energies

MDPI

Article

Thermal Data Assimilation into a Real-Time Digital Twin of Liquid-Cooled Power Electronics via an Edge-Resident Particle Swarm Framework †

Braden Priddy, Josiah Worch, Kerry Sado, Richard Hainey, Austin R. J. Downey *, Jamil Khan and Kristen Booth

Molinaroli College of Engineering and Computing, University of South Carolina, Columbia, SC 29208, USA; bpriddy@email.sc.edu (B.P.); jworch@email.sc.edu (J.W.); ksado@email.sc.edu (K.S.); khan@cec.sc.edu (J.K.); kristen.booth@sc.edu (K.B.)

* Correspondence: austindowney@sc.edu

† The views expressed are those of the author and do not reflect the official policy or position of the Department of War or the U.S. Government.

Abstract

The next generation of naval and defense systems will strain current naval ship cooling systems. Throughout its life-cycle, this strain will alter the behavior of the physical system, and any virtual representation of the system will become outdated due to component aging. Digital twins are a trending tool that can assimilate real-time sensor data to tailor a virtual representation to its physical counterpart. The online faithful virtual representation of the physical system provided by digital twins can be used for real-time system optimizations and proactive fault detection, diagnostics, and control adjustments, alleviating the stress of component aging. To support these complex power systems throughout their lifecycles, data-driven solutions for digital twin tuning will become essential. This paper investigates the application of a parameter-tuning digital twin framework to enhance the performance of a multi-physics model. The digital twin framework comprises a digital twin tuning scheme, a physical testbed designed to emulate the cooling system of a ship, and a multi-physics representation of that system. The digital twin tuning scheme leverages a swarm of particles and online sensor data to evaluate permutations of parameters to update the digital representation periodically. The digital twin framework was applied to a physical system to provide experimental data results demonstrating the usefulness of the tuning system. The physical system was designed and constructed to emulate the heat generation and dissipation from 6 liquid-cooled power converters under loads ranging from 10–15 kW at 99% efficiency. Two scenarios were applied to evaluate the performance of the digital twin framework. Results demonstrate that the digital twin framework can adapt to system changes in real-time and improve the accuracy of the related virtual representation by more than 90% when measured at four points of the system under test.

Keywords: digital twin; particle swarm optimization; edge computing; data assimilation; thermal modeling

Academic Editor: Boyu Qin

Received: 12 February 2026
Revised: 30 April 2026
Accepted: 14 May 2026
Published: 20 May 2026

1. Introduction

As threats against the next generation of defense systems increase in scale and complexity, modernizing its existing components is imperative. To combat these novel threats, the Navy is shifting towards directed energy weapons such as air defense radar, Unmanned

aerial systems jamming, and laser weapons [1]. However, these new systems have more significant power demands than legacy weapon systems [2]. To address this issue, the Navy has begun switching its fleet to fully electric ships. Integrating new electric propulsion systems will ensure that the loads of next-generation weapon systems are met and offer additional benefits, such as enhanced fuel efficiency, lower noise levels, and reduced emissions [3]. Fully integrated electric propulsion systems are certainly viable, and have been utilized on commercial vessels. However, the implementation of this technology onto Navy vessels introduces a new set of engineering challenges. Naval weapon and mission systems require reliable power distribution, which, if not properly managed, could necessitate the shutdown of the systems for indeterminate periods. Moreover, a lack of stability in the power grid of a ship can lead to thermal, mechanical, and electrical stress on the vessel, jeopardizing the safety of its crew. Consequently, modernizing the power and energy systems on ships has become more imperative than ever to support new combat systems [4].

To modernize naval fleets and support directed energy weapons, a revolution in control systems is necessary. Traditional naval power and thermal management systems are unsuitable for the large and complex load profiles of modern directed energy weapons [5]. These systems operate reactively, adjusting to power fluctuations only after they occur, which is incompatible with the real-time requirements of modern vessels [6]. Additionally, legacy control systems lack integrated fault detection, making it difficult to anticipate or address potential system failures needed to proactively manage rapid power fluctuations of next-generation weapon systems [7]. Moreover, traditional naval systems are compartmentalized, leading to communication gaps between subsystems, requiring input from human operators [8]. In contrast, a fully electric ship requires an interconnected autonomous control architecture capable of managing ship resources, optimizing performance based on desired operation, and adapting to faults, all in real time. To achieve this, significant advances in modeling and simulation are necessary to form a virtual representation of a ship that can predict system behavior under various conditions [9].

Digital twins are a tool that can enable better control and performance of a system by continuously assimilating real-time sensor data into faithful system representations to link the virtual and physical spaces together [10]. As such, a digital twin comprises two key elements: a virtual representation and a feedback loop between the virtual and physical counterparts [10]. The digital twin is continuously updated using real-time sensor data to accurately replicate and adapt to the behavior of the physical counterpart during its operational state [11]. While the concept of digital twins has been around since the early 2000s, their implementation was hindered by the technology of that time [12]. Recently, digital twins have gained popularity with the advent of the fourth industrial revolution, as evidenced by the increased number of publications and patents in recent years [13]. As engineering problems become more complex and improving the efficiency of existing systems becomes more challenging, digital twins have emerged as an attractive tool for performing functions such as prognostication, optimization, testing, and control [14]. For instance, a digital twin updated using fuzzy logic and operational data from a power plant has demonstrated the feasibility of creating an accurate virtual representation of the plant [15]. This approach allowed for safer training of new control room operators and studying different control schemes. Other digital twin applications include online parameter optimization to increase the efficiency of a system. The cooling water system of a district plant behavior was reproduced and helped in optimization efforts, achieving 2–3% in energy saving from the previous year [16]. Another example of increased system performance boosted by digital twins, is using online measurements from a heat pump to calibrate a model to reduce fouling and unplanned downtime [17]. Real-time health monitoring is also an area that significantly benefits from digital twin implementation. A

digital twin of cooling fans achieved a 95% fault detection success rate, informing a user of a fault before it had occurred [18]. Furthermore, a digital twin modeled on a system designed to emulate liquid-cooled power converters detected faults in the cooling system with a 100% success rate in under 1 min [19]. A complete digital twin implementation enables real-time monitoring, improves system reliability, enhances risk management, and increases system efficiency [20]. In naval applications, a digital twin can be coupled with other existing combat systems and electrical models to enhance the survivability of a ship. Digital twins allow operators to address dynamic changes in the system and support strategic decision-making. Most importantly, this tool can aid in the operational management of the next generation of naval electric ships and high-power systems.

Digital twins have become increasingly important for use in advancing the design, operation, and maintenance of complex engineering systems. While many publications suggest different frameworks for the development of digital twins, the National Academy of Engineering (NAE) has identified and the general characteristics and implementation procedures for a useful digital twin [10]. Using the framework information provided by the NAE, digital twins can be developed and implemented for a variety of unique physical systems. The standardization provided by the NAE allows for consistency in the production of quality digital twins for a variety of physical systems. This consistency is especially critical for mission-driven systems such as naval vessels, where the digital twin framework must operate in real-time, be computationally efficient, and accurate within an acceptable margin of error. Meeting these conditions requires continuous digital twin tuning techniques utilizing reliable data streams and processing capabilities to update the virtual representation to reflect changes in its physical counterpart.

Population-based optimization algorithms have recently gained considerable traction as a framework for updating digital twin faithful representations. Data-driven algorithms can learn from a large amount of data to model the complex behavior of a system that other techniques fail to capture [21]. These meta-heuristic algorithms seek an optimal solution to a cost function by utilizing a swarm of model instances to systematically explore a search-space [22]. Specifically, particles in the swarm will each represent different variations of the same virtual representation. The swarm iteratively evaluates them based on real-time sensor data and seeks to minimize discrepancy between the simulation and actual system behavior. The manner in which these particles move throughout the search space is defined by an algorithm specific to its swarm-like behavior, discussed in more detail later in this paper. A comprehensive review of the most popular swarm optimization algorithms was done by testing them on 30 different benchmark functions [23]. It was concluded that in an environment where computation cost and speed are necessary, Particle Swarm Optimization (PSO) performed the best. Population-based algorithms have shown successful implementation in updating multi-physics model parameters for heat management, such as optimizing cooling strategies in a data center [24], enhancing prediction performance of CNC machining technology spindle thermal error [25], digital twin controller of HVAC systems [26]. These publications focus on the performance of the digital twin rather than the speed at which it is updated. In the context of updating a virtual representation of a cooling system to form a digital twin, speed is of the utmost importance [27]. Among population-based methods, PSO was chosen as the model updating method selected for the present framework because bounded-time optimization on an edge computer requires low implementation complexity, direct handling of bounded parameters, and repeated restart between successive data windows.

This work makes three contributions. First, it formulates a parameter-tuning digital twin framework for a liquid-cooled power electronics thermal system using a coupled electrical, fluid, and thermal representation. Second, it implements an edge-resident particle

swarm framework that executes on the same local computing environment as the data acquisition and simulation workflow, which removes network latency from the update loop and constrains calibration to fixed ten-minute data windows. Third, it introduces recovery, calibrated, and stall window logic so that the swarm can switch between rapid state recovery and fine parameter refinement after a physical discontinuity. Relative to prior digital twin studies that emphasize offline calibration or accuracy improvement under nominal conditions [24–27], the present framework emphasizes bounded update latency, computational feasibility at the edge, and recovery from abrupt system changes under experimentally relevant thermal loads. The framework is demonstrated on a liquid-cooled testbed representative of shipboard thermal management for six power converter modules, and the continuously updated representation improves root mean squared error by more than 90% at four measurement locations. The remainder of this paper is organized as follows. Section 2 discusses the methodology behind the proposed digital twin, the PSO framework, the cooling system, and modeling methods. The results and discussion are presented in Section 3. The final section provides a conclusion to this work.

2. Materials and Methods

This section provides an overview of the digital twin framework, the cooling system of the power and energy testbed, the faithful representation, data acquisition, and updating scheme that form the digital twin framework.

2.1. Digital Twin Framework

The digital twin framework aims to assimilate real-time sensor data to link a physical counterpart to its virtual representation. An overview of the digital twin framework is showcased in Figure 1. During the operation of a system, operators control various mechanical and electrical components to meet certain objectives. Throughout the lifecycle of the system, these actions will cause the system to degrade, forever changing its behavior. These changes are captured by sensors instrumented throughout the system, and data is relayed to the digital twin via data acquisition equipment. Running alongside the physical system, the digital twin utilizes the acquired data to tailor the virtual representation to its physical counterpart. To accomplish this goal, a digital twin tuning using an updating technique within the digital twin autonomously tunes parameters in the faithful representation. The tuned faithful representation outputs simulation data and is evaluated against the real-world data. The process of assessing permutations of virtual representation parameters happens continuously until the optimal parameters are found. Once the optimum is found, the virtual representation is deemed calibrated, and the projected digital twin response is returned to the operator. The operator can then leverage the digital twin to perform strategic planning, system optimization, or increase operational efficiency.

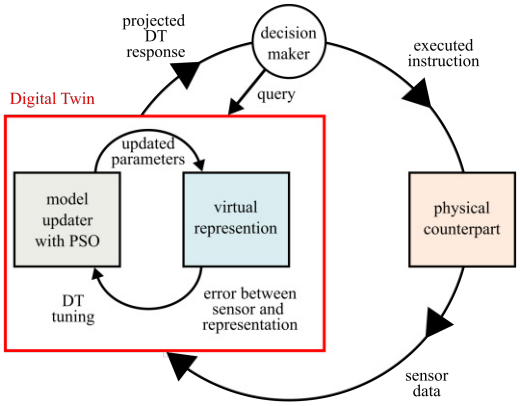

Figure 1. Diagram of the digital twin framework.

2.2. Testbed Configuration Overview

In this study, a notional shipboard power system is used as an example to validate and study the proposed digital twin framework. The power system is replicated by a testbed composed of six power electronic converters that interface with various loads as illustrated in the electrical setup section of the diagram in Figure 2. These converters are part of a microgrid designed to emulate the power and energy systems of a naval ship, effectively replicating the onboard power system [28]. Each module is designed to manage specific voltage inputs and outputs while supporting a load. The converters in use typically produce power losses ranging from approximately 100–150 W, which is comparable to power losses generated by many Naval power electronics [29]. Typically, these power losses would be handled by the cooling system. However, the power modules were not operational in this work and only act as heat sinks. Instead, the dissipated heat from the power modules was emulated by a water heater.

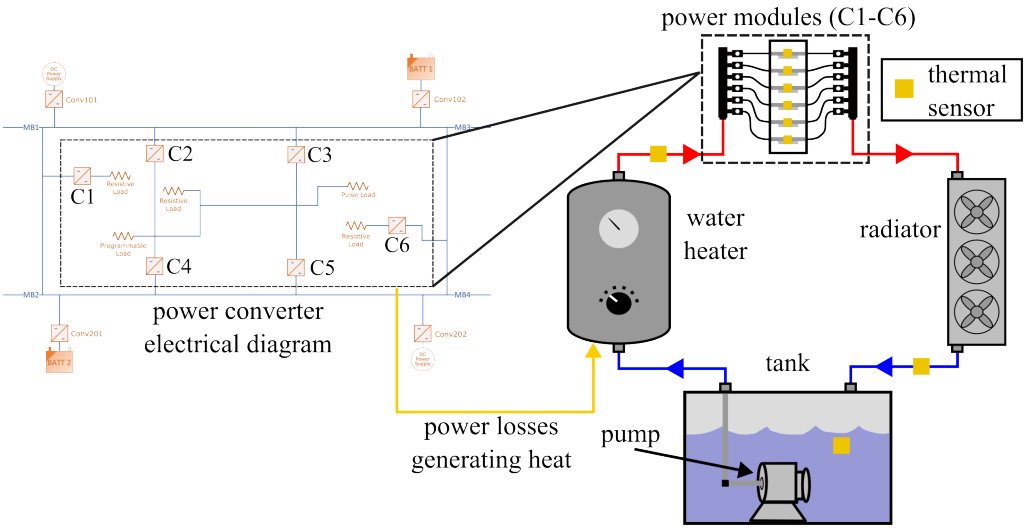

Figure 2. Electro-Thermal diagram of testbed system showing heater operation emulating general power converter module heat production.

The testbed is outfitted with a cooling loop for heat management and designed to mimic the cooling system of a naval ship. The cooling loop is composed of a submersible pump, tank, water heater, a three-fan radiator, as shown in Figure 3, and the six power modules shown in detail in Figure 4. The testbed is also instrumented with nine temperature sensors that record temperature data during experiments. There are four points where these temperature sensors are located: six on the heat sinks of the power modules, one submerged in the tank, and one after both the heater and radiator. The direction in which the water flows through the cooling loop is shown in Figure 2. Distilled water is circulated throughout the loop at 3.785 L/min by a submersible pump inside a 37.85-L tank. Water first flows through a water heater coupled with a power supply where the dissipated heat produced by ship systems is emulated based on the approximate power losses from the six labeled power converters, C1–C6, during standard operation. This heated water then flows through the coolant plates of six power modules. While the power modules can produce heat, they were not operational during testing and only act as thermal masses. After the power modules, water enters the fan radiator and the chilled water returns to the tank. These components help form the cooling system that deploys a 40-h load profile simulating a battle-time scenario.

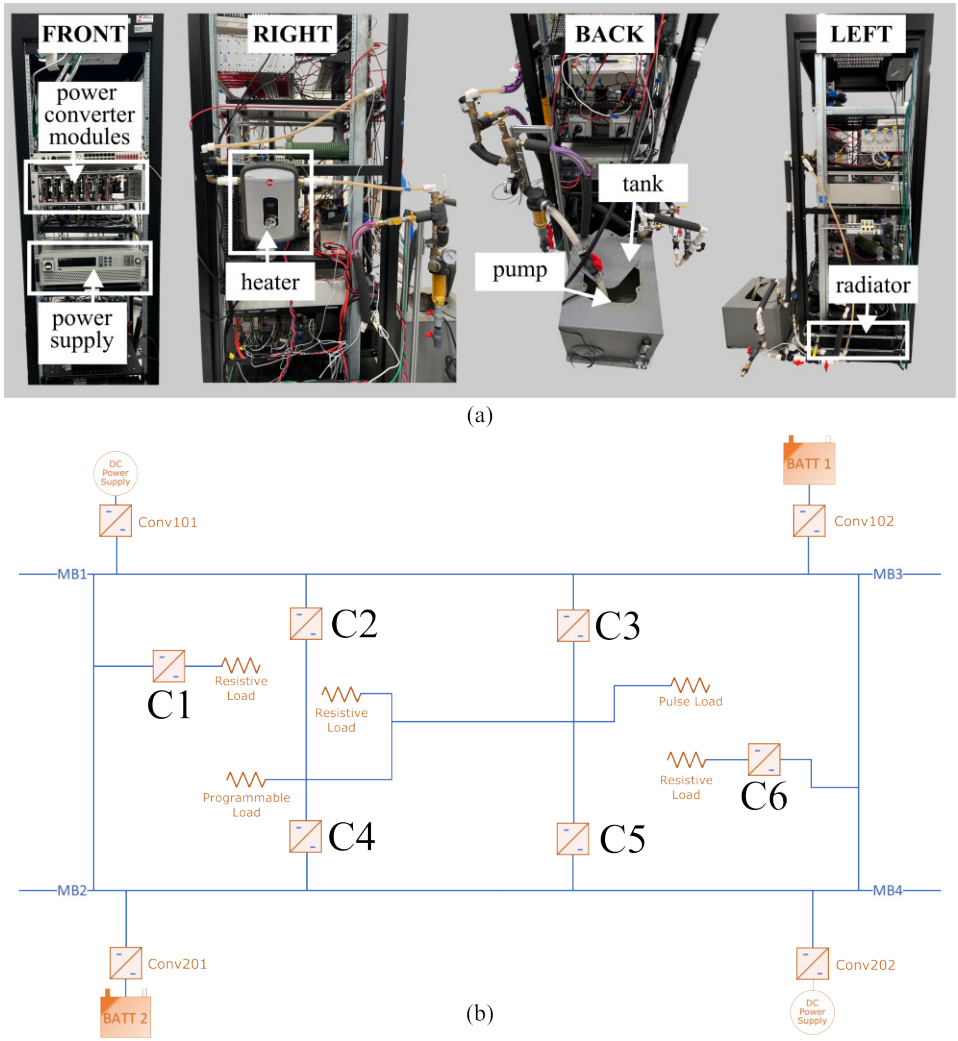

Figure 3. Electronic power converter full testbed, showing: (**a**) physical testbed views from the front, left, right, and back, with labeled components, and; (**b**) power converter system electrical diagram.

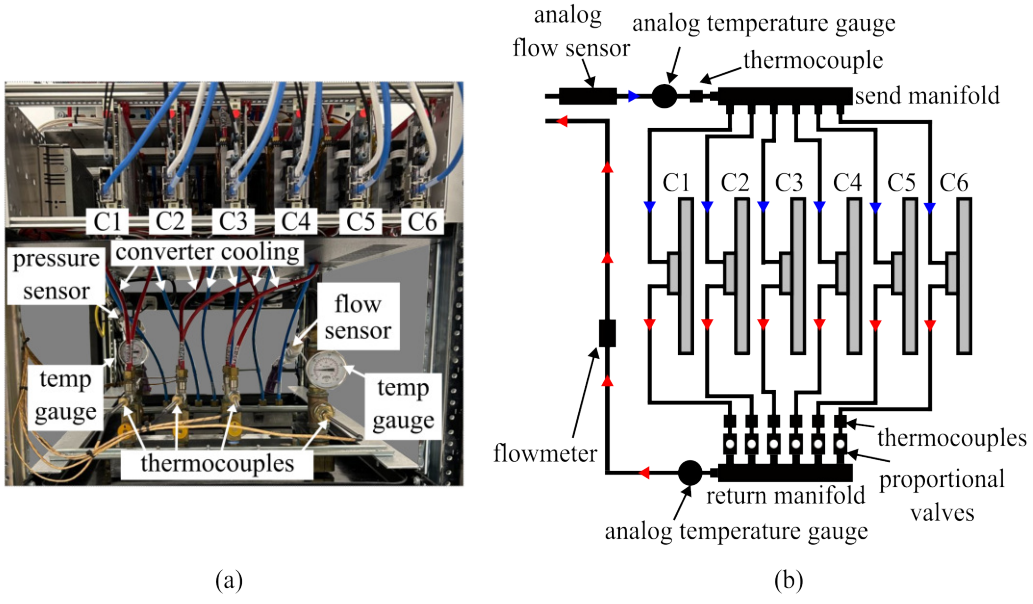

Figure 4. Electronic power converter liquid cooling, showing: (**a**) physical cooling system for power converters during testing, and; (**b**) power converting cooling system diagram.

2.3. Characterization of the Multi-Physics Representation

A multi-physics approach was used to capture the behavior of the testbed cooling system. The virtual representation couples three physical fields: electrical, fluid, and thermal. The electrical field maps the measured heater power input to the heat source applied to the loop. The fluid field represents coolant transport through the heater, power modules, radiator, and tank at the measured nominal flow rate of 3.785 L/min. The thermal field predicts the transient temperatures of the coolant and the effective solid thermal masses associated with these components. This decomposition was selected because the dominant experimental inputs are electrical power and coolant transport, whereas the dominant outputs are temperatures measured at four locations.

Each major cooling-loop element was represented as a lumped control volume governed by the energy balance

$$C_j \frac{dT_j}{dt} = \dot{m} c_p \left(T_{in,j} - T_j\right) + Q_j - \sum_k \frac{T_j - T_k}{R_{jk}} - (hA)_j \left(T_j - T_{amb}\right) \quad (1)$$

where C_j is the effective thermal capacitance of component j, $\dot{m}$ is the coolant mass flow rate, c_p is the coolant specific heat, $T_{in,j}$ is inlet coolant temperature, Q_j is internal heat generation, R_{jk} is the effective thermal resistance between adjacent control volumes, and $(hA)_j$ is the effective ambient heat rejection coefficient. For the heater, Q_j is prescribed by the measured electrical power supplied to the water heater. For the power modules, $Q_j = 0$ during the baseline experiment because the modules act as thermal masses rather than active heat sources. For the radiator and tank, the ambient heat rejection term dominates the sink behavior.

Before implementing the representation into the digital twin framework, it must be characterized to ensure that it faithfully represents the physical system. While the updating scheme can fit an inaccurate representation to a physical system, that approach reduces the accuracy of future predictions. To ensure the highest level of accuracy in the virtual representation, the cooling loop components were characterized and their parameters were determined using four categories: measurement, experiment, estimation, and online tuning. Measured parameters include geometric quantities and directly observed operating conditions, such as pipe lengths, pipe diameters, tank volume, coolant flow rate, and heater power input. Experimentally identified parameters were obtained by fitting the representation to recorded temperature data when direct measurements were insufficient to match the transient response. Estimated parameters include effective thermal resistance and ambient heat rejection terms when closed-form calculation would require assumptions not supported by the available hardware information. The subset selected for online updating consists of the parameters that are both weakly identifiable by direct measurement and strongly influential on transient temperature error. Table 1 summarizes the parameters of the lumped multi-physics representation, the associated component, the determination method, and whether each quantity was fixed, a fixed boundary condition (BC) or updated online.

Table 1. Parameters of the lumped multi-physics representation and identification method.

Parameter	Symbol	Component	Determination Method	Status
Pipe lengths	L_p	Piping network	Measured from hardware	Fixed
Pipe inner diameters	D_p	Piping network	Measured from hardware	Fixed
Tank volume	V_t	Tank	Measured from hardware	Fixed

Table 1. *Cont.*

Parameter	Symbol	Component	Determination Method	Status
Nominal coolant mass flow rate	$\dot{m}$	Cooling loop	Measured during operation; reference value in convective term	Fixed BC
Heater input power	Q_{h}	Heater	Measured from power supply; baseline heat input	Fixed BC
Ambient temperature	T_{amb}	Laboratory boundary	Measured during experiments	Fixed
Effective thermal capacitance	C_{h}	Heater	Lumped parameter used in $C_j\dot{T}_j$ term	Fixed
Effective thermal capacitance	C_{pm}	Power modules	Lumped parameter used in $C_j\dot{T}_j$ term	Fixed
Effective thermal capacitance	C_{r}	Radiator	Lumped parameter used in $C_j\dot{T}_j$ term	Fixed
Effective thermal capacitance	C_{t}	Tank	Lumped parameter used in $C_j\dot{T}_j$ term	Fixed
Thermal resistance	$R_{\mathrm{h,pm}}$	Heater to power modules	Effective resistance identified from inter-node heat transfer term	Tunable
Thermal resistance	$R_{\mathrm{pm,r}}$	Power modules to radiator	Effective resistance identified from inter-node heat transfer term	Tunable
Thermal resistance	$R_{\mathrm{r,t}}$	Radiator to tank	Effective resistance identified from inter-node heat transfer term	Tunable
Ambient heat rejection coefficient	$(hA)_{\mathrm{h}}$	Heater	Effective coefficient identified in ambient heat loss term	Tunable
Ambient heat rejection coefficient	$(hA)_{\mathrm{pm}}$	Power modules	Effective coefficient identified in ambient heat loss term	Tunable
Ambient heat rejection coefficient	$(hA)_{\mathrm{r}}$	Radiator	Effective coefficient identified in ambient heat loss term	Tunable
Ambient heat rejection coefficient	$(hA)_{\mathrm{t}}$	Tank	Effective coefficient identified in ambient heat loss term	Tunable
Power-module heat input	Q_{pm}	Power modules	Effective internal heat generation identified within Q_j term	Tunable
Effective coolant flow factor	$\alpha_{\dot{m},\mathrm{pm}}$	Power modules	Identified multiplier applied to $\dot{m}c_{\mathrm{p}}(T_{\mathrm{in},j} - T_j)$ term	Tunable

2.4. Digital Twin Updating Scheme

The objective of the updating scheme is to continuously recalibrate the bounded tunable parameter vector so that the multi-physics representation remains synchronized with the cooling-loop measurements. A new optimization cycle is triggered when a complete ten-minute window of sensor data becomes available. Let window k contain the temperature measurements from the heater, power modules, radiator, and tank sensors over that interval. The representation is not re-identified at every sample because each optimization cycle requires repeated evaluation of the full multi-physics model. Periodic updating was therefore adopted to provide a bounded optimization time, reduce sensitivity to instantaneous sensor noise, and prevent incomplete transients from dominating the search. Continuous updating could reduce latency only if the representation and optimizer were reformulated for streaming operation. That was not the objective of the present study.

Each particle position represents one candidate value of the tunable parameter vector, and the corresponding velocity governs how that candidate moves through the feasible set [30]. The particle position update is written as

$$X_i^{t+1} = V_i^{t+1} + X_i^t \tag{2}$$

and the velocity update is

$$V_i^{t+1} = WV_i^t + r_1\phi_1(P_i - X_i^t) + r_2\phi_2(P_g - X_i^t) \tag{3}$$

where W is the inertia factor, ϕ_1 and ϕ_2 are the cognitive and social acceleration coefficients, and r_1 and r_2 are random scalars sampled from $[0, 1]$. The feasible set is defined by lower and upper bounds on every entry of θ together with the physical constraints $C_j > 0$, $R_{jk} > 0$,

$(hA)_j \geq 0$, and $Q_{\text{pm}} \geq 0$. Consequently, the objective function is window-dependent cost defined below, and the constraints are the parameter bounds and positivity requirements that keep the search within physically admissible values.

The cost function is selected according to the mismatch between the measured and simulated temperatures at the beginning of each new data window. If the initial mismatch is large, a recovery window is triggered and the cost emphasizes state alignment at the start of the window,

$$w_{\text{recovery}} = \left\| \mathbf{T}_{\text{exp},1} - \mathbf{T}_{\text{sim},1} \right\|_2 \tag{4}$$

where $\mathbf{T}_{\text{exp},1}$ and $\mathbf{T}_{\text{sim},1}$ are the four-sensor temperature vectors at the first sample of the current window. If the initial mismatch is small, a calibrated window is used and the cost is the root mean squared error over the full window,

$$w_{\text{calibrated}} = \sqrt{\frac{1}{n_k} \sum_{i=1}^{n_k} \left\| \mathbf{T}_{\text{exp},i} - \mathbf{T}_{\text{sim},i} \right\|_2^2} \tag{5}$$

where n_k is the number of samples in the current window. A stall window combines both objectives when the initial mismatch and the full-window error are simultaneously non-negligible,

$$w_{\text{stall}} = w_{\text{recovery}} + w_{\text{calibrated}} \tag{6}$$

The selected window type determines both the cost function and the level of particle exploration. Recovery windows emphasize rapid state correction, calibrated windows emphasize fine parameter refinement, and stall windows combine both behaviors to avoid premature stagnation near the previous optimum as seen in Figure 5.

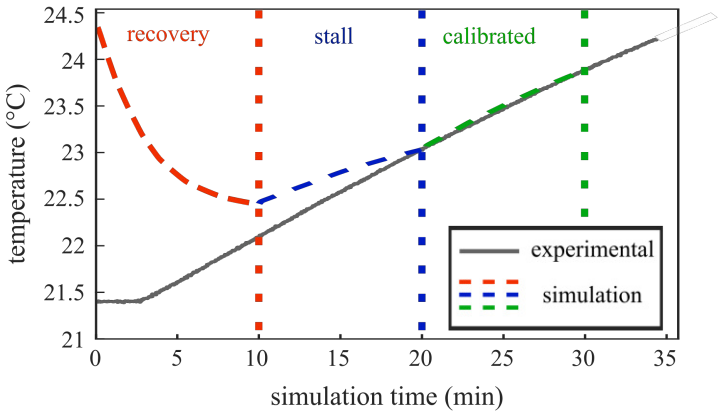

Figure 5. The edge-resident particle swarm updating the hyper-parameters according to the performance of the previous window.

Choosing an appropriate window size is paramount for the performance of the updating scheme. If the window is too small, the updating scheme may not have enough time to converge before the next window becomes available. If the window is too large, the latency between the physical system and the virtual representation reduces the ability of the digital twin to adapt to dynamic changes in the physical system. To determine the appropriate update interval, the simulation displayed in Figure 6 was run 10,000 times on the same five-minute window. The average runtime of the simulation was 4.636 s, with a standard deviation of 0.634 s. The maximum runtime was 16.925 s, but a runtime of 16 s or more occurred only twice out of 10,000 instances. As a result of this investigation, a ten-minute window was chosen.

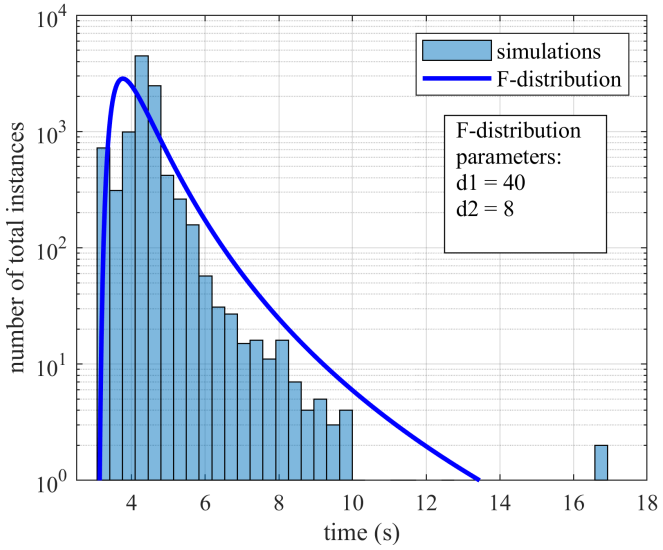

Figure 6. Timing distribution of 10,000 model instances with F-distribution fit.

Using the measured average runtime, one ten-minute window supports approximately $600/4.636 \approx 129$ serial model evaluations. For a swarm with N_p particles, this corresponds to approximately $\lfloor 129/N_p \rfloor$ full particle-swarm iterations per window on average. As a result, a 10-particle swarm allows about 12 full iterations, a 20-particle swarm about 6, and a 30-particle swarm about 4. This runtime budget bounds the practical number of configurable parameters that can be processed in real time. In the present implementation, the 12-parameter vector remained computationally feasible within the ten-minute deadline, whereas larger swarms would reduce the number of full iterations available before the next data window arrives. Small swarms reduce search diversity and can delay recovery after a discontinuity, while very large swarms increase computational cost without a proportional reduction in window error.

The function of the updating algorithm is shown in Figure 7. The algorithm starts by initializing a swarm of particles with random positions and velocities. Once enough data is collected to fill a window, the main loop begins. The same dissipated heat load profile deployed onto the water heater, measured in watts, is also applied to the digital twin, linking the virtual representation to the physical system. After the load profile is configured, the PSO update loop begins. The values of the representation parameters are tuned using the position of each particle, and the configured virtual representation is run and evaluated against the current sensor window using the cost function associated with that window type. When the global-best calibrated-window root mean squared error

$$\text{Cost} = \text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(T_{\text{sim},i} - T_{\text{exp},i}\right)^2} \tag{7}$$

is less than 0.1 °C, the representation is flagged as calibrated for that window. When a new window is acquired, the global best solution from the previous window is retained, the acceleration factors are reset, and an additive perturbation sampled from $[-1, 1]$ is applied to each particle position to prevent entrapment at the previous optimum before the next window begins.

The manually tuned model is the static baseline representation. It was identified once offline from the 40-h experimental data and then held fixed unless the updating scheme was activated. The online updating scheme therefore does not replace the baseline model. It starts from that manually tuned representation and recalibrates the tunable subset of parameters whenever a new sensor window becomes available.

The pertinent values used to initialize and govern the PSO algorithm are given in Table 2, and have been included in this manuscript to allow for reproducibility. Particle positions were randomly initialized using a uniform distribution within the ranges listed in Table 2. All parameters were constrained to the global bounds [0, 8], and the optimization objective was to minimize the temperature–pressure RMSE between simulated and measured system responses.

Table 2. Particle Swarm Optimization (PSO) hyperparameters and effective model parameters.

Parameter	Symbol/Range	Description
Swarm size	$n_{\text{pop}} = 5$	Number of particles
Maximum iterations	$N_{\text{it}} = 5$	Optimization iterations
Inertia weight	$w = 1.0$	Initial inertia coefficient
Inertia damping	$w_d = 0.99$	Inertia decay per iteration
Cognitive coefficient	$c_1 = 0.5$	Personal best influence
Social coefficient	$c_2 = 0.7$	Global best influence
Velocity limit	± 1.6	20% of variable range
Radiator–tank thermal resistance	$R_{\text{r,t}} \in [0, 3]$	Inter-node heat transfer term $\sum_k (T_j - T_k)/R_{jk}$
Heater effective heat input	$Q_h^{\text{eff}} \in [0, 1]$	Internal heat generation term Q_j
Coolant flow effectiveness factor	$\alpha_{\dot{m},\text{pm}} \in [0, 2]$	Convective transport term $\dot{m} c_p (T_{\text{in},j} - T_j)$
Radiator ambient heat rejection	$(hA)_r \in [0, 1]$	Ambient loss term $(hA)_j (T_j - T_{\text{amb}})$

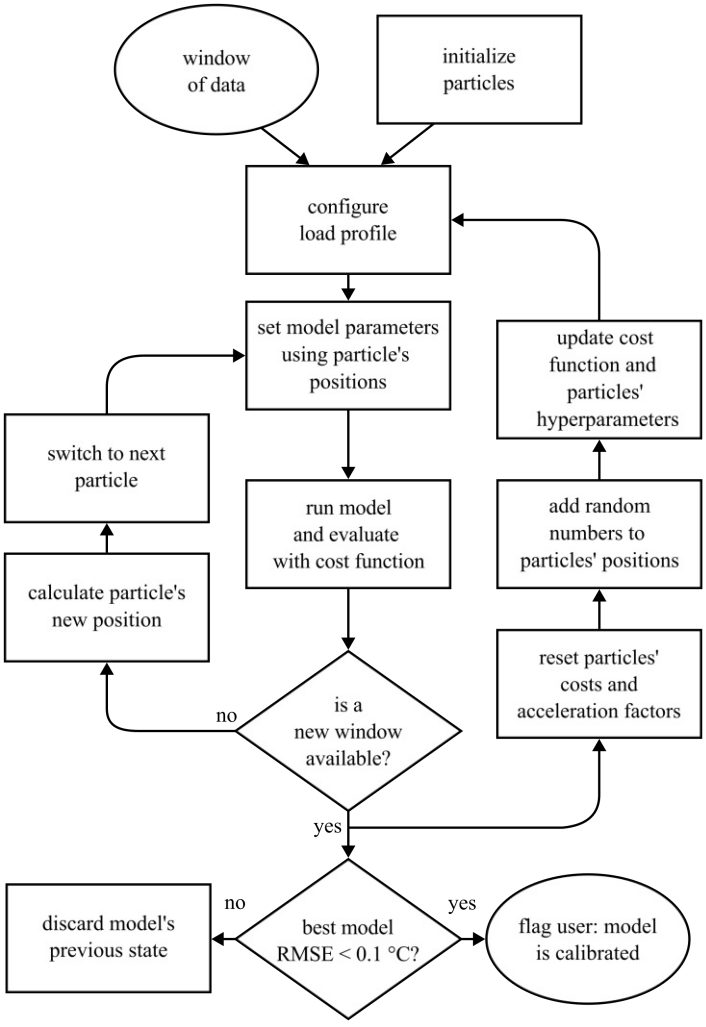

Figure 7. Flowchart of the multi-physics virtual representation updating scheme.

2.5. Test Scenario

An arbitrary ten-hour test scenario was deployed onto the testbed to gather data for characterization of the virtual representation. The scenario was created to replicate a battle-time scenario while keeping in mind the limitations of the testbed cooling system. Many ship systems were considered while designing a load profile. The electrical loads of the systems were calculated, and dissipated heat was dumped into the system via the power supply and water heater. However, not every system was modeled in the same way. Hotel and communication systems are air-cooled and are only modeled electrically. The water-cooled, propulsion, navigational, and directed energy systems are modeled electrically and thermally. The propulsion system is throttled throughout the test and represents the majority of the load for the duration of the test. The navigational, communication, and radar systems have a constant load throughout the test. Unlike these systems, hotel and directed energy systems are applied intermittently, depending on the time. The load of each system is combined to create a ten-hour load profile deployed onto the power supply Figure 8.

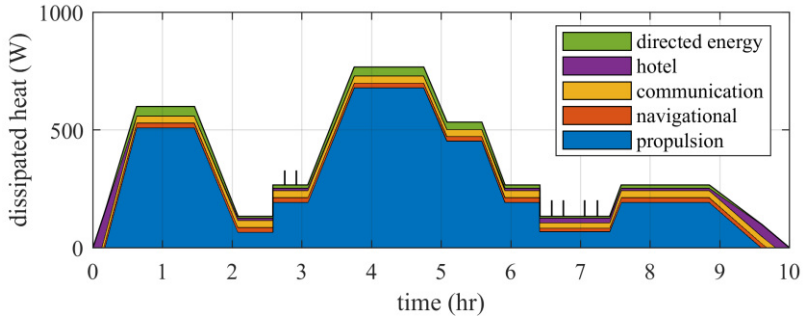

Figure 8. Distribution of dissipated heat per ship system during a ten-hour scenario.

The ten-hour load profile was repeated four times, creating a 40-h experiment; the results are shown in Figure 9. During the experiment, temperature readings were taken at four points along the cooling system (heater, power modules, radiator, and tank). As previously mentioned in Section 2.3, experimental data was used to fit the virtual representation. Results of this manually tuned representation are shown in Figure 10 and its performance metrics in Table 3. The simulation results show that the manually tuned virtual representation is relatively accurate.

One significant benefit of a calibrated representation is that it can be leveraged to conduct test scenarios without physically altering the testbed. In later sections, this manually tuned virtual representation will synthesize experimental data to assess the performance of the updating scheme when handling discontinuities.

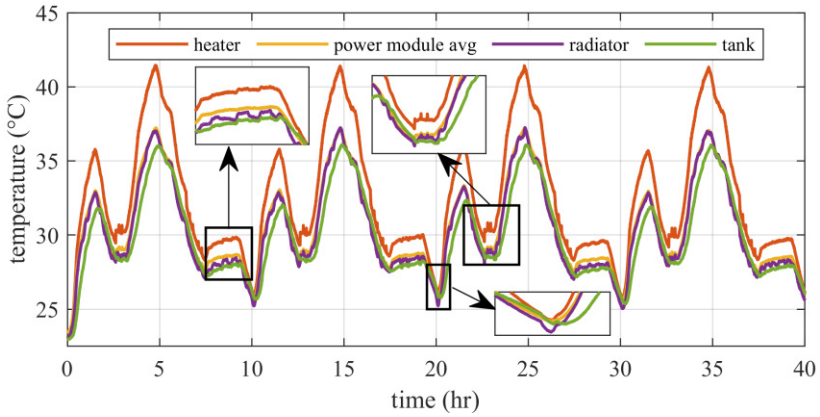

Figure 9. Recorded temperature at the tank, heater, power modules, and radiator, during the 40-h experiment.

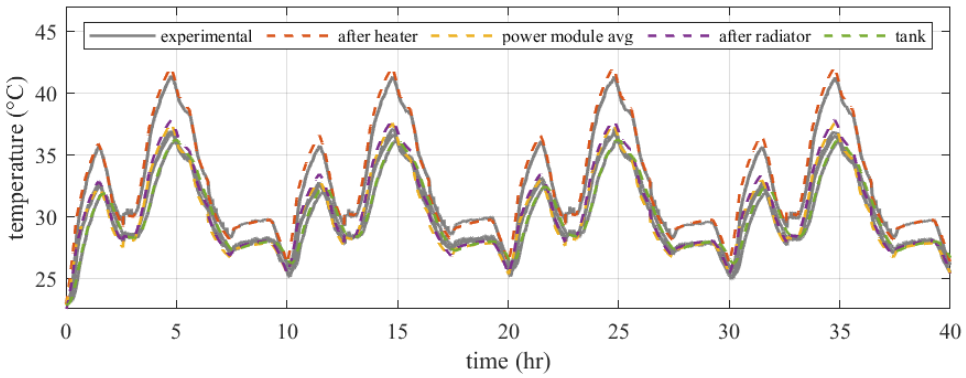

Figure 10. Manually tuned virtual representation simulation results of the 40-h experiment compared to the experimentally gathered data.

Table 3. Metric results for the manually tuned virtual representation at four points, numerical case study.

Metric	After Heater	Average of Power Modules	After Radiator	Inside Tank
mean absolute error (°C)	0.500	0.420	0.447	0.249
mean squared error (°C)	0.501	0.260	0.315	0.101
root mean squared error (°C)	0.708	0.511	0.562	0.318
normalized root mean squared error	0.030	0.028	0.032	0.012

3. Results and Discussion

This section investigates the robustness of the proposed digital twin updating scheme and its ability to handle physical discontinuities.

3.1. Updating Scheme Investigation

To understand how much the updating scheme improves the virtual representation, it was directly compared against a representation without updating in Figure 11. The virtual representation was run for the first three and a half hours of the experiment without updating any parameters. At the three-and-a-half-hour mark, the representation parameters were randomized; this was done to determine if the updating scheme could recover from a bad initial state. The randomization of the virtual representation parameters causes a drastic increase in the simulation temperature. After two ten-minute windows of data, the updating scheme improves the virtual representation and returns it to its initial accuracy. After these two windows, the accuracy of the updating virtual representation continues to improve, exceeding the performance of the static representation, best shown by RMSE percent improvement in Table 4 and by comparison in Figure 12.

Table 4. Metric results for a virtual representation continuously updated at four locations, numerical case study.

Metric	After Heater	Average of Power Modules	After Radiator	Inside Tank
mean absolute error (°C)	0.032	0.016	0.048	0.014
mean square error (°C)	0.003	0.001	0.005	0.001
root mean squared error (°C)	0.051	0.024	0.024	0.021
normalized root mean squared error	0.003	0.001	0.005	0.001
percent improvement RMSE (°C)	93%	95%	96%	93%

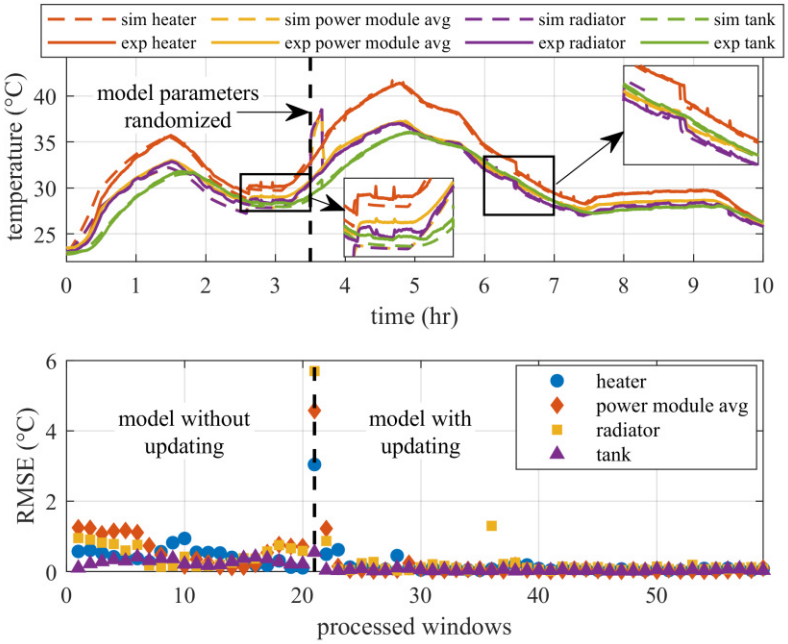

Figure 11. Virtual representation without updating compared to representation with updating scheme with parameter randomization mid-test.

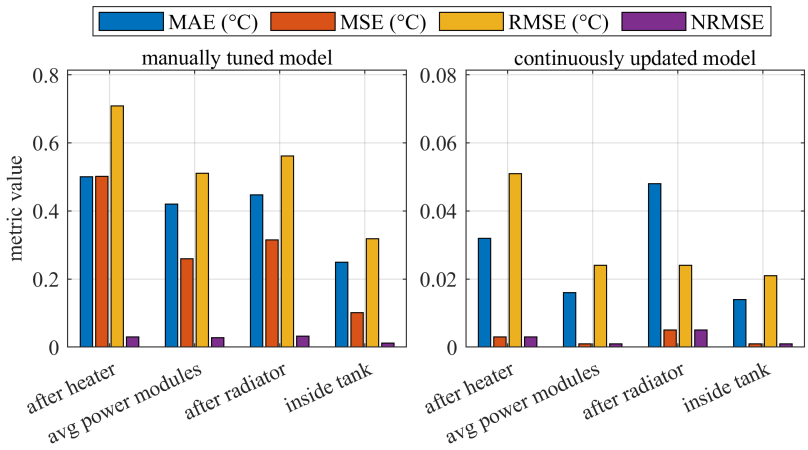

Figure 12. Comparison of virtual representation metrics with and without an automatic representation updating scheme.

3.2. Test Case Scenarios

The baseline multi-physics representation was first validated against the 40-h experimental data shown in Figures 9 and 10 and Table 3. The two discontinuity scenarios in this subsection were then synthesized from that experimentally validated baseline. Synthetic scenarios were used intentionally so that the timing and magnitude of each imposed change were known exactly and the updating scheme could be evaluated under controlled conditions without introducing confounding hardware changes.

Two scenarios were deployed to test the adaptability of the updating scheme in handling discontinuities in the physical system. In each scenario, the static representation was run for the first 3.5 h. A physical change was then introduced to alter the system behavior, and the representation updating scheme began evaluating permutations of the tunable parameter vector to recover the new operating condition.

In the first scenario, the power modules begin dissipating heat into the cooling loop at 3.5 h. The results from this scenario are shown in Figure 13. The temperature rise after the discontinuity is initially underpredicted because the pre-change parameter set

assumes that the modules act only as thermal masses. After the first few windows, the updating scheme reduces the mismatch and recovers the new operating condition. The temporary discrepancy near the 4.5 h mark is attributable to stochastic exploration between consecutive windows. In the present implementation, an additive perturbation is applied to the particle positions at each window transition to avoid entrapment near the previous optimum. When that perturbation is large relative to the local sensitivity of the radiator states, it can force needless exploration and produce a short-lived temperature mismatch. A practical remedy is to scale the inter-window perturbation with the previous-window error so that larger perturbations are used only during recovery windows and smaller perturbations are used during calibrated windows. A second remedy is to restrict particle re-initialization to a reduced neighborhood around the global best solution after repeated evidence of convergence.

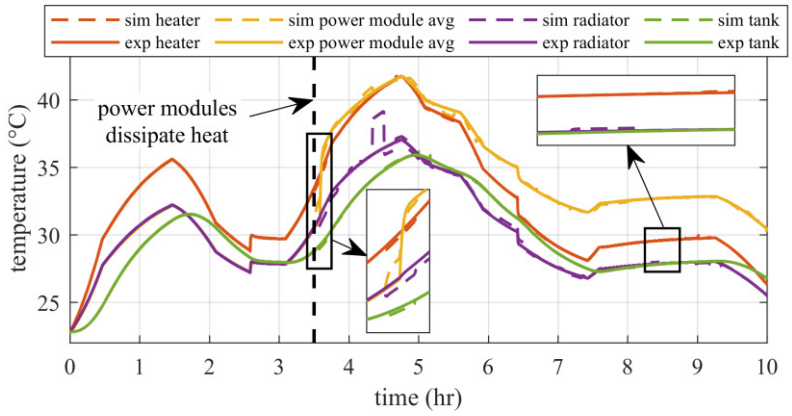

Figure 13. Scenario 1: Power modules dissipate heat into the cooling loop.

The second scenario, shown in Figure 14, removes insulation from the tank. Again, the updating scheme recovers after a few windows. However, the largest physical change occurs at the tank, while the largest early prediction error appears at the heater. This occurs because the cost function aggregates error across all four measurement locations and the upstream heater state responds first to the altered loop energy balance. During the first recovery windows, the swarm can therefore find parameter sets that reduce the global cost while still misallocating part of the disturbance between tank and heater losses. Once the recovery window is triggered, the added emphasis on initial-state mismatch drives the swarm away from that local minimum and toward a parameter set consistent with the new boundary condition. This result shows that the framework can recover from a structural thermal change, but it also indicates that a scalar aggregate cost function can temporarily mask the location of the disturbance.

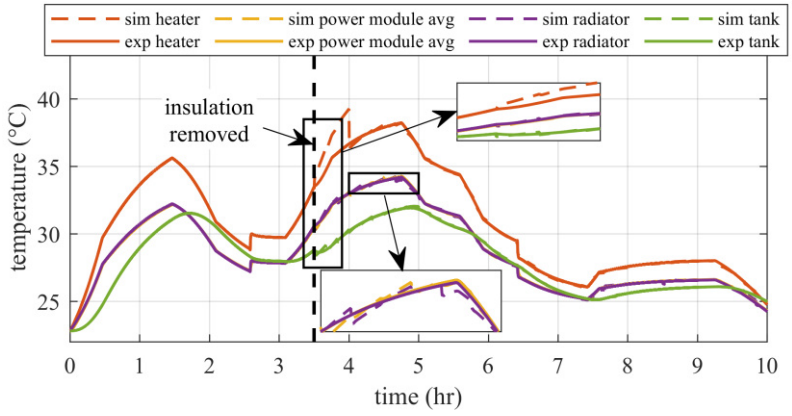

Figure 14. Scenario 2: Insulation removed from the tank.

Although the present study focuses on temperature tracking rather than closed-loop control or automated fault isolation, the observed reduction in root mean squared error has direct operational implications. A representation that remains calibrated after a discontinuity can support supervisory actions such as thermal margin assessment, load scheduling, coolant temperature management, and early recognition of abnormal thermal resistance or heat rejection. Scenario 1 corresponds to an unmodeled internal heat source, while Scenario 2 corresponds to degraded insulation and increased environmental heat loss. In both cases, the updating scheme restores agreement between the model and the sensor data within a few windows, which is a necessary precursor to using the digital twin for optimization or fault detection. A dedicated fault classifier or control optimizer was not implemented in this study, so these operational benefits should be interpreted as enabled capabilities rather than experimentally validated closed-loop functions.

This study is bounded by the scope of the experimental platform and the intended objective of online thermal-model calibration. The liquid-cooled testbed captures the dominant thermal interactions needed to evaluate the proposed digital twin framework under controlled transient conditions, although larger installations may include additional components and coupling effects. Temperature measurements were used as the primary updating signal, while flow rate and boundary conditions were prescribed from the experiment, which is consistent with the focus on thermal-state agreement. The scalar cost function and ten-minute update window were selected to balance calibration accuracy, numerical stability, and computational feasibility for edge implementation. The discontinuity scenarios in this subsection were derived from the validated baseline representation so that the imposed changes were known exactly and the recovery behavior of the updating scheme could be assessed in a repeatable manner. These modeling and experimental choices are appropriate for the present study and provide a clear basis for future extension to broader hardware conditions and larger thermal management architectures.

4. Conclusions

As modern systems naval and defense become increasingly complex, maintaining accurate and up-to-date virtual representations of these systems poses a significant challenge. Throughout the lifecycle of a system, discrepancies between itself and its virtual counterpart can lead to reduced performance and inaccuracies. To address this, physics-based approaches have emerged as an option for enabling real-time virtual representation updating. This paper investigates the application of a particle swarm optimization framework to enhance the accuracy of a digital twin through proper tuning. The digital twin framework comprises an updating scheme, a physical testbed designed to emulate a cooling system of a ship, and a multi-physics virtual representation of that system. The updating scheme leverages a swarm of particles and online sensor data to evaluate permutations of representation parameters. Two scenarios were considered to evaluate the performance of the digital twin framework where a physical change in the cooling system would occur. The methods developed, tested, and discussed in this paper provide a foundation for the development of digital twin updating frameworks using particle swarm optimization. Results demonstrate that the digital twin framework can adapt to system changes across ten-minute windows, and improve the accuracy of its static virtual representation by more than 90% when measured at four points of the system under test. This allows for the application of accurate digital twins in conditions with rapidly changing parameters which would quickly render a model without updating obsolete. This method additionally provides adaptation to account for degradation in systems with significant wear and aging effects.

Author Contributions: Data curation, J.W.; Formal analysis, K.S., A.R.J.D. and J.K.; Supervision, K.S., A.R.J.D. and J.K.; Writing—original draft, B.P.; Writing—review & editing, J.W., K.S., R.H., A.R.J.D., J.K. and K.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Office of Naval Research under contracts N00014-20-C-1106, N00014-22-C-1003, and N00014-23-C-1012.

Data Availability Statement: The datasets generated and/or analyzed during the current study are not publicly available due to funding resource but are available from the corresponding author on reasonable request.

Acknowledgments: The support of the Office of Naval Research is gratefully acknowledged. Approved, DCN# 2026-4-2-2155. DISTRIBUTION STATEMENT A. Approved for public release: distribution is unlimited. Approved: 24 April 2026. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Navy. The views expressed are those of the author and do not reflect the official policy or position of the Department of War or the U.S. Government.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

CNC	Computer Numerical Control
HVAC	Heating, Ventilation, and Air Conditioning
MAE	Mean Absolute Error
MSE	Mean Squared Error
NAE	National Academy of Engineering
NRMSE	Normalized Root Mean Squared Error
OS	Operating System
PSO	Particle Swarm Optimization
RMSE	Root Mean Squared Error

References

1. Guitton, M. Fighting the locusts: Implementing military countermeasures against drones and drone swarms. *Scand. J. Mil. Stud.* **2021**, *4*, 26–36. [CrossRef]
2. Sado, K.; Hannum, J.; Skinner, E.; Ginn, H.L.; Booth, K. Hierarchical digital twin of a naval power system. In *2023 IEEE Energy Conversion Congress and Exposition (ECCE)*; IEEE: Washington, DC, USA, 2023; pp. 1514–1521.
3. Sulligoi, G.; Vicenzutti, A.; Menis, R. All-electric ship design: From electrical propulsion to integrated electrical and electronic power systems. *IEEE Trans. Transp. Electrif.* **2016**, *2*, 507–521. [CrossRef]
4. Monti, A.; Boroyevich, D.; Cartes, D.; Dougal, R.; Ginn, H.; Monnat, G.; Pekarek, S.; Ponci, F.; Santi, E.; Sudhoff, S.; et al. Ship power system control: A technology assessment. In *IEEE Electric Ship Technologies Symposium, 2005*; IEEE: Philadelphia, PA, USA, 2005; pp. 292–297.
5. O'Rourke, R. *Navy Lasers, Railgun, and Gun-Launched Guided Projectile: Background and Issues for Congress*; Report No. R44175; Congressional Research Service: Washington, DC, USA, 2022. Available online: https://crsreports.congress.gov/product/pdf/R/R44175 (accessed on 13 May 2026).
6. Booth, K.; Sado, K.; Hannum, J.; Knight, J.; Dougal, R. Introduction of posture-based pre-alignment for naval applications. In *2023 IEEE Electric Ship Technologies Symposium (ESTS)*; IEEE: Washington, DC, USA, 2023; pp. 464–468.
7. Abbaspour, A.; Mokhtari, S.; Sargolzaei, A.; Yen, K.K. A survey on active fault-tolerant control systems. *Electronics* **2020**, *9*, 1513. [CrossRef]
8. Naval Research Advisory Committee. *How Autonomy Can Transform Naval Operations*; Arlington, VA, USA, 2013. Available online: https://apps.dtic.mil/sti/pdfs/ADA591983.pdf (accessed on 13 May 2026).
9. Office of Naval Research. *Naval Power and Energy Systems Technology Development Roadmap*; Department of the Navy: Arlington, VA, USA, 2019. Available online: https://apps.dtic.mil/sti/pdfs/AD1093333.pdf (accessed on 13 May 2026).
10. National Academy of Engineering and National Academies of Sciences, Engineering, and Medicine. *Foundational Research Gaps and Future Directions for Digital Twins*; The National Academies Press: Washington, DC, USA, 2023.

11. Sado, K.; Hannum, J.; Booth, K. Digital twin modeling of power electronic converters. In *2023 IEEE Electric Ship Technologies Symposium (ESTS)*; IEEE: Washington, DC, USA, 2023; pp. 86–90.
12. Wilsdon, K.; Hansel, J.; Kunz, M.R.; Browning, J. Autonomous control of heat pipes through digital twins: Application to fission batteries. *Prog. Nucl. Energy* **2023**, *163*, 104813. [CrossRef]
13. Tao, F.; Zhang, H.; Liu, A.; Nee, A.Y.C. Digital twin in industry: State-of-the-art. *IEEE Trans. Ind. Inform.* **2019**, *15*, 2405–2415. [CrossRef]
14. Tao, F.; Xiao, B.; Qi, Q.; Cheng, J.; Ji, P. Digital twin modeling. *J. Manuf. Syst.* **2022**, *64*, 372–389. [CrossRef]
15. Alves de Araujo Junior, C.A.; Villanueva, J.M.; Almeida, R.J.S.; Azevedo de Medeiros, I.E. Digital twins of the water cooling system in a power plant based on fuzzy logic. *Sensors* **2021**, *21*, 6737. [CrossRef] [PubMed]
16. Matsuoka, K.; Hill, D. Online optimization of cooling water system in a district cooling plant by using digital twin. *ASHRAE Trans.* **2020**, *126*, 427.
17. Aguilera, J.J.; Meesenburg, W.; Markussen, W.B.; Zühlsdorf, B.; Elmegaard, B. Real-time monitoring and optimization of a large-scale heat pump prone to fouling—Towards a digital twin framework. *Appl. Energy* **2024**, *365*, 123274. [CrossRef]
18. Peng, C.-C.; Su, C.-Y. Modeling and parameter identification of a cooling fan for online monitoring. *IEEE Trans. Instrum. Meas.* **2021**, *70*, 1–14. [CrossRef]
19. Worch, J.; Sado, K.; Downey, A.R.; Khan, J.; Santi, E. Real-time blockage detection and autonomous recovery in liquid-cooled systems using digital twins. In *2025 IEEE Electric Ship Technologies Symposium (ESTS)*; IEEE: Alexandria, VA, USA, 2025; pp. 543–549.
20. Abo-Khalil, A.G. Digital twin real-time hybrid simulation platform for power system stability. *Case Stud. Therm. Eng.* **2023**, *49*, 103237. [CrossRef]
21. Duran-Limon, H.A.; Chavoya, A.; Hernández-Ochoa, M. The role of machine learning in big data analytics: Current practices and challenges. In *Development Methodologies for Big Data Analytics Systems*; Mora, M., Wang, F., Marx Gomez, J., Duran-Limon, H., Eds.; Springer: Cham, Switzerland, 2024; pp. 47–74.
22. Yin, P.-Y.; Yu, S.-S.; Wang, P.-P.; Wang, Y.-T. A hybrid particle swarm optimization algorithm for optimal task assignment in distributed systems. *Comput. Stand. Interfaces* **2006**, *28*, 441–450. [CrossRef]
23. Ab Wahab, M.N.; Nefti-Meziani, S.; Atyabi, A. A comprehensive review of swarm optimization algorithms. *PLoS ONE* **2015**, *10*, e0122827. [CrossRef] [PubMed]
24. Zohdi, T.I. A digital-twin and machine-learning framework for precise heat and energy management of data-centers. *Comput. Mech.* **2022**, *69*, 1501–1516. [CrossRef]
25. Lu, Q.; Zhu, D.; Wang, M.; Li, M. Digital twin-driven thermal error prediction for CNC machine tool spindle. *Lubricants* **2023**, *11*, 219. [CrossRef]
26. Mohseni, S.-R.; Zeitouni, M.J.; Parvaresh, A.; Abrazeh, S.; Gheisarnejad, M.; Khooban, M.-H. FMI real-time co-simulation-based machine deep learning control of HVAC systems in smart buildings: Digital-twins technology. *Trans. Inst. Meas. Control* **2023**, *45*, 661–673. [CrossRef]
27. Antonopoulos, A.; Karakostas, B.; Katsoulakos, T.; Mavrakos, A.; Tsaousis, T.; Zavvos, S. A digital twin enabled decision support framework for ship operational optimisation towards decarbonisation. In *Proceedings of Eighth International Congress on Information and Communication Technology*; Yang, X.-S., Sherratt, R.S., Dey, N., Joshi, A., Eds.; Springer Nature: Singapore, 2023; pp. 467–476.
28. Cronin, J.; Santi, E.; Wunderlich, A.; Knight, J. Fast system level model for digital twin based optimization of naval power and energy system. In *IEEE Electric Ship Technologies Symposium (ESTS)*; IEEE: Alexandria, VA, USA, 2023; pp. 1–7.
29. Imperix Datasheet. PEB 8038 Half-Bridge SiC Power Module. Imperix Datasheet. 2023. Available online: https://imperix.com/products/power/half-bridge-module/ (accessed on 1 February 2024).
30. Kennedy, J.; Eberhart, R. Particle swarm optimization. In *Proceedings of ICNN'95—International Conference on Neural Networks*; IEEE: Perth, WA, Australia, 1995; Volume 4, pp. 1942–1948.