MITSUBISHI ELECTRIC RESEARCH LABORATORIES
https://www.merl.com

Digital Twins for Vapor Compression Cycles: Challenges &Opportunities

Laughman, Christopher R.; Deshpande, Vedang M.; Qiao, Hongtao; Bortoff, Scott A.; Chakrabarty, Ankush

TR2023-103 August 23, 2023

Abstract

Digital twins are a promising technology for vapor compression cycles because their capabilities can enable new approaches to system analysis, design, and control. As the computational models utilized by these digital twins must capture the observed dynamics of these physical systems, we describe important characteristics of these systems that impact the structure and implementation of these models. Specific attributes of these systems that govern model development include large-scale structures with tens or hundreds of thousands of equations, time constants that range over 10 orders of magnitude, and derivative discontinuities that affect the performance of solvers. We then describe candidate modeling, calibration, and estimation techniques that leverage modern mathematical and computational methods to meet these requirements, and present use cases demonstrating their efficacy.

International Congress of Refrigeration (ICR) 2023

Mitsubishi Electric Research Laboratories, Inc.
201 Broadway, Cambridge, Massachusetts 02139

DOI: 10.18462/iir.icr.2023.0765

Digital Twins for Vapor Compression Cycles: Challenges & Opportunities

Christopher R. LAUGHMAN*, Vedang DESHPANDE, Hongtao QIAO, Scott A. BORTOFF, Ankush CHAKRABARTY

Mitsubishi Electric Research Laboratories
Cambridge, MA, USA
{laughman, deshpande, qiao, bortoff, chakrabarty}@merl.com

ABSTRACT

Digital twins are a promising technology for vapor compression cycles because their capabilities can enable new approaches to system analysis, design, and control. As the computational models utilized by these digital twins must capture the observed dynamics of these physical systems, we describe important characteristics of these systems that impact the structure and implementation of these models. Specific attributes of these systems that govern model development include large-scale structures with tens or hundreds of thousands of equations, time constants that range over 10 orders of magnitude, and derivative discontinuities that affect the performance of solvers. We then describe candidate modeling, calibration, and estimation techniques that leverage modern mathematical and computational methods to meet these requirements, and present use cases demonstrating their efficacy.

Keywords: Digital twin, Vapor compression cycles, Simulation, Calibration, Estimation, Control, Modelica, Differential algebraic equations

1. INTRODUCTION

A host of concerns related to climate impact, energy efficiency, and comfort performance are converging to make this a time of challenge and opportunity for vapor compression cycles, as one of the key technologies used in today's building energy systems. While the energy efficiency of buildings is a perennial concern, the climate-related implications of traditional carbon-based energy sources and refrigerants with a high global warming impact have amplified the need for more sustainable approaches to space heating and cooling. New technologies that enable high levels of system integration promise to increase electrification without sacrificing equipment performance, but these methods rely on a high bandwidth of information exchange between equipment, the building envelope, and the electrical grid. As greater system integration is enabled by exposing more granular performance information that can be used to develop advanced control and optimization methods, system architectures for vapor compression cycles that can produce a wide range of data hold significant promise in the pursuit of sustainable design objectives.

Digital twins are emerging as a potent new technology for enabling such new capabilities for multiphysical systems. While the phrase "digital twin" has been rapidly adopted in the marketing domain and is consequently attached to a variety of overlapping meanings, one helpful definition proposed by the AIAA Digital Engineering Integration Committee (2020) is that a digital twin is "a set of virtual information constructs that mimics the structure, context and behavior of an individual / unique physical asset, or a group of physical assets, is dynamically updated with data from its physical twin throughout its life cycle, and informs decisions that realize value." As simulation-oriented representations that track the behavior of physical systems, these tools are amenable to a wider range of performance monitoring and control analyses than is possible with strictly physical sensing approaches.

As data can be used in virtual representations of system behavior in a variety of modalities, we explicitly distinguish between a *digital model*, which is a virtual representation that can be used to predict the behavior of the physical system but does not itself change in response to measurement data, and a *digital twin*, in which the output of the virtual representation is dependent upon current and past data obtained from the physical system by automatically evolving in response to measurement data. This evolution ensures that the output of the digital twin automatically tracks the output of the physical system and enables actionable decisions to be made for the physical system on the basis of the behavior of the virtual representation.

Thelen et al. (2022) provides an excellent review of digital twin technology across a wide variety of domains. More specialized treatments of digital twins in the building domain are covered by Das et al. (2022), who studies the application of machine-learning technology to digital twins of smart buildings, Vering et al. (2019), who investigates the use of digital twins in the design phase of building systems, and Xie et al. (2023), who applies digital twins to fault detection and diagnosis for HVAC systems.

Whereas this prior work is focused almost exclusively on the application of digital twins to building-level behavior, digital twin technologies stand to provide important benefits to vapor compression cycles by providing a variety of new applications and capabilities, and by enabling computational models of these systems to adapt to the characteristics of specific physical assets which may vary across a fleet of equipment. For example, accurate digital twins may facilitate the development of advanced performance monitoring applications of such quantities as the total system refrigerant mass without being subject to conventional sensor limitations. This technology could also be used in model predictive control applications to optimize the constrained closed-loop behavior of buildings and equipment by using accurate predictions of the system performance over long time horizons to determine control strategies. Digital twins can thus be used to robustly integrate vapor compression cycles with other smart building/grid applications by leveraging their enhanced predictive capabilities.

While the value of digital twins for vapor compression cycles is apparent, the complexity of these physical systems suggests a need to characterize the requirements on digital twins to ensure that their virtual representation can accurately and parsimoniously describe the behavior of the physical system. We are thus motivated to understand the requirements imposed upon of virtual representations of vapor compression cycles, and seek to identify a set of candidate approaches that have these properties.

The structure of this paper is as follows: Section 2 describes fundamental characteristics of vapor compression cycles that circumscribe the behavior of the models used in digital twins. Section 3 then surveys candidate approaches that satisfy these requirements and enable the digital twin behavior to track that of the physical system as it changes over time, after which Section 4 enumerates two use cases that demonstrate a prototype application of these methods in practice. Section 5 concludes the paper with a brief summary and pointers to directions for future research.

2. CHALLENGES

Computational models are the core of any digital twin, and must have two attributes to be successful: they must describe the expected behavior of the corresponding physical system, and adapt to changes in the physical system over time. We therefore turn our attention first to the numerical and computational properties of vapor compression cycle models that are required to describe experimentally observed behavior. We group the essential traits of these cycles into the following 4 categories: the systems exhibit dynamic behavior, are data rich but information poor, have large-scale structure, and possess nonlinear characteristics.

2.1. Dynamic behavior

All vapor compression cycles exhibit dynamic behavior during operation. While steady-state models are valuable during equipment design and specification processes because the models are often simple and can be simulated rapidly, their predictions do not characterize the transient behavior of equipment in real-world scenarios. State-of-the-art equipment relies upon automatic control systems to regulate variable actuators to achieve superior part-load system operating efficiency, but even simple window-mounted room air-

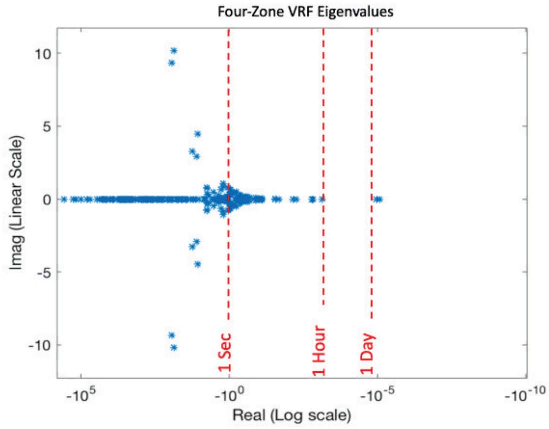

Figure 1: Eigenvalues of the Jacobian for a four-zone VRF system, with the real-axis on a logarithmic scale.

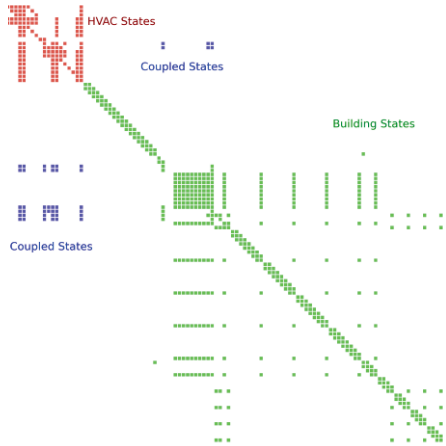

Figure 2: Jacobian sparsity pattern for a system model of a vapor compression cycle in a building (Chakrabarty et al., 2021).

conditioners exhibit many dynamic phenomena due to disturbances and the effects of on/off cycling for temperature control.

These systems also exhibit coupled multi-input, multi-output dynamic behavior due to the multiphysical interactions that take place in the equipment. These system-level behaviors are most accurately described in a computational representation as sets of nonlinear differential algebraic equations (DAEs), resulting from the discretization of the partial differential equations (PDEs) governing multiphase fluid flow and conjugate heat transfer. The hygrothermal behavior of the building is also described by a distinct set of DAEs that can be coupled to the cycle through fluid flow and heat transfer. The resulting system behavior can then be studied by solving the coupled cycle/building system forward over the time horizon of interest (Li, 2014b).

2.2. Measurements: Data rich but Information poor

Although the continuing growth and expansion of the IoT market has provided new resources for monitoring system behavior in the field, most vapor compression cycles typically only operate over a limited set of conditions and do not produce sufficient information needed to characterize system operation for all scenarios of interest. As it is incumbent upon equipment manufacturers to guarantee performance over all practical operating conditions, the robustness of these models to potential behavior is paramount. For example, accurate predictions of cycle behavior during a heat wave may be valuable when planning demand response events and load curtailment procedures, but the infrequent occurrence of such high temperature events suggests that information about them will be relatively scarce.

Physics-based modeling approaches are often advantageous in such information-poor scenarios because of their established extrapolative abilities and generalization properties (Bhattacharya et al., 2022), which ensure high predictive quality. While data describing the observed performance of these systems is essential to overcome modeling limitations due to various sources of uncertainty, hybrid approaches that combine physics- and data-based approaches often yield more practical and robust results with contemporary tools than purely data-driven techniques. Methods to calibrate and automatically adapt the structure of system models to account for uncertainties associated with the installation process will furthermore serve to constructively blend both physics- and data-based methods (Li, 2014a).

2.3. Large-scale structure

Contemporary HVAC systems built on vapor compression cycles can be spatially extensive, with tens of heat exchangers connected by hundreds of meters of pipe. As each of these heat exchangers is described by a large set of DAEs, the size of an overall system model can become very large, with potential for thousands of states and millions of equations.

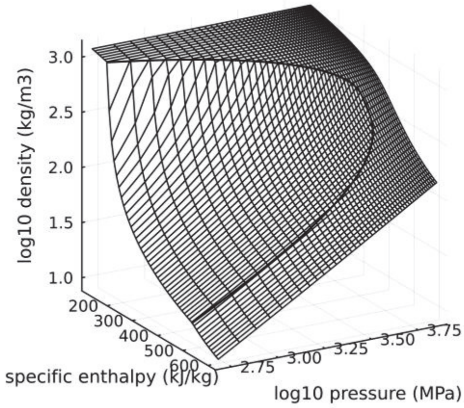

Figure 3: Density surface of refrigerant R32 as a function of pressure and specific enthalpy.

In addition to the large spatial scale needed to describe these systems, they are also governed by a wide range of temporal scales. The internal dynamics of the refrigerant flow typically evolve over times of milliseconds to seconds, while the hygrothermal behavior of a building often has time constants that range from days to weeks. This can be seen in Figure 1, which illustrates the eigenvalues for a 4-zone VRF system coupled to a building; the fastest time constants for this system are on the order of 10^{-5} seconds, while the slowest time constants are on the order of a day. As these time constants are spread over more than 10 orders of magnitude, these models are numerically stiff and require solvers that can manage these widely spaced time constants (Cellier, 2006). Couplings between variables at the system scale also often confound popular assumptions; for example, refrigerant pressures are often incorrectly assumed to exhibit only fast dynamics, but their coupling to air-side volumes results in fast and slow modes that are both apparent during system operation.

Fortunately, the structure of the overall system model is often quite sparse due to the local nature of the interactions between control volumes. This can be seen in Figure 2, which shows the sparsity structure of the incidence matrix for a system model that comprises a small vapor compression cycle and a building. The terms in the model corresponding to each subsystem are explicitly notated, as are the coupling terms relating the airflow and heat transfer between the two subsystems. The high level of sparsity for these systems (typically >98%) thus makes the use of sparse solvers advantageous for these applications.

2.4. Nonlinear characteristics

Since many of the main performance advantages of vapor compression cycles stem from the large variations in the heat transfer coefficients associated with evaporating and condensing fluids, it is essential that cycle models for digital twins capture the nonlinear behavior of these systems. Many physical processes of a cycle are also nonlinear, including the mass flow/pressure drop relations inside refrigerant tubes and expansion devices, the interactions between mass flow, pressure drop, and heat transfer during the compression process, and the process of water condensation or frost formation on the outside of refrigerant-to-air heat exchangers.

These systems also exhibit hybrid and discontinuous dynamics that affect their behavior. One example is illustrated in Figure 3, which shows the density of the popular refrigerant R32 as a function of pressure and specific enthalpy over a wide range of conditions. The density derivatives with respect to both of these variables are discontinuous at the liquid and vapor saturation lines, imposing important requirements for any numerical solver used for these models. Models that do not satisfy these constraints may produce results that have non-physical variations in the total refrigerant mass and exhibit attendant deviations in the overall cycle behavior (Laughman, 2017).

3. APPROACHES

Practical realizations of digital twins for vapor compression cycles must satisfy an array of requirements to ensure that the behavior of the virtual representation conforms to that of the physical system. In this section, we describe candidate technologies for this application by describing model representations that meet the performance requirements, parameter calibration techniques that are applicable to these large-scale and nonlinear systems, and state estimation methods that produce model output that is statistically consistent with the measured data to ensure that the virtual behavior tracks the physical behavior over time.

3.1. System representation

Many standard modeling tools and approaches are unable to correctly model vapor compression cycles because of their dynamic, large-scale, and nonlinear behavior. Of the many extant open-source and commercial modeling approaches for model creation, Wetter, et al (2016) describes how component-based, equation-oriented modeling languages such as Modelica (Modelica Association, 2021) play a prominent role for these applications by enabling users to manage model complexity via the systematic construction of large-scale physics-based models from component-based building blocks. Such equation-oriented modeling languages automatically compile model code into a software representation of a DAE, and use state-of-the-art DAE solvers to execute the model when running simulations. This compiled DAE code can also be interfaced to other computational environments for the use in calibration or state estimation via the Functional Mockup Interface (Junghanns et al., 2021), which is particularly valuable for digital twins.

A new tool named ModelingToolkit.jl (Ma et al., 2021) based on the Julia language (Bezanson et al., 2017). has recently emerged to further address the need for more advanced tools for equation-oriented acausal system modeling. ModelingToolkit uses a symbolic computational algebra framework that enables the construction of large acausal system models from smaller component models and generates imperative Julia code; unlike Modelica, whose compilers often target low-level languages such as C or Java, the Julia model code emitted by the ModelingToolkit compiler can be easily interfaced with many of Julia's numerical programming capabilities, such as new DAE solvers for numerically stiff systems, machine learning frameworks, and other state-of-the-art computational tools.

One example of ModelingToolkit's capabilities can be seen in its ability to interface with existing automatic differentiation (AD) tools (Baydin et al., 2018), which computationally generates derivatives of complex models. Jacobians and gradients of these models are valuable in digital twins they enable the sensitivity of parameters and/or states of the models to be calculated for the purposes of model corrections or updates. Conventional approaches for these derivative calculations use finite difference methods to calculate numerical Jacobians, but these methods can be inaccurate near the derivative discontinuities that are common in vapor compression cycles. In comparison, AD methods can symbolically calculate these derivatives for compiled model code at a given operating point, leading to derivatives of a much higher accuracy for similar or lower computational expense than numerical Jacobians.

3.2. Calibration

Models of vapor compression cycles typically have a wide range of geometric, performance (e.g., heat transfer coefficients), and equipment (e.g., compressor maps) parameters that need to be calibrated to ensure a good correspondence with physical systems. While some of these values may be obtained from manufacturing specifications, there may be significant uncertainty for many other parameters due to installation- and equipment-specific variation. System parameters may also vary slowly over the equipment life-cycle and affect the performance. Online methods for determining parameter values for the model thus have significant value to improve the performance of the digital twin over the operational life of the cycle.

These calibration processes typically proceed via sensitivity analyses, either by directly calculating the derivatives of specific outputs of the cycle model with respect to the parameters of interest, or by calculating the sensitivity of simulation outputs with respect to those same parameters. The computational difficulty of employing the former approach has motivated an emphasis on the latter, in which the map between parameter variations and the residual between model outputs and measured data is used to drive the calibration process. Because physics-based models are only valid for specific parameter ranges, calibration algorithms cannot apply arbitrary parameter variations to simulation models without occasionally causing the models to fail. The significant computational expense of solving systems of DAEs also tends to make simulations of cycle behavior computationally expensive, so that Monte Carlo-based calibration methods that rely on thousands of simulations are often computationally impractical.

Bayesian optimization-based methods have been demonstrated to be a promising technique for such calibration processes. These methods are based on the use of Gaussian processes to determine the mean and uncertainty of an acquisition function over a sample space of candidate parameters, after which an optimization algorithm identifies candidate regions of the parameter space to sample next to characterize the uncertainty and identify parameters which minimize a calibration-cost function. These methods have been demonstrated to be sample efficient, resulting in much faster convergence than Monte Carlo-based methods, and requiring fewer model simulations. Prior work by Chakrabarty et al. (2022) has demonstrated that these methods are effective at calibrating simulation models of vapor compression cycles that are only valid for limited parameter ranges. Furthermore, they can efficiently identify parameter values that result in a good fit between simulation output and measurement data for an application on a model that comprises a vapor compression cycle under closed-loop control in a building (Chakrabarty et al., 2021) in a purely data-driven manner. These methods are agnostic to the type of simulation model used and treat the model as a black-box, allowing the methods to generalize to various digital twin architectures and enable learning from multiple building data sources (Zhan et al, 2022).

3.3. State estimation

Despite the significant effort that is often invested in model development and calibration, no model of a cycle will perfectly reproduce the observed behavior of a real system due to simplifications, unmodeled physics, and other factors. Statistically consistent methods for correcting model state trajectories given a set of experimental observations, known as state estimators, ensure that digital twin predictions optimally trade-off between the information available from measurements and the information encoded in the model structure.

Kalman-based state estimators represent a popular and powerful approach to solving these problems, and function by first propagating the model forward between measurement times, and then calculating the state corrections to the model at measurement times on these basis of the mean and covariance of the corrections. These methods can either be implemented either online as a filter, or in an acausal manner as a smoother. A key advantage of these methods is that state estimates constructed from available measurements can be used to analyze other system variables that are otherwise impractical to measure. These methods can also be extended to nonlinear systems via linearization and are known as extended Kalman filters (EKF) or smoothers (EKS) (Simon, 2006). Such approaches are required for vapor compression cycles due to the nonlinear behavior of the models.

While these methods have been successfully used for vapor compression cycles (Bortoff et al., 2019; Cheng et al., 2005), these prior applications have generally been focused on low-order component and system models. Large-scale models for digital twins pose particular challenges for extended Kalman methods because of memory constraints on size of the covariance matrix, which increases with the square of the number of state variables. In addition, the finite difference methods used to calculate the Jacobian for the extended Kalman estimators can be inaccurate due to discontinuities in the underlying nonlinear model. Finally, the calculated state corrections for these estimators are not guaranteed to satisfy physics-based constraints that are assumed by models. Such constraint violations can result in erroneous model behavior, ranging from non-physical output values to computational model failure.

Ensemble Kalman filters (EnKF) and smoothers (EnKS) (Evensen, 2009) represent alternative state estimation methods that address the limitations of the EKF and EKS. Instead of integrating the covariance matrices forward over each time interval between measurements, ensemble Kalman estimators use a sequential

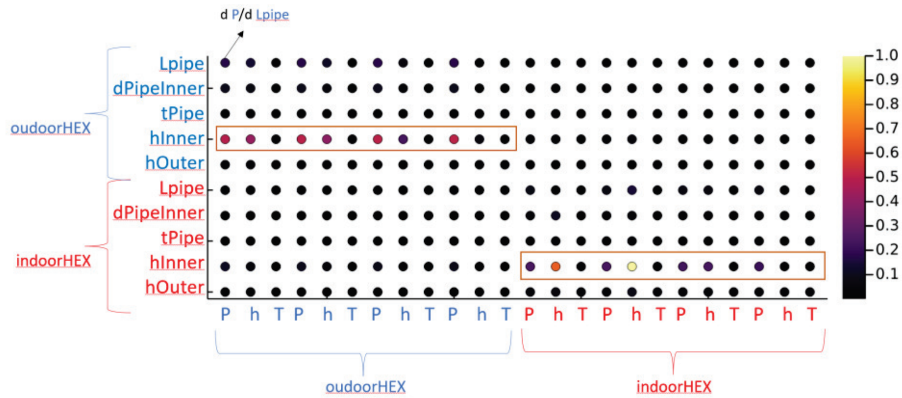

Figure 4: Local sensitivities of cycle states with respect to model parameters. Sensitivities are normalized and scaled between 0 and 1.

Monte Carlo approach to directly estimate the covariance matrix under an assumption of normally distributed state variables from an ensemble of sampled state vectors, or particles, that are propagated through the system dynamics. This dramatically reduces the memory storage and computational methods in comparison to the EKF and EKS, and even a small number of particles can result in state estimates with accuracies comparable to those of the extended Kalman methods. These ensemble-based methods can also incorporate physics-based constraints (Deshpande et al, 2022) by recognizing that the classical state correction computation can be formulated as an optimization problem, and thus can be reformulated as a constrained optimization problem to manage such constraints.

Beyond ensemble Kalman methods, alternative Monte Carlo-based estimation methods such as particle filters and smoothers that are not formulated under assumptions of normally-distributed variables may better represent the effect of model nonlinearities. Such methods require large number of particles for good estimation accuracy, which incurs significant computational cost when full-order physics-based models are used. Reduced-order or surrogate models which can efficiently simulate large sets of particles thus have potential to further improve the efficiency of particle-based estimation methods.

4. USE CASES

While digital twins in vapor compression cycles have a wide range of potential uses that employ these modeling, calibration, and state estimation methods, we focus here on two brief case studies that build on the cited literature to demonstrate the range of possibilities for this technology. We first identify a set of digital twin model parameters to calibrate against measured data by performing a sensitivity analysis on a cycle model. This uses the symbolic AD-compatible properties of the ModelingToolkit cycle models to efficiently calculate parametric derivatives of the cycle model, rather than using time-consuming and inaccurate numerical derivatives. In the second case study, we estimate the refrigerant mass and pipe length for the same cycle for performance monitoring or diagnostics by using an EKF with temperature measurements located at the middle and outlet of each HEX as well as pressure measurements at both compressor ports. We thus combine the information encoded in the structure of the physics-based models with data from system measurements to study the behavior of variables that are valuable but impractical to observe from direct experimental measurements.

We developed these case studies on a full-scale dynamic ModelingToolkit model of a air-source vapor compression cycle with an evaporating and a condensing heat exchanger (HEX), as well as a variable speed compressor, variable speed fans, and a variable position expansion device (Deshpande et al., 2022). Since the HEX dynamics dominate the overall behavior of the system, dynamic models for these components were constructed using finite volume discretizations with four volumes for each HEX that describe the one-dimensional refrigerant flow, thermal behavior of the tube wall, and the airflow across the HEX. Algebraic

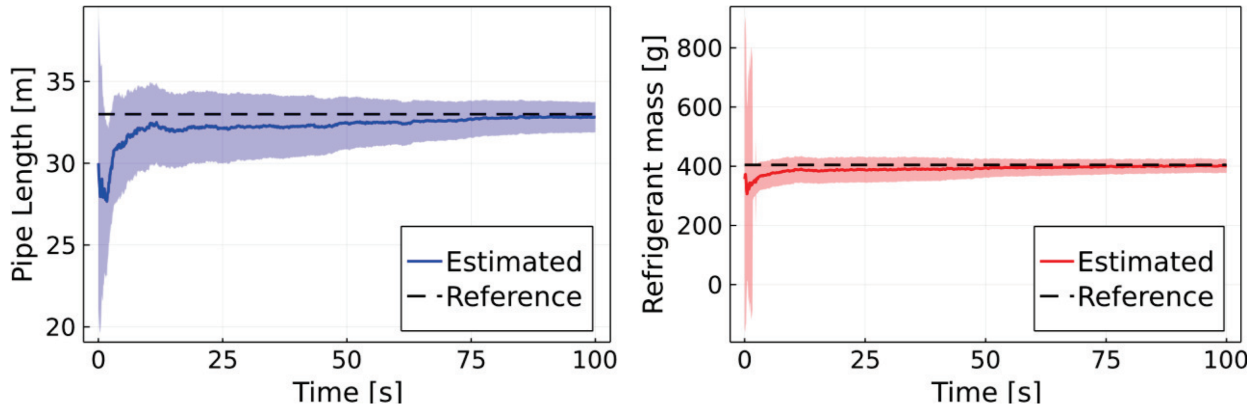

Figure 5: Solid lines show estimated pipe length (left) and refrigerant mass (right) while shaded areas show 3σ bounds. Dashed lines show reference parameter values.

models were used for the compressor and expansion valve. The resulting cycle model consists of an index-1 set of 278 DAEs, which was reduced via index reduction to a 24-dimensional set of nonlinear ODEs. The 24-dimensional state vector consists of pressures (P), specific enthalpies (h), and wall temperatures (T) for each control volume in both HEXs.

Figure 4 illustrates a heatmap of local sensitivities of the cycle states (on the horizontal axis) to a set of model parameters (on the vertical axis) that were calculated directly from the cycle model via AD at a given operating point. Since values corresponding to different state variables span several orders of magnitude, these sensitivities are normalized and scaled between 0 and 1 for the purposes of comparison. As is evident from the figure, the dynamics of both HEXs are most sensitive to the refrigerant-side heat transfer coefficient `hInner` at the inner wall of the HEX pipes. This type of analysis can automatically improve the speed and computational efficiency of the model calibration process, as it indicates that this effort should be focused on only two of the ten parameters in this set.

The AD compatibility of this model was also used for joint state and parameter estimation using the extended Kalman framework, which uses model Jacobians and local sensitivities for model linearizations. This joint estimation is accomplished by first augmenting the model's state vector with unknown or uncertain model parameters of the model, and then estimating the augmented state vector using the EKS. The accuracy of this approach is evident in Figure 5, where the left plot demonstrates the efficacy of this method by accurately estimating the pipe length of a cycle model from a limited set of temperature and pressure measurements. Similarly, the right plot of this same figure demonstrates the ability of this method to simultaneously estimate the refrigerant mass as calculated from the calibrated pipe length and the estimated state variables. For this simulated case, the total refrigerant mass of this system was estimated with less than 1% error. As changes in the refrigerant mass can have a significant effect on the energy performance and direct climate impact of vapor compression cycles, the ability of this technology to effectively provide “virtual sensors” for otherwise unobservable phenomena has high potential value.

These demonstrations suggest that digital twins can add to the advantages of digital models by assimilating updated data that reflect important changes in the vapor compression cycle behavior over its lifecycle. Whereas digital models are generally only designed to represent the system during the design stage, the additional structure of digital twins of vapor compression cycles that satisfy the stated requirements can create new system-level capabilities that provide actionable performance information. This information can be used in next-generation integrated building control and planning algorithms, as well as enable the development of advanced equipment control and optimization methods to enhance both system-level and fleet-level energy and climate-related performance.

5. CONCLUSIONS

Digital twins have an opportunity to play a key role as we strive to convert the flood of data from low-cost sensors into actionable information. While vapor compression cycles have characteristics that impose certain implementation challenges for digital twins, technology and methods are readily available to further develop these tools for performance monitoring, control, and diagnostic applications. Future efforts to refine these methods and test them on experimental data promise to provide valuable insights and new avenues of investigation, as well as enable new applications that we hope will play some small part in achieving our larger climate and energy objectives.

REFERENCES

AIAA Digital Engineering Integration Committee. Digital twin: Defnition & value. An AIAA and AIA position paper, American Institute of Aeronautics and Astronautics (AIAA) and Aerospace Industries Association (AIA), 2020.

Baydin, A., Pearlmutter, B., Radul, A., & Siskind, J. (2018). Automatic differentiation in machine learning: A survey. Journal of Machine Learning Research, 18, 1-43.

Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A fresh approach to numerical computing. SIAM review, 59(1), 65–98. doi: 10.1137/141000671

Bhattacharya, C., Chakrabarty, A., Laughman, C., & Qiao, H. (2022). Modeling nonlinear heat exchanger dynamics with convolutional recurrent networks. IFAC-PapersOnLine, 55(37), 99-106. (2nd Modeling, Estimation and Control Conference MECC 2022) doi: 10.1016/j.ifacol.2022.11.168

Bortoff, S., Schwerdtner, P., Danielson, C., Di Cairano, S., & Burns, D. (2022). H-infinity loop-shaped model predictive control with HVAC application. IEEE Transactions on Control Systems Technology, 30(5), 2188–2203

Cellier, F.E., & Kofman, E. (2006). Continuous System Simulation. Springer.

Chakrabarty, A., Bortoff, S. A., & Laughman, C. R. (2022). Simulation failure-robust Bayesian optimization for data-driven parameter estimation. IEEE Transactions on Systems, Man, and Cybernetics: Systems, 1-12. doi: 10.1109/TSMC.2022.3216790

Chakrabarty, A., Maddalena, E., Qiao, H., & Laughman, C. (2021). Scalable Bayesian optimization for model calibration: Case study on coupled building and HVAC dynamics. Energy and Buildings, 253, 111460. doi: 10.1016/j.enbuild.2021.111460

Cheng, T., He, X.-D., & Asada, H. (2005). Nonlinear observer design and experimental verification for heat exchangers during the start-up process. In 2005 American Control Conference. doi: 10.1109/ACC.2005 .1470757

Das, H. P., Lin, Y.-W., Agwan, U., Spangher, L., Devonport, A., Yang, Y., ... Spanos, C. J. (2022). Machine learning for smart and energy-efficient buildings. arXiv preprint 2211.14889.

Deshpande, V. M., Laughman, C. R., Ma, Y., & Rackauckas, C. (2022). Constrained smoothers for state estimation of vapor compression cycles. In 2022 American Control Conference (ACC) (p. 2333-2340). doi: 10.23919/ACC53348.2022.9867269

Evensen, G. (2009). The ensemble Kalman filter for combined state and parameter estimation. IEEE Control Systems Magazine, 29(3), 83–104.

Junghanns, A., Gomes, C., Schulze, C., Schuch, K., Pierre, R., Blaesken, M., Sommer, T. (2021). The Functional Mock-up Interface 3.0 - New features enabling new applications. In Modelica Conferences (pp. 17–26).

Laughman, C.R., & Qiao, H. (2017). On the influence of state selection on mass conservation in dynamic vapor compression cycle models. Mathematical and Computer Modelling of Dynamical Systems, 23 (3), 262-283.

Li, P., Qiao, H., Li, Y., Seem, J.E., Winkler, J., & Li, X. (2014a). Recent advances in dynamic modeling of HVAC equipment. Part 1: Equipment modelling. HVAC&R Research, 20 (4), 136-149.

Li, P., Li, Y., Seem, J.E., Qiao, H., Li, X., & Winkler, J. (2014b). Recent advances in dynamic modeling of HVAC equipment. Part 2: Modelica-based modelling. HVAC&R Research, 20 (4), 150-161.

Ma, Y., Gowda, S., Anantharaman, R., Laughman, C., Shah, V., & Rackauckas, C. (2021). ModelingToolkit: A composable graph transformation system for equation-based modeling. arXiv preprint 2103.05244.

Modelica Association. (2021). Modelica specification, Version 3.5. Retrieved from www.modelica.org on 2023/02/09.

Simon, D. (2006). Optimal state estimation: Kalman, H-infinity, and nonlinear approaches. John Wiley & Sons.

Thelen, A., Zhang, X., Fink, O., Lu, Y., Ghosh, S., Youn, B. D., Hu, Z. (2022). A comprehensive review of digital twin - Part 1: Modeling and twinning enabling technologies. Structural and Multidisciplinary Optimization, 65(12), 354. doi: 10.1007/s00158-022-03425-4

Vering, C., Mehrfeld, P., Nurenberg, M., Coakley, D., Lauster, M., & Müller, D.¨ (2019). Unlocking potentials of building energy systems' operational efficiency: Application of digital twin design for HVAC systems. 16th International Building Performance Simulation Association (IBPSA), 1304–1310.

Wetter, M., Bonvini, M., and Nouidui, T. (2016). Equation-based languages – A new paradigm for building energy modelling, simulation, and optimization. Energy and Buildings, 117, 290-300. doi: 10.1016/j.enbuild.2015.10.017

Xie, X., Merino, J., Moretti, N., Pauwels, P., Chang, J. Y., & Parlikad, A. (2023). Digital twin enabled fault detection and diagnosis process for building HVAC systems. Automation in Construction, 146, 104695.

Zhan, S., Wichern, G., Laughman, C., Chong, A., & Chakrabarty, A. (2022). Calibrating building simulation models using multi-source datasets and meta-learned Bayesian optimization. Energy and Buildings, 270, 112278.