Rapid Energy Optimization of Vapor Compression Systems Using Probabilistic Machine Learning and Extremum Seeking Control

Chakrabarty, Ankush; Burns, Daniel J.; Guay, Martin; Laughman, Christopher R.

TR2022-172 May 24, 2023

Abstract

Extremum seeking control (ESC) is a popular datadriven approach for optimizing the energy consumption of vapor compression systems (VCS). Tuning ESC control parameters can present a challenge to implementation, especially in advanced variants of ESC, because time-consuming and problemspecific manual tuning is often required to eliminate numerical and dynamical instabilities. In this paper, we propose an automatic ESC tuning mechanism based on a Bayesian optimization framework that systematically leverages closedloop ESC experiments to compute highperforming ESC parameters. We validate the proposed Bayesian-optimized ESC on a physicsbased Modelica model of a VCS. This new approach is six times faster and yields a 9% higher coefficient of performance than a stateoftheart timevarying ESC method under identical experimental conditions.

International Refrigeration and Air Conditioning Conference at Purdue 2022

Rapid Energy Optimization Of Vapor Compression Systems Using Probabilistic Machine Learning And Extremum Seeking Control

Ankush CHAKRABARTY[1*], Daniel J. BURNS[2], Martin GUAY[3], Christopher R. LAUGHMAN[1]

[1] Mitsubishi Electric Research Laboratories
Cambridge, MA, USA
Email: {chakrabarty, laughman}@merl.com

[2] Email: danburns@ieee.org

[3] Department of Chemical Engineering
Queens University, ON, Canada
Email: guaym@queensu.ca

* Corresponding Author

ABSTRACT

Extremum seeking control (ESC) is a popular data-driven approach for optimizing the energy consumption of vapor compression systems (VCS). Tuning ESC control parameters can present a challenge to implementation, especially in advanced variants of ESC, because time-consuming and problem-specific manual tuning is often required to eliminate numerical and dynamical instabilities. In this paper, we propose an automatic ESC tuning mechanism based on a Bayesian optimization framework that systematically leverages closed-loop ESC experiments to compute high-performing ESC parameters. We validate the proposed Bayesian-optimized ESC on a physics-based Modelica model of a VCS. This new approach is six times faster and yields a 9% higher coefficient of performance than a state-of-the-art time-varying ESC method under identical experimental conditions.

1. INTRODUCTION

Vapor compression systems (VCSs) provide essential functionality in many energy transfer systems due to their cost-effectiveness and reliability (She et al., 2018). As the design of VCSs has continued to evolve to meet increasingly stringent specifications, one key trend that characterizes this development is the use of variable actuators to enable equipment to adapt to variations in the operating conditions, disturbances, and plant uncertainties (Chua, Chou, & Yang, 2010). While the inclusion of these actuators can result in lower energy consumption or higher thermal comfort over a wide range of conditions, these variable actuators require the use of feedback controllers to regulate the system's dynamic behavior to prescribed actuator setpoints.

The design, tuning, and validation of feedback controls in VCSs can be an expensive and time-consuming process for equipment manufacturers, which must ensure proper closed-loop performance over a wide range of environmental and installation conditions. Because the system energy efficiency depends on the selection of setpoints provided to these feedback controllers, the careful assignment of setpoints in an outer-loop enables the energy optimization of performance without reconfiguring system architectures or re-tuning inner-loop controller parameters (e.g., PID gains) (Jäschke & Skogestad, 2011). Unfortunately, the task of assigning a set of optimal setpoints as a function of the regulated inputs and the driving conditions is often challenging due to the nonlinear, multivariable, and frequently unmodeled dynamics of the closed-loop system.

Extremum-seeking control (ESC) design techniques have proven to be effective in computing energy-efficient setpoints of complex VCSs (Li, Li, & Seem, 2010; Wang & Li, 2019; Burns, Laughman, & Guay, 2020; Chakrabarty, Danielson, Bortoff, & Laughman, 2021). ESC algorithms are model-free techniques that perform a gradient descent on an unknown convex map representing the steady-state relationship between manipulated inputs and a performance output. Since ESC works without an *a priori* characterization of this map, the approach is inherently robust to disturbances and the wide variation of environments in which vapor compression systems are deployed. However, convergence to the optimizer using the common perturbation-based extremum seeking control occurs at a rate two timescales slower than the dominant plant dynamics, which is a severe constraint for thermal systems with large time constants. To address this limitation, a proportional-integral extremum seeking control (PI-ESC) algorithm was developed that converges to the optimizer at the same time-scale as the dominant plant dynamics (Guay & Burns, 2017).

While PI-ESC can offer significant convergence speedup compared with other ESC algorithms (Burns, Laughman, & Guay, 2016), its behavior is governed by multiple parameters that interact in non-intuitive ways; in some instances, certain combinations of these parameters can even render the closed-loop system unstable. This characteristic makes the process of tuning PI-ESC controllers rather challenging. In this paper, we address this limitation by developing a Bayesian optimization (BO) method that uses VCS data directly to discover combinations of PI-ESC gains that achieve rapid convergence while ensuring the closed-loop VCS is numerically and dynamically stable. Since BO is a global, derivative-free optimization methodology that is designed to obtain solutions without a large number of function evaluations/iterations, it has been reported to perform well on controller tuning problems for many industrial applications (Neumann-Brosig, Marco, Schwarzmann, & Trimpe, 2019). BO's data-driven nature implies that it is agnostic to the control architecture under consideration and insensitive to unmodeled dynamics of the closed-loop system (Khosravi et al., 2021; Lu, Kumar, & Zavala, 2020).

Contributions: (i) we tune PI-ESC parameters in the outer-loop and do not modify the inner-loop PID controllers, as the inner-loop controller gains are designed to ensure that the dynamic behavior of the equipment satisfied design requirements and operational constraints; and, (ii) prior work such as (Khosravi, Eichler, Schmid, Smith, & Heer, 2019) has not explicitly considered potential instabilities that arise during the PI-ESC tuning process. In comparison, our recently proposed failure-robust BO (FRBO) algorithm (Chakrabarty, Bortoff, & Laughman, 2021) is designed specifically to learn unstable regions from unstable combinations of PI-ESC parameters, and avoid those regions in future BO iterations. Since the region in the PI-ESC parameter space which yields 'good' parameters is problem-specific, we use problem-specific data to provide a suitable option to attack this problem via a novel modification of Bayesian optimization, which is data-driven, to automatically identify regions of safe combinations of PI-ESC parameters and efficiently search within such safe regions when tuning PI-ESC controller parameters. The success of this failure-robust BO (FRBO) framework is demonstrated by outperforming a PI-ESC controller with hand-tuned parameters on a VCS application.

Organization: In Section 2, we provide an overview of how PI-ESC is used for energy optimization in vapor compression systems. Section 3 delineates our novel FRBO framework for tuning PI-ESC parameters despite failures. Section 4 describes our experimental setup, results, and discussion. We present our conclusions in Section 5.

2. ENERGY OPTIMIZATION USING PI-EXTREMUM SEEKING CONTROL

The problem of energy minimization of the VCS can be abstracted as follows. The underlying VCS dynamical system can be modeled by the closed-loop system

$$x_{t+1} = f(x_t, r_t) \tag{1a}$$

$$y_t = h(x_t, r_t), \tag{1b}$$

where t is the time index, $x_t \in \mathbb{R}^{n_x}$ is the vector of state variables at time t, r_t is the setpoint variable at time t taking values in $\mathcal{R} \subset \mathbb{R}^{n_r}$ and $y_t \in \mathbb{R}_{\geq 0}$ is a scalar power output of the VCS at time t; our objective is to minimize y. We assume that the dynamics f and the setpoint-to-energy function h are both unmodeled (therefore, unknown at design time). Empirically, we have observed that the energy function h exhibits sufficient smoothness to warrant the use of data-driven gradient estimates for energy optimization (Burns et al., 2020).

The principle of ESC is based on obtaining a sequence of setpoints r_t for $t \geq 0$ such that each r_t moves along a direction of negative gradient of the function h. A first-order ESC control law consequently has the form $r_{t+1} = -k_g g_t + d_t$, where g_t is an estimate of the gradient of h w.r.t. r, k_g is a control gain or step-size, and d_t is a persistently exciting dither signal (Chakrabarty, Danielson, et al., 2021). By incorporating integral action to the aforementioned control law, its convergence rate has been significantly accelerated, as demonstrated in (Guay & Burns, 2017). The proportional-integral ESC (PI-ESC) law is given, in velocity form, by

$$r_{t+1} = r_t - k_g(g_{t+1} - g_t) - \frac{1}{\tau_I} g_t + d_t, \tag{2}$$

where τ_I is a time constant of the integral term. The gradient estimate g_t and its update g_{t+1} are not known because f and h are unmodeled.

To this end, we exploit the smoothness of h to write a Taylor series approximation of the difference in cost values, ignoring second- and higher-order terms, as

$$\Delta y_t := y_{t+1} - y_t = \underbrace{\frac{\partial h}{\partial x}}_{\triangleq g_t^x} \Delta x_t + \underbrace{\frac{\partial h}{\partial r}}_{\triangleq g_t^r} \Delta r_t. \tag{3}$$

Since f and h represent unmodeled functions, we cannot obtain a numerical or analytical derivative of $\partial h/\partial x$ or $\partial h/\partial r$. Instead, we estimate the unknown quantities in (3) directly from data $\{(r_t, y_t)\}$ obtained during experiments with the closed-loop system.

One can then use the data to formulate a linear regression problem that involves solving the matrix equation

$$\begin{bmatrix} \Delta y_{t-N_\ell+1} \\ \vdots \\ \Delta y_t \end{bmatrix} = \begin{bmatrix} 1 & \Delta r_{t-N_\ell+1}^\top \\ \vdots & \vdots \\ 1 & \Delta r_t^\top \end{bmatrix} \underbrace{\begin{bmatrix} g_t^x \\ g_t^r \end{bmatrix}}_{g_t}.$$

from which estimates of the gradient g_t can be obtained efficiently.

Denoting $\hat{g}_t^r$ and $\hat{g}_t^x$ as estimates of g_t^x and g_t^r, respectively, and $\hat{g}_t$ as an estimate of g_t, we use the following recursive least-squares (RLS) estimator at time $t-1$ to estimate the gradient at time t:

$$w_t = w_{t-1} - F w_{t-1} + \varphi_{t-1}, \qquad P_t = \alpha P_{t-1} + w_{t-1} w_{t-1}^\top + \varepsilon I, \tag{4a}$$

$$K_t = \frac{P_t^{-1} w_{t-1}}{\alpha + w_{t-1}^\top P_t^{-1} w_{t-1}}, \qquad \hat{g}_t = \hat{g}_{t-1} + K_t e_t, \tag{4b}$$

where $e_t = y_t - \hat{y}_t$, $\varphi_{t-1} \triangleq \begin{bmatrix} 1 & r_{t-1} \end{bmatrix}$, ε is a small scalar-valued term that seeks to ensure good numerical conditioning of P_t, and F is a scalar filter gain coefficient. Subsequently, we use

$$\hat{y}_{t+1} = \hat{y}_t + F e_t + \varphi_t^\top \hat{g}_{t-1} + w_t^\top (\hat{g}_t - \hat{g}_{t-1}).$$

to predict the cost $\hat{y}_t$ at time t based on the estimated gradient (4b). A dither signal is also required to ensure that the system is persistently excited. We employ a dither signal of the form $d_t = D \sin(\omega t + \varphi_0)$, where ω is a vector of unique frequencies of the sinusoidal dither, φ_0 is a vector of unique initial phases of the sinusoidal dither, and D is a user-defined small amplitude.

In summary, the following parameters must be selected to design a PI-ESC controller: (i) the integral time constant τ_I; (ii) the control gain k_g; (iii) the forgetting factor α; (iv) the filter coefficient F; (v) the dither magnitude D; (vi) the dither frequency ω; and, (vii) the dither phase φ. We have empirically observed that the PI-ESC performance is strongly correlated with the selection of each of these variables. In fact, if these variables are not carefully selected, the PI-ESC can exhibit numerical instability which results in unstable dynamics for the closed-loop system.

3. FAILURE-ROBUST BAYESIAN OPTIMIZATION FOR CONTROLLER TUNING

For the ensuing discussion, we will search for the following PI-ESC parameters:

$$\theta \triangleq \begin{bmatrix} \tau_I & k_g & \alpha & F & D^\top & \omega^\top & \varphi^\top \end{bmatrix}^\top,$$

where the dimension of $\theta = 1+n_r+1+1+n_r+n_r+n_r = 3+4n_r$. We assume that the set of admissible parameters Θ, which forms the search space for BO, is known to the designer. Clearly, even for small n_r, the search space is not sufficiently low-dimensional to justify trying to tune the PI-ESC parameters manually. We have observed that attempts to hand tune these parameters for VCS applications often results in unstable closed-loop dynamics or frequent triggering of failsafe mechanisms due to impractically large PI-ESC control actions.

We begin by describing classical Bayesian optimization methods, and then describe details of the failure-robust BO approach. For the j-th BO iteration, let θ_j denote the candidate set of PI-ESC parameters. With these parameters, an experiment is performed in which the PI-ESC closed-loop system parameterized by θ_j is observed on a time interval of interest $[\mathcal{T}_0, \mathcal{T}_f]$, and measurements $y_{\mathcal{T}_0:\mathcal{T}_f} := \{y_{\mathcal{T}_0}, y_{\mathcal{T}_0+1}, \cdots, y_{\mathcal{T}_f}\}$ obtained from the system over this time interval are used to compute a performance value J_j. By learning a probabilistic surrogate model from θ to J and exploiting the statistics of the learned surrogate model, BO generates a sequence of θ candidates that converge to the optimal (in the sense of the performance metric chosen) PI-ESC parameters $\theta^\star$.

For performance-driven BO, an objective function (to be minimized) we have found to be useful has the form

$$J(x) = \eta_1 J_{\text{avg}}(x) + \eta_2 J_{\text{osc}}(x),$$

where x is an initial condition of the system and η_1 and η_2 are positive weights on each component. Here, J_{avg} is designed to filter out the dithering effect in steady state and promote lower steady state J values. This component of the cost is obtained by computing the mean of the final $\mathcal{T}'$ cost values, that is $J_{\text{avg}} \triangleq \mathbb{E}[J_{\mathcal{T}_f-\mathcal{T}':\mathcal{T}_f}]$, where $\mathbb{E}$ is the expectation operator. The second component of the cost $J_{\text{osc}} \triangleq \frac{1}{(\mathcal{T}_f - \mathcal{T}_0)} \sum_{t=\mathcal{T}_0+1}^{\mathcal{T}_f} |J_t|$ is designed to penalize oscillations in the closed-loop PI-ESC response. Since we aim to evaluate performance over multiple initial conditions within an admissible set of initial conditions, $\mathbb{X}$, we repeat these simulations for a set of $N_x \in \mathbb{N}$ initial conditions $\{x_i\}_{i=1}^{N_x}$ and compute the total cost over this set. That is, the cost assigned to the parameter θ_j is given by

$$J_j = \sum_{i=1}^{N_x} J_{\text{avg}}(x_i) + J_{\text{osc}}(x_i), \tag{5}$$

where the two components are defined above.

Classical BO methods assume the presence of a single global optimum, and smoothness of the θ to J map. Since J is typically assumed to be continuous, one can leverage the data at the j-th iteration to construct a surrogate Gaussian Process (GP) model of the cost, given by

$$\hat{J}_j := \mathsf{GP}\left(\mu(\theta; \mathcal{D}_j), \sigma(\theta, \theta'; \mathcal{D}_j)\right), \tag{6}$$

where $\mu(\cdot)$ is the predictive mean function, and $\sigma(\cdot,\cdot)$ is the predictive variance function, and $\mathcal{D}_j$ containing $\{\theta_{[0:j]}, J_{[0:j]}\}$ is the dataset collected thus far. Typically, the variance is expressed through the use of kernel functions (Snoek, Larochelle, & Adams, 2012).

At the j-th learning iteration, for a new query sample $\theta \in \Theta$, the GP model predicts the mean and variance of the reward to be

$$\mu(\theta) = k_j(\theta)^\top \mathfrak{K}_j^{-1} J_{0:j} \quad \text{and} \quad \sigma(\theta) = \mathcal{K}(\theta,\theta) - k_j(\theta)\mathfrak{K}_j^{-1} k_j(\theta)^\top,$$

where

$$k_j(\theta) = \begin{bmatrix} \mathcal{K}(\theta_0,\theta) & \mathcal{K}(\theta_1,\theta) & \cdots & \mathcal{K}(\theta_j,\theta) \end{bmatrix}, \quad \text{and} \quad \mathfrak{K}_j = \begin{bmatrix} \mathcal{K}(\theta_0,\theta_0) & \cdots & \mathcal{K}(\theta_0,\theta_j) \\ \vdots & \ddots & \vdots \\ \mathcal{K}(\theta_j,\theta_0) & \cdots & \mathcal{K}(\theta_j,\theta_j) \end{bmatrix},$$

The accuracy of predicted mean and variance is strongly linked to the selection of the kernel and the best (in some sense) set of hyperparameters such as length scales and variance parameters of the kernels and estimated noise. We obtain these hyperparameters by maximizing the log marginal likelihood function.

In Bayesian optimization, we use the mean and variance of the surrogate model $\hat{J}_j$ in (6) to construct an acquisition function to inform the selection of a θ_j that increases the likelihood of minimizing the current best cost. To this end, we compute the incumbent $\hat{J}_j^\star := \min_{\theta\in\Theta} \mu(\theta; \mathcal{D}_j)$ and use it to define an expected improvement (EI) acquisition function that has the form

$$\mathsf{EI}(\theta, j) = \begin{cases} \sigma(\theta)\gamma(z) + (\hat{J}_j^\star - \mu(\theta))\Gamma(z), & \text{if } \sigma(\theta) > 0, \\ 0 & \text{if } \sigma(\theta) = 0. \end{cases}$$

where $z = \frac{\hat{J}_j^\star - \mu(\theta)}{\sigma(\theta)}$, and $\gamma(\cdot)$, $\Gamma(\cdot)$ are the PDF and the CDF of the zero-mean unit-variance normal distribution, respectively. In the j-th iteration of learning, we use the data $\mathcal{D}_j$ to construct the EI acquisition function using the surrogate $\hat{J}_j$. Subsequently, we compute the optimizer candidate

$$\theta_{j+1} = \arg\max_{\theta\in\Theta} \mathsf{EI}(\theta, j), \tag{7}$$

which serves as the parameter estimate θ in the next BO iteration.

3.1 Failure Robust Bayesian Optimization

While one could, in theory, apply classical Bayesian optimization directly to find an optimal set of PI-ESC parameters $\theta^\star$, we have empirically found that this approach has severe limitations. The classical BO procedure tends to select θ candidates frequently that make the PI-ESC closed-loop system unstable, causing a loss of cooling capacity. In hardware, this event triggers failsafe operation; in simulated systems, this event results in non-physical states that violate model assumptions. These failed simulations also often result in arbitrarily large and nonsensical cost values that render the (θ, J) datapoint useless for consequent BO iterations. Since seemingly arbitrary combinations of the components of θ result in these instabilities, classical BO wastes many iterations evaluating candidates that cause instability, resulting in very slow convergence rates. Rather than employing heuristics to avoid the regions in parameter space that are likely to result in simulation failures, we adopt a data-driven approach to estimate these regions and avoid them by means of a modified acquisition function tailored to promote 'failure robustness'. We refer to this algorithm as failure-robust BO, or FRBO, (Chakrabarty, Bortoff, & Laughman, 2021). A few exemplar iterations of FRBO are shown in Fig. 1 to accompany this mathematical description.

The first step of the FRBO algorithm involves constructing a dataset for failure region estimation, which we will pose as a supervised learning problem. We begin by randomly sampling a set of initial conditions $\mathbb{X}_{N_x} := \{x_i\}_{i=1}^{N_x}$ within the space of admissible initial conditions $\mathbb{X}$; this set is then kept fixed. At the j-th FRBO iteration, we simulate the PI-ESC closed-loop system on $[\mathcal{T}_0, \mathcal{T}_f]$ where the PI-ESC is parameterized by the candidate parameter θ_j. If the simulation fails, we assign a label $\ell_j = 1$ and store a nonsense value in the cost $J_j = \mathtt{NaN}$. If the simulation is successful, we assign a label $\ell_j = -1$ and store the real-valued cost J_j described in (5). Thus, at the end of the j-th FRBO iteration, we have a dataset $\{(\theta_k, \ell_k, J_k)\}_{k=0}^{j}$.

We can use the (θ, ℓ) components of this dataset to estimate the failure region boundary by casting the estimation problem as a supervised learning problem. In particular, we need to construct a probabilistic learning machine $\mathcal{F} : \Theta \to [0, 1]$, where the output of the learner is the probability that $\theta \in \Theta$ is inside the failure region. That is, $\mathcal{F}(\theta)$ is the learned probability that the closed-loop PI-ESC parameterized by θ will be unstable (and therefore, fail). A few considerations go into the selection of a learning algorithm suited for the task of failure region estimation. First, $\mathcal{F}$ should be able to generate meaningful estimates of the failure region despite being trained on limited (θ, ℓ) data, since the FRBO algorithm is designed to converge without requiring a large number of iteration. Second, $\mathcal{F}$ needs to be retrained often, so a learner that requires a large number of training iterations before yielding accurate predictions is not ideal for FRBO. Third, $\mathcal{F}$ needs to be able to generate decision boundaries exhibiting complex geometries, since the failure region may have irregular contours. One or more of these pre-requisites restrict the utility of deep neural networks (needs large datasets) or linear classifiers (cannot generate nonlinear decision boundaries). Instead, we have found that a nonparametric kernel-based probabilistic classifier such as a variational Gaussian process classifier (VGPCs) (Chakrabarty, Bortoff, & Laughman, 2021; Hensman, Matthews, & Ghahramani, 2015) works well in practice for failure region estimation in the context of FRBO. Over comparable classifiers like a probabilistic support vector machine, the VGPC has the advantage that its outputs—a mean and variance—is easily interpretable since the variational proxy distribution is taken to be Gaussian.

Once the failure region estimator is trained, the FRBO algorithm utilizes a failure-robust expected-improvement (FREI) acquisition function of the form

$$\mathsf{FREI}(\theta,j) = \mathsf{EI}(\theta,j) \cdot (1 - \mathsf{PF}(\theta,j)), \tag{8}$$

where $\mathsf{PF}(\theta,j) = \mathcal{F}(\theta_j)$ is the probability of failure calculated by training $\mathcal{F}$ using (θ, ℓ) data up to the j-th iteration, and EI is described in (7). Note that the GP surrogate of the cost (6) is trained on (θ, J) data which resulted in stable closed-loop trajectories for which $\ell_j = +1$. Upon maximizing this acquisition function, FR-BO selects the next optimizer candidate

$$\theta_{j+1}^{*} = \arg\max_{\theta\in\Theta} \ \mathsf{FREI}(\theta,j). \tag{9}$$

This maximization problem is often solved in low dimensions by sampling on Θ, evaluating the acquisition function on those samples, and choosing the maximizer on that finite set of samples. In such an approach, the samples on Θ are selected randomly at each FRBO iteration to encourage exploration.

We note that maximizing FREI indicates that the component EI should be large, and that the component PF should be small, i.e., near zero. The former indicates that the considered θ is likely to minimize the cost function (5), and the latter implies that such a θ is likely to result in a successful simulation. By combining both these beneficial qualities

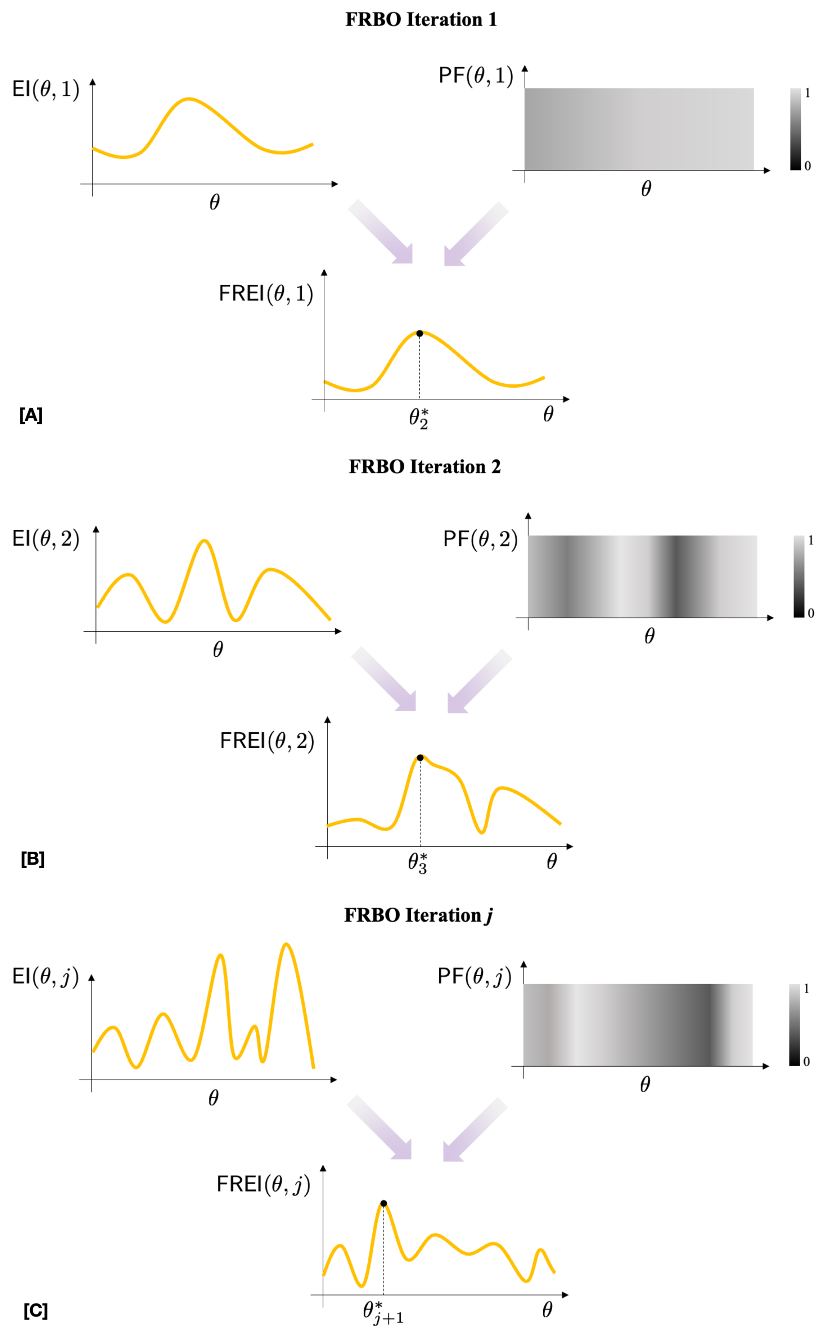

Figure 1: Iterations of the FRBO approach demonstrated on a simple 1-D objective function. In [A], we have no knowledge of the failure region, so the FREI acquisition function is nearly identical to the EI acquisition function. In later iterations, our understanding of the failure region improves and regions of potential successes (black bands in PF) weigh the EI function more, and enable us to quickly find good feasible solutions. As shown in [B] and [C], improved estimates of the failure region in later iterations can significantly change our candidate solutions.

into the selection of the next candidate θ, FRBO automatically increases the likelihood of choosing θ values that do not result in instabilities of the closed-loop system, while remaining likely to optimize the cost function.

4. RESULTS

We demonstrate the practical application of the FRBO algorithm for the parameter tuning of PI-ESC controllers to reduce the energy consumption of a variable capacity VCS under standard operating conditions (Burns et al., 2020). Contemporary VCSs often have many variable actuators, such as compressor speed, expansion valve position, and fan speeds; since the number of actuators often exceeds the number of variables regulated by feedback controllers, the remainder of these actuators can be used to optimize the system performance according to a given metric. In this case study, we regulate the cooling capacity provided by the equipment to a user-provided setpoint by controlling the compressor speed, and seek to identify the values of the other actuators that minimize the energy consumption.

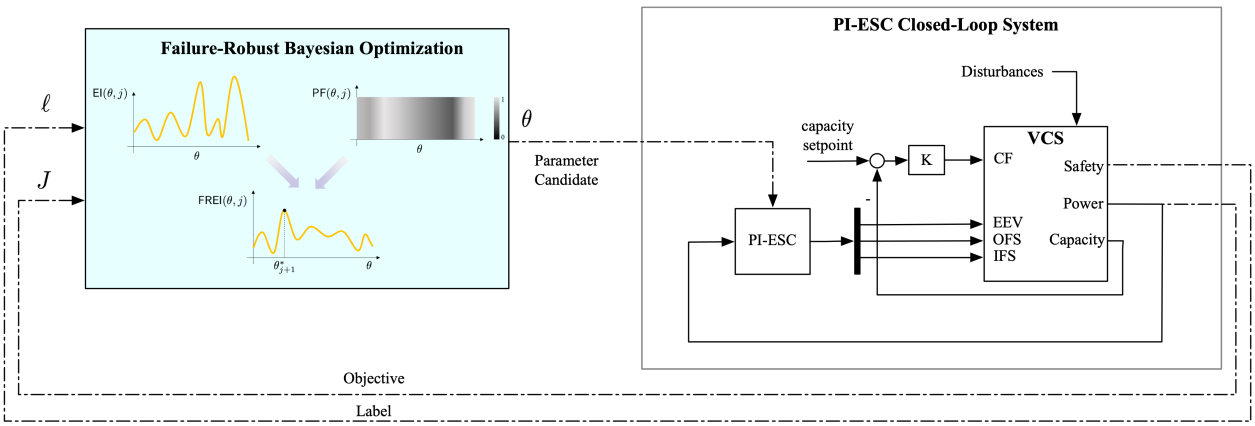

Figure 2: Block diagram of FRBO-tuned PI-ESC for optimizing energy in VCS.

4.1 Experimental Setup

A high-fidelity model for the dynamics of a prototype VCS with the above architecture was developed in the Modelica language (Modelica-Association, 2017). Equation-based models of the compressor, heat exchangers, valves, and fans were constructed using an object-oriented approach and interconnected to form a complete cycle model. For full details about this model, we refer the reader to (Chakrabarty, Danielson, et al., 2021).

As multiple combinations of actuator positions produce the same cooling capacity but differing values of electrical power consumption (Burns & Laughman, 2012), we configured the vapor-compression cycle for realtime optimization via extremum seeking control to identify actuator values that minimize the energy consumption as shown in Figure 2. The value of the compressor frequency is computed by a proportional-integral (PI) controller acting on the difference between a setpoint of 2 kW and the measured value of the cooling capacity; the PI controller gains are designed and fixed offline and demonstrate good regulatory performance under regular operating conditions. This controller was also implemented in Modelica, as the vapor compression system and compressor feedback are thus treated as the optimization target for a proportional-integral extremum seeking controller (PI-ESC) (Burns et al., 2020). The PI-ESC algorithm assigns the EEV position, outdoor fan speed (OFS) and indoor fan speed (IFS) such that a measurement of equipment power is minimized. Assuming that zone setpoint and system disturbances (heat load and ambient temperature) are held constant, the combination of EEV, IFS, and OFS values at steady state are energy-optimal. A block diagram schematic of the FRBO-tuned PI-ESC for energy optimization is provided in Figure 2.

The closed-loop Modelica system model was interfaced to the PI-ESC code using the Functional Mockup Interface (FMI) standard (Modelica-Association, 2019). The Dymola (Dassault Systemes, 2019) environment was used for the initial development of this model, which was then exported as a Functional Mockup Unit (FMU) containing executable simulation code as well as a DAE solver. An advantage of this FMU-based approach is that the original model can be implemented in Modelica, which can efficiently solve large sets of stiff nonlinear differential equations and preserves physics-informed dynamics, while the PI-ESC and FRBO code can be written in Python to leverage existing machine learning tools.

For this example, $n_r = 3$ as the values of three setpoints (EEV, IFS, and OFS) are to be determined by the PI-ESC.

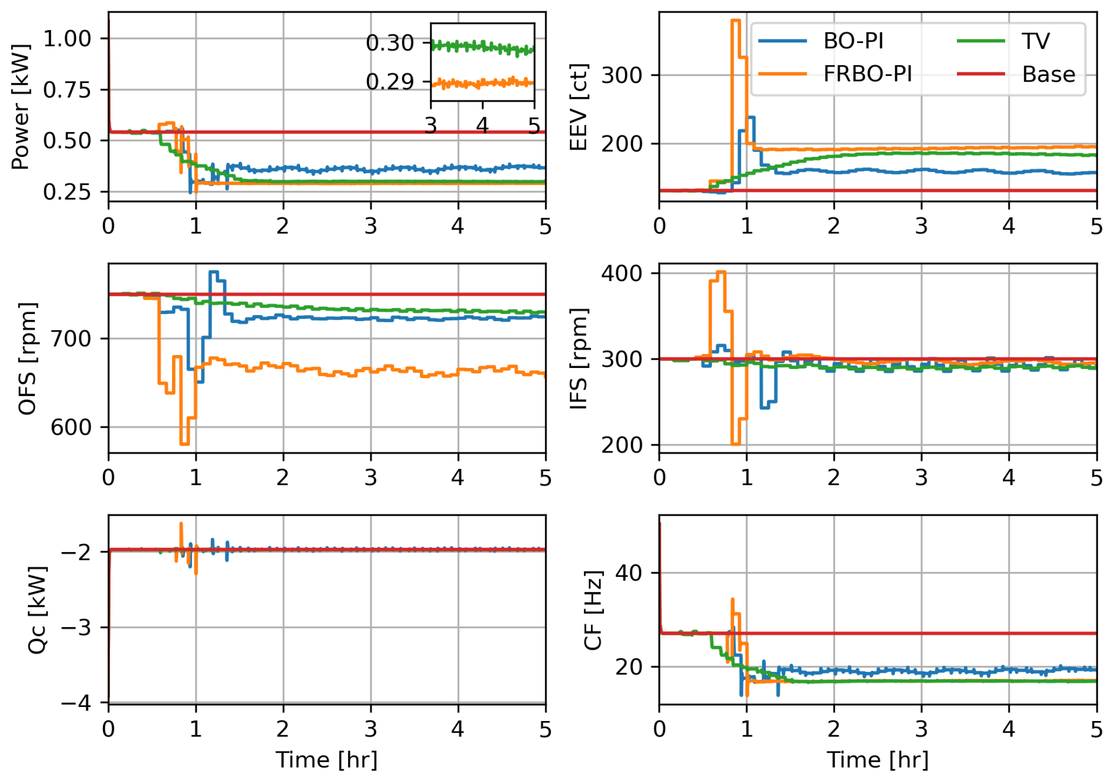

Figure 3: Comparison of closed-loop performance of PI-ESC, TV-ESC, and baseline, for the multivariable space cooling application. (EEV = electronic expansion valve position; OFS = outdoor fan speed; IFS = indoor fan speed; Qc = capacity; CF = compressor frequency.)

Thus, θ is 12-dimensional, as the phase φ is fixed at zero. The search space Θ is identical to that in Example 1, with the exception that the final two intervals for D and ω are $[0.01, 10]^3$ and $[1, 10^3]^3$ because $n_r = 3$, and the search space for k_g is $[0.1, 10^2]^3$. The initial set of states is fixed for this system in order to start the experiment on the VCS at an equilibrium point, so all FRBO costs are calculated for the same initial condition. From this initial time $\mathcal{T}_0 = 0$, we simulate forward to $\mathcal{T}_f = 120$ min to allow the system to enter a 95% settling zone. Each run of 120 min simulation time requires a wall time of 10–15 min due to the large number of internal states in the Modelica model. We select $\mathcal{T}'$ to be the final 20 min, and assume that power measurements are obtained every 60 s. We also use the same hyperparameters for the VGPC failure classifier and GP regressor as in Example 1, except that both are trained for 2000 iterations with a learning rate of 0.01, and allow the FRBO to terminate after 500 iterations. At termination, the parameter set generated by FRBO is given by $\tau_I = 61.3$, $k_g = [4.15, 10.55, 25.42]$, $\alpha = 0.92$, $F = 0.99$, $D = [1.31, 2.15, 4.85]$, $\omega = [10.21, 13.55, 17.96]$.

4.2 Results and Discussion

We demonstrate the potential of the FRBO-tuned PI-ESC by comparing its performance to three different scenarios: a constant baseline set of actuator setpoints, a time-varying ESC (Guay, 2014) that has been previously validated on the heat pump energy optimization problem (Chakrabarty, Danielson, et al., 2021), and a PI-ESC tuned by classical BO. We did not compare the performance of the FRBO-tuned PI-ESC method to a hand-tuned PI-ESC method because we were unable to manually identify parameters that did not result in closed-loop instabilities. Figure 3 illustrates the results of this comparison. While both classical BO and FRBO are allowed the same number of optimization iterations, FRBO converges to a better solution within those iterations while BO wastes many iterations searching for candidate parameters that result in failure. In addition, the power is reduced significantly from the baseline power by both the TV-ESC and our proposed FRBO-tuned PI-ESC, which converges more quickly than the TV-ESC method. This comparison was performed by switching on both ESCs at the 1 hour mark, after the initial transient has disappeared. The focused plot within the top left subplot also illustrates that the FRBO-PI-ESC converges to the optimal power

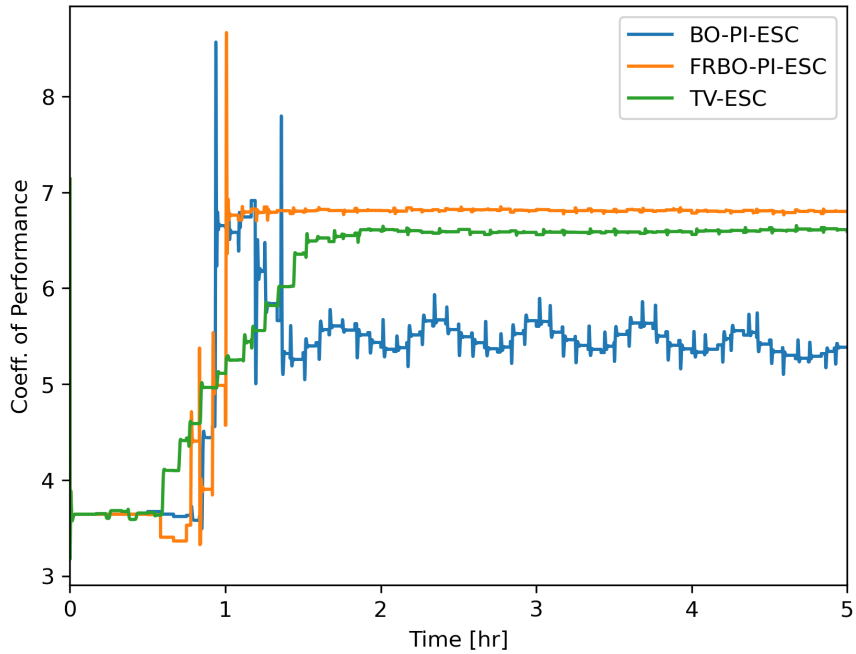

Figure 4: Comparison of coefficient of performance for PI-ESC and TV-ESC. Both classical BO and FRBO solutions have been obtained after allowing 1000 optimization iterations.

value[1] within 1 hour, while the TV-ESC continues to gradually move towards the optimizer and reaches it around 10 hours. In comparison, the BO-PI-ESC does not reach an optimal energy level. This is supported by the time-variation of the EEV, IFS, and OFS: for PI-ESC, the values quickly converge to their final locations after an initial oscillation, whereas the trajectory is much more gradual for the TV-ESC method. Note that the capacity is maintained at 2 kW despite setpoint changes in the ESC outer loop.

The efficiency of the vapor compression system is described by the coefficient of performance (COP), defined as the instantaneous thermal capacity divided by electrical energy consumed by the system. The COP for both PI-ESC and TV-ESC are shown in Figure 4. FRBO-PI-ESC converges to a higher COP compared to TV-ESC (6.81 v. 6.64) and BO-PI-ESC (6.81 v. 5.44). Since the compressor frequency is directly under capacity control, the extremum seeking controllers manipulate the fan speeds and valve position. The outdoor fan speed is reduced, the indoor fan speed largely settles near its initial setting, and the electronic expansion valve opens. This change in the EEV, in particular, will reduce the pressure difference across the compressor and correspondingly reduce its apparent load and increase the cooling capacity. Because the compressor is under capacity feedback control, the higher cooling capacity allows the controller to reduce the compressor speed while continuing to meet the desired capacity setpoint, which accounts for the large increase in COP.

In addition to identifying an improved operating point, PI-ESC converges in less time than the alternate methods. The actuators manipulated by PI-ESC initially experience large swings in amplitude, but this excitation does not cause a reduction in capacity in this case. Scenarios in which these swings are too large for practical considerations can be addressed by reducing their amplitude via appropriate adjustments to PI-ESC parameters, or allowing them to be intercepted by protection logic. The improvement in convergence rate to about 1 hour overcomes important obstacles to wide scale deployments of ESC in VCS applications. Disturbances acting on the system are typically associated with building dynamics or diurnal weather patterns and therefore have timescales slower than 1 hour, implying that PI-ESC could be used to track optimal energy performance while rejecting disturbances in this frequency band.

[1] This was confirmed by exhaustively sampling the setpoint space.

5. CONCLUSIONS

Parameter tuning is an important consideration in the application of ESC methods, as the success of advanced variants turns on their proper identification. The proposed BO technique is a valuable tool in the broader context of model-free real-time optimization system design using ESC. Future work will be focused on the application of the BO tuning approach over a wide range of competing ESC methodologies. This is particularly important in the application of fast ESC techniques for optimizing the operation of vapor compression systems.

REFERENCES

Burns, D. J., & Laughman, C. R. (2012). Extremum seeking control for energy optimization of vapor compression systems. In *International refrigeration and air conditioning conference at purdue.*

Burns, D. J., Laughman, C. R., & Guay, M. (2016). Proportional-integral extremum seeking for vapor compression systems. In *2016 american control conference (acc)* (pp. 2352–2357).

Burns, D. J., Laughman, C. R., & Guay, M. (2020). Proportional–integral extremum seeking for vapor compression systems. *IEEE Transactions on Control Systems Technology*, *28*(2), 403-412. doi: 10.1109/TCST.2018.2882772

Chakrabarty, A., Bortoff, S. A., & Laughman, C. R. (2021). Simulation failure robust Bayesian optimization for estimating black-box model parameters. In *Proc. IEEE International Conference on Systems, Man, and Cybernetics* (pp. 1533–1538).

Chakrabarty, A., Danielson, C., Bortoff, S. A., & Laughman, C. R. (2021). Accelerating self-optimization control of refrigerant cycles with Bayesian optimization and adaptive moment estimation. *Applied Thermal Engineering*, *197*, 117335.

Chua, K. J., Chou, S. K., & Yang, W. M. (2010, Dec). Advances in heat pump systems: A review. *Applied Energy*, *87*(12), 3611-3624. doi: {10.1016/j.apenergy.2010.06.014}

Dassault Systemes. (2019). *Dymola.* Retrieved from `https://www.3ds.com/products-services/catia/products/dymola/`

Guay, M. (2014). A time-varying extremum-seeking control approach for discrete-time systems. *Journal of Process Control*, *24*(3), 98–112.

Guay, M., & Burns, D. J. (2017). A proportional integral extremum-seeking control approach for discrete-time nonlinear systems. *International Journal of Control*, *90*(8), 1543–1554.

Hensman, J., Matthews, A. G., & Ghahramani, Z. (2015). Scalable variational Gaussian process classification. *J. Mach. Learn. Res.*, *38*, 351–360.

Jäschke, J., & Skogestad, S. (2011). Nco tracking and self-optimizing control in the context of real-time optimization. *Journal of Process Control*, *21*(10), 1407–1416.

Khosravi, M., Behrunani, V. N., Myszkorowski, P., Smith, R. S., Rupenyan, A., & Lygeros, J. (2021). Performance-driven cascade controller tuning with bayesian optimization. *IEEE Transactions on Industrial Electronics*, *69*(1), 1032–1042.

Khosravi, M., Eichler, A., Schmid, N., Smith, R. S., & Heer, P. (2019). Controller tuning by bayesian optimization an application to a heat pump. In *2019 18th European Control Conference* (pp. 1467–1472).

Li, P., Li, Y., & Seem, J. E. (2010, May). Efficient operation of air-side economizer using extremum seeking control. *Journal of Dynamic Systems, Measurement, and Control*, *132*(3).

Lu, Q., Kumar, R., & Zavala, V. M. (2020). Mpc controller tuning using bayesian optimization techniques. *arXiv preprint arXiv:2009.14175*.

Modelica-Association. (2017). Modelica specification, Version 3.4.

Modelica-Association. (2019). Functional Mockup Interface for Model Exchange and Co-Simulation, Version 2.0.1.

Neumann-Brosig, M., Marco, A., Schwarzmann, D., & Trimpe, S. (2019). Data-efficient autotuning with bayesian optimization: An industrial control study. *IEEE Transactions on Control Systems Technology*, *28*(3), 730–740.

She, X., Cong, L., Nie, B., Leng, G., Peng, H., Chen, Y., … Luo, Y. (2018). Energy-efficient and-economic technologies for air conditioning with vapor compression refrigeration: A comprehensive review. *Applied Energy*, *232*, 157–186.

Snoek, J., Larochelle, H., & Adams, R. P. (2012). Practical Bayesian optimization of machine learning algorithms. In *Proc. neurips* (pp. 2951–2959).

Wang, W., & Li, Y. (2019). Intermediate pressure optimization for two-stage air-source heat pump with flash tank cycle vapor injection via extremum seeking. *Applied Energy*, *238*, 612–626.