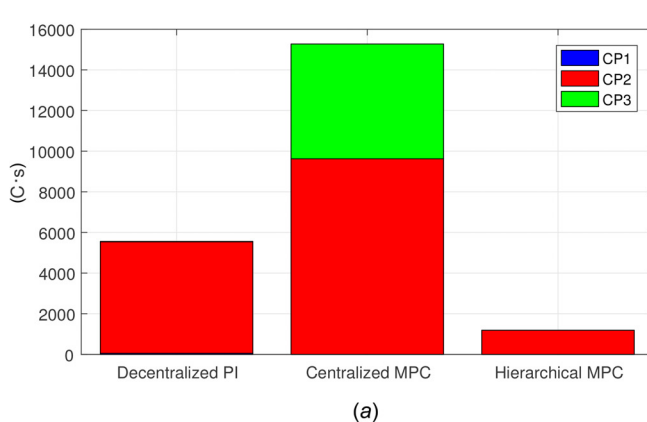


Fig. 18 Comparison of constraint violations for hierarchical and baseline controllers: (a) total temperature violations and (b) peak cold plate wall temperature violations




Fig. 19 Comparison of total pump energy consumption for hierarchical and baseline controllers

Future work will extend the model-based control design to better capture nonlinear regimes in the system dynamics by performing a local linearization at each controller update and/or representing the system as a set of switched linear models with modes specific to operating regions. Formal procedures will also

be explored for choosing the number of layers of the hierarchy, the partitioning of dynamics within each layer, and the update rate of each layer. Future work will also include control of other energy domains present in vehicle energy systems, using the graph-based models for electrical and turbomachinery components presented in Ref. [32]. Lastly, analytical techniques for ensuring stability will be incorporated into the hierarchical control demonstration, such as the passivity approach for graph-based models in Ref. [33], which has been extended to switched graph-based models in Ref. [34].

Funding Data

- National Science Foundation Graduate Research Fellowship (Grant No. DGE-1144245).
- National Science Foundation Engineering Research Center for Power Optimization of Electro-Thermal Systems (POETS) with cooperative agreement EEC-1449548.
- Air Force Research Laboratory (Contract No. FA8650-14-C-2517).

Appendix A: Experimental Testbed Overview

This experimental testbed was developed to emulate features of fluid-based thermal management systems while being rapidly reconfigurable to allow for numerous configurations. Table 4 and

Table 4 Testbed component descriptions

Component	Specifications	Number supported
(a) 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) Brazed-plate HX	Koolance HXP-193 12 plates 4.0 kW at 5 LPM and 20 °C inlet temperature difference	4
(c) Cold plate HX	Wakefield-Vette 6-pass, 6" exposed cold plate Vishay LPS1100H47R0JB thick film resistors, 47 Ω, 1100 W max power each Crydom 10PCV2415 solid-state relay Echun Electronic Co. ECS1030-L72 noninvasive current sensor	4
(d) Pipe	Koolance HOS-13CL Clear PVC 13 mm × 16 mm	—
(e) Reservoir	Koolance 80 × 240 mm Acrylic 8" eTape liquid-level sensor	4
(f) Chiller	Polyscience 6000 Series Up to 2900 W at 20 °C −10 °C to +70 °C	1
(g) Temp. sensor	Koolance SEN-AP008B (fluid) Koolance SEN-AP007P (surface) 10 KΩ thermistor	16
(h) Pressure sensor	Measurement Specialties US300 Up to 310 kPa gauge	7
(i) Flow rate sensor	Aqua computer high flow 0.5–25 LPM	8




Fig. 20 Individual components as labeled in Table 4 with a 6" ruler for scale. A fluid temperature sensor is pictured in (g).




Fig. 21 Infrared image of CP 1 and reservoir 1 from example testbed configuration in Fig. 4

Fig. 20 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. Components are connected via flexible tubing.

Centrifugal pumps are the primary fluid movers in the system. Speed is controlled via a PWM % 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%.

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

Each CP heat exchanger consists of several 47 Ω resistive heaters mounted to an aluminum cold plate that has copper tubing passing through. The heaters on each cold plate are wired to a solid-state relay actuating the heater power output. Up to four heaters can be mounted on each cold plate; however, in this paper just two are used, allowing a maximum heat load of 1.7 kW to be applied to each cold plate.

The reservoirs act as thermal storage elements. A liquid-level sensor inside each reservoir is used to calculate its liquid mass and therefore its thermal capacitance.

A 1.5 HP (1.12 kW) industrial chiller acts as a thermal energy sink (e.g., a vapor compression system). With variable temperature control from -10°C to 70°C , the chiller can emulate a wide range of sink conditions.

Infrared cameras were used to identify locations on the HX and CP walls that closely represent the average wall temperature, at which surface temperature sensors are affixed. The infrared image in Fig. 21 shows CP 1 and reservoir 1 of the example testbed configuration in Fig. 4. The cable for the CP 1 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. From within LabVIEW, signals can be exchanged with MATLAB/SIMULINK either by running the two programs simultaneously and communicating via the user datagram protocol or by embedding MATLAB code in LabVIEW using a MATLAB script node.

Appendix B: Hydraulic Coupling Constraints

To determine the constraints on mass flow rate for closed-loop control, the nonlinear hydraulic model is simulated to steady-state at all combinations of pump speeds in the range of 20–65% PWM in increments of 0.25%. As a safety margin against complete 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 resulting envelope of mass flow rates through pumps 3 and 4 is shown in Fig. 22(a).

Figure 22(b) shows the envelope of pump commands generating mass flow rates in the envelope of Fig. 22(a). This demonstrates that coupling between the two pump commands must be taken into account to avoid reverse flow.

The envelope in Fig. 22(a) is assumed to be a polyhedron, defined by the linear inequality

$$E_{\text{pri}} = \{\dot{m}_{\text{pri}} | H_{\text{pri}} \dot{m}_{\text{pri}} \leq z_{\text{pri}}\} \quad (\text{B1})$$

where H_{pri} is a matrix of appropriate dimensions and z_{pri} is a vector. Vertices used to define this polyhedron are circled in Fig. 22(a).

The envelope of achievable mass flow rates in the secondary loop is defined similarly by

$$E_{\text{sec}} = \{\dot{m}_{\text{sec}} | H_{\text{sec}} \dot{m}_{\text{sec}} \leq z_{\text{sec}}\} \quad (\text{B2})$$

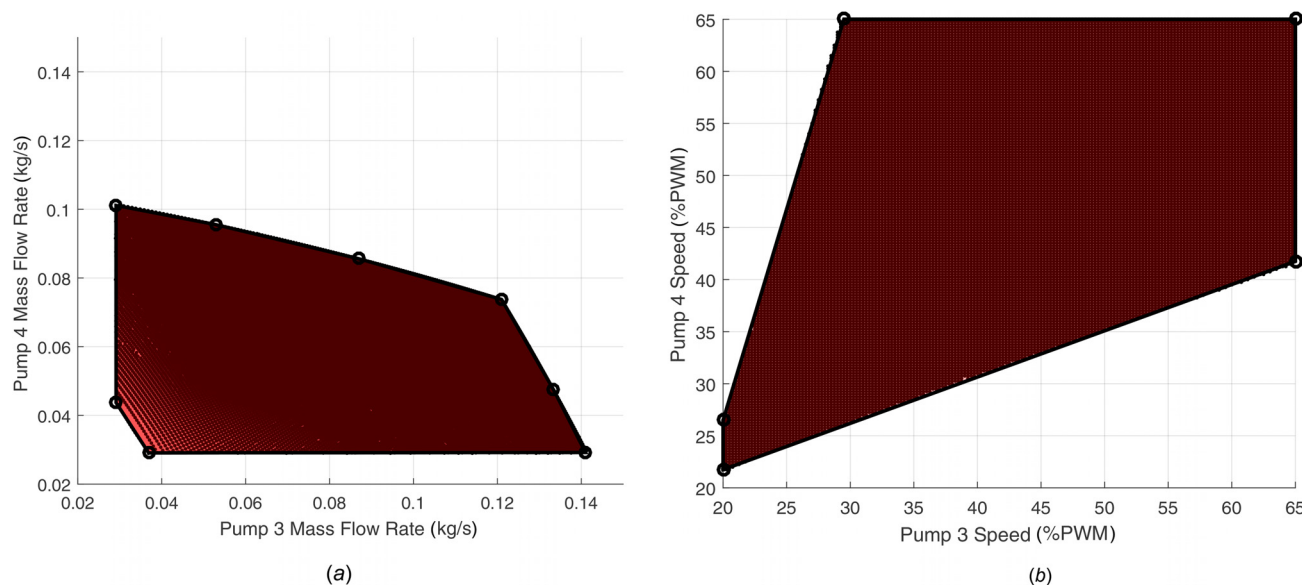


Fig. 22 Hydraulic coupling in pumps 3 and 4: (a) envelope of mass flow rates above 0.03 kg/s and (b) pump commands generating mass flow rates in Fig. 22(a)

References

- [1] Dooley, M., Lui, C., and Newman, R. W., 2013, "Efficient Propulsion, Power, and Thermal Management Integration," *AIAA Paper No. 2013-3681*.
- [2] Williams, M., 2014, "A Hierarchical Control Strategy for Aircraft Thermal Systems," *M.S. thesis*, University of Illinois at Urbana-Champaign, Champaign, IL.
- [3] Bodie, M., Russell, G., McCarthy, K., Lucas, E., Zumbeke, J., and Wolff, M., 2010, "Thermal Analysis of an Integrated Aircraft Model," *AIAA Paper No. 2010-288*.
- [4] Doty, J., Yerkes, K., Byrd, L., Murthy, J., Alleyne, A., Wolff, M., Heister, S., and Fisher, T. S., 2017, "Dynamic Thermal Management for Aerospace Technology: Review and Outlook," *J. Thermophys. Heat Transfer*, **31**(1), pp. 86–98.
- [5] Yerkes, K. L., Scofield, J. D., Courson, D. L., and Jiang, H., 2014, "Steady-Periodic Acceleration Effects on the Performance of a Loop Heat Pipe," *J. Thermophys. Heat Transfer*, **28**(3), pp. 440–454.
- [6] Thomas, S. K., and Yerkes, K. L., 1997, "Quasi-Steady-State Performance of a Heat Pipe Subjected to Transient Acceleration Loadings," *J. Thermophys. Heat Transfer*, **11**(2), pp. 306–308.
- [7] Bodie, M., and Wolff, M., 2010, "Robust Optimization of an Aircraft Power Thermal Management System," *AIAA Paper No. 2010-7086*.
- [8] Doman, D. B., 2016, "Fuel Flow Control for Extending Aircraft Thermal Endurance—Part I: Underlying Principles," *AIAA Paper No. AIAA 2016-1621*.
- [9] Doman, D. B., 2015, "Rapid Mission Planning for Aircraft Thermal Management," *AIAA Paper No. AIAA 2015-1076*.
- [10] Doman, D. B., 2016, "Fuel Flow Control for Extending Aircraft Thermal Endurance—Part II: Closed Loop Control Closed Loop Control of Dual Tank Systems," *AIAA Paper No. AIAA 2016-1622*.
- [11] Deppen, T. O., Hey, J. E., Alleyne, A. G., and Fisher, T. S., 2015, "A Model Predictive Framework for Thermal Management of Aircraft," *ASME Paper No. DSCC2015-9771*.
- [12] Pangborn, H., Hey, J. E., Deppen, T. O., Alleyne, A. G., and Fisher, T. S., 2016, "Hardware-in-the-Loop Validation of Advanced Fuel Thermal Management Control," 46th AIAA Thermophysics Conference, Washington, DC, June 13–17, pp. 1–8.
- [13] Pangborn, H., Williams, M., Koeln, J., and Alleyne, A., 2017, "Graph-Based Hierarchical Control of Thermal-Fluid Power Flow Systems," *American Control Conference (ACC)*, Seattle, WA, May 24–26, pp. 2099–2105.
- [14] Moore, K. L., Vincent, T. L., Lashhab, F., and Liu, C., 2011, "Dynamic Consensus Networks With Application to the Analysis of Building Thermal Processes," *IFAC proc.*, **44**(1), pp. 3078–3083.
- [15] Preisig, H. A., 2009, "A Graph-Theory-Based Approach to the Analysis of Large-Scale Plants," *Comput. Chem. Eng.*, **33**(3), pp. 598–604.
- [16] Koeln, J. P., Williams, M. A., and Alleyne, A. G., 2015, "Hierarchical Control of Multi-Domain Power Flow in Mobile Systems—Part I: Framework Development and Demonstration," *ASME Paper No. DSCC2015-9908*.
- [17] Rasmussen, B., 2005, "Dynamic Modeling and Advanced Control of Air Conditioning and Refrigeration Systems," *Ph.D. dissertation*, University of Illinois at Urbana-Champaign, Champaign, IL.
- [18] Puntel, A., Emo, S., Michalak, T. E., Ervin, J., Byrd, L., Tsao, V., and Reitz, T., 2013, "Refrigerant Charge Management and Control for Next-Generation Aircraft Vapor Compression Systems," *SAE Paper No. 2013-01-2241*.
- [19] Deppen, T. O., 2013, "Optimal Energy Use in Mobile Applications With Storage," *Ph.D. dissertation*, University of Illinois at Urbana-Champaign, Champaign, IL.
- [20] Srivastava, S. K., Cartes, D. A., Maturana, F., Ferrese, F., Pekala, M., Zink, M., Meeker, R., Carnahan, D., Staron, R., Scheidt, D., and Huang, K., 2008, "A Control System Test Bed for Demonstration of Distributed Computational Intelligence Applied to Reconfiguring Heterogeneous Systems," *IEEE Instrum. Meas. Mag.*, **11**(1), pp. 30–37.
- [21] Scattolini, R., and Colaneri, P., 2007, "Hierarchical Model Predictive Control," 46th IEEE Conference on Decision and Control, New Orleans, LA, Dec. 12–14, pp. 4803–4808.
- [22] Scattolini, R., 2009, "Architectures for Distributed and Hierarchical Model Predictive Control—A Review," *J. Process Control*, **19**(5), pp. 723–731.
- [23] Hill, R. C., Cury, J. E. R., De Queiroz, M. H., Tilbury, D. M., and Lafortune, S., 2010, "Multi-Level Hierarchical Interface-Based Supervisory Control," *Automatica*, **46**(7), pp. 1152–1164.
- [24] Bendtsen, J., Trangbaek, K., and Stoustrup, J., 2012, "Hierarchical Model Predictive Control for Plug-and-Play Resource Distribution," *Distributed Decision Making and Control*, Springer, London, pp. 339–358.
- [25] Rahnama, S., Bendtsen, J. D., Stoustrup, J., Member, S., and Rasmussen, H., 2017, "Robust Aggregator Design for Industrial Thermal Energy Storages in Smart Grid," *IEEE Trans. Smart Grid*, **8**(2), pp. 902–916.
- [26] Koeln, J. P., Williams, M. A., Pangborn, H. C., and Alleyne, A. G., 2016, "Experimental Validation of Graph-Based Modeling for Thermal Fluid Power Flow Systems," *ASME Paper No. DSCC2016-9782*.
- [27] West, D. B., 2001, *Introduction to Graph Theory*, 2nd ed., Pearson, London.
- [28] Roberts, R. A., and Eastbourn, S. M., 2014, "Vehicle Level Tip-to-Tail Modeling of an Aircraft," *Int. J. Thermodyn.*, **17**(2), pp. 107–115.
- [29] Lofberg, J., 2004, "YALMIP: A Toolbox for Modeling and Optimization in MATLAB," *IEEE International Symposium on Computer Aided Control Systems Design*, New Orleans, LA, Sept. 2–4, pp. 284–289.
- [30] Gurobi Optimization, 2016, *Gurobi Optimizer Reference Manual*, Gurobi Optimization, Houston, TX.
- [31] Wächter, A., and Biegler, T. L., 2006, "On the Implementation of an Interior-Point Filter Line-Search Algorithm for Large-Scale Nonlinear Programming," *Math. Programming*, **106**(1), pp. 25–57.
- [32] Williams, M. A., Koeln, J. P., Pangborn, H. C., and Alleyne, A. G., 2018, "Dynamical Graph Models of Aircraft Electrical, Thermal, and Turbomachinery Components," *ASME J. Dyn. Syst., Meas., Control*, **140**(4), p. 041013.
- [33] Koeln, J. P., and Alleyne, A. G., 2017, "Stability of Decentralized Model Predictive Control of Graph-Based Power Flow Systems Via Passivity," *Automatica*, **82**, pp. 29–34.
- [34] Pangborn, H. C., Koeln, J. P., and Alleyne, A. G., 2018, "Passivity and Decentralized MPC of Switched Graph-Based Power Flow Systems," *American Control Conference*, Milwaukee, WI, June 27–29 (in press).