

Accelerating self-optimization control of refrigerant cycles with Bayesian optimization and adaptive moment estimation

Chakrabarty, Ankush; Danielson, Claus; Bortoff, Scott A.; Laughman, Christopher R.

TR2022-010 February 16, 2022

Abstract

This paper presents a model-free self-optimization control algorithm for modulating multiple inputs simultaneously to minimize the power consumption of a Vapor Compression System (VCS). We propose the use of Bayesian Optimization (BO) to warm-start a state-of-the-art Extremum Seeking Control (ESC) algorithm and then accelerate the ESC on-line with Adam, a well-studied algorithm used to train deep neural networks. BO initializes the ESC at conditions favorable for rapid convergence while concurrently learning a surrogate map of VCS power consumption as a function of the inputs. In addition, the warm-start increases the likelihood of attaining a global optimum for locally convex, but globally non-convex, objective functions by identifying regions where the global optimum most likely resides. The proposed algorithm is evaluated using a Modelica model of an air conditioning system with variable compressor speed, an electronic expansion valve and two variable speed fans. We demonstrate the acceleration of this algorithm in simulations of an occupied space with a realistic heat pump model with realistic ambient temperature profiles, variation in heat-load, and different actuation rates. We also show that, in spite of the presence of unknown exogenous disturbances, the proposed algorithm computes better set-points, faster than the ESC. We also observe that the proposed method improves transient performance compared to the state-of-the-art.

Applied Thermal Engineering 2022

© 2022 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Accelerating self-optimization control of refrigerant cycles with Bayesian optimization and adaptive moment estimation

Ankush Chakrabarty^{1,a}, Claus Danielson^b, Scott A. Bortoff^a, Christopher R. Laughman^a

^a*Mitsubishi Electric Research Laboratories, Cambridge, MA*

^b*Department of Mechanical Engineering, University of New Mexico, Albuquerque, NM*

Abstract

This paper presents a model-free self-optimization control algorithm for modulating multiple inputs simultaneously to minimize the power consumption of a Vapor Compression System (VCS). We propose the use of Bayesian Optimization (BO) to warm-start a state-of-the-art Extremum Seeking Control (ESC) algorithm and then accelerate the ESC online with Adam, a well-studied algorithm used to train deep neural networks. BO initializes the ESC at conditions favorable for rapid convergence while concurrently learning a surrogate map of VCS power consumption as a function of the inputs. In addition, the warm-start increases the likelihood of attaining a global optimum for locally convex, but globally non-convex, objective functions by identifying regions where the global optimum most likely resides. The proposed algorithm is evaluated using a Modelica model of an air conditioning system with variable compressor speed, an electronic expansion valve and two variable speed fans. We demonstrate the acceleration of this algorithm in simulations of an occupied space with a realistic heat pump model with realistic ambient temperature profiles, variation in heat-load, and different actuation rates. We also show that, in spite of the presence of unknown exogenous disturbances, the proposed algorithm computes better set-points, faster than the ESC. We also observe that the proposed method improves transient performance compared to the state-of-the-art.

Keywords: Energy efficiency, Refrigerant cycles, Extremum seeking, Data-driven methods, Optimization

1. Introduction

The significant energy consumption of HVAC systems, as a fraction of both the average total building energy consumption and of energy usage on global scales, has long motivated efforts to improve their energy efficiency [1–4]. As a means of reducing energy consumption, optimization of internal control set-points of HVAC systems provides a promising path that does not require a complete redesign of the building infrastructure. In particular, *self-optimizing control* aims to find optimal set-points such that feedback control loops can provide near-optimal operation under a range of operating conditions and exogenous inputs [5]. Machine learning and data-driven optimization algorithms have a central role to play in self-optimization of building/HVAC set-points due to their ability to represent and exploit surrogate models of complex cost functions that arise from multi-physical interactions. While machine learning methods have been used previously to construct building and equipment models for performance analysis and prediction [6–11], they have seldom been used

for self-optimization. These methods can profitably be used to optimize the internal set-points of energy systems by leveraging updated operational data to adapt efficiently the behavior of the system to site-specific conditions, and can accordingly improve the energy efficiency of individual systems without labor intensive tuning that might be performed by either the manufacturer or a field technician.

HVAC control systems are largely designed to regulate process variables to set-points, so appropriate set-points for these variables must be identified to achieve reliable and energy-efficient operation. We can therefore formulate a real-time set-point optimization problem with the objective of finding set-points that optimize a measurable performance function. As it is typically difficult and expensive to build accurate behavioral models of field-installed equipment, we consider a so-called ‘model-free’ setting in which there is no given model of the optimization cost, so that the gradients can only be computed numerically. One class of model-free algorithms is Extremum Seeking Control (ESC), which can be considered as a data-driven gradient descent algorithm that estimates a local gradient online using excitation signals and filtering [12, 13]. ESC has a major advantage in that one does not have to wait for the measured output to reach its steady-state; the algorithm can operate as long as the system satisfies some theoretical properties. This implies that ESC has the potential to converge quickly, compared to an algorithm that

Email addresses: chakrabarty@merl.com (Ankush Chakrabarty), cdanielson@unm.edu (Claus Danielson), bortoff@merl.com (Scott A. Bortoff), laughman@merl.com (Christopher R. Laughman)

¹Corresponding author. Phone: +1 (617) 758-6175.

would require waiting for the system to reach steady-state. These advantages have resulted in ESC being proposed for energy-saving in many HVAC applications including vapor compression cycles [14–17], air handling units [18], chilled water plants [19], and variable refrigerant flow systems [20]. Most of these results deal with manipulating single set-points, although [21] is an exception, where the authors consider multivariate ESC with three inputs.

The designer faces two challenges in applying ESC to HVAC set-point optimization. First, tuning the algorithm parameters, such as the gradient step-size, is difficult because by definition a model is lacking. In the experience of the authors, step-sizes tend to have bounded ranges in which closed-loop convergence is achieved, and values outside these ranges can result in divergence. Small step-sizes can result in slow convergence, while large step-sizes can cause oscillation around the optimum. Furthermore, the objective function is nonlinear and non-convex in the domain of allowable set-points. Thus, selecting a single fixed step-size could be overly conservative. Second, for the HVAC problem the objective functions is at least locally convex in a neighborhood of the optimal set-points. This implies that if the ESC algorithm is initialized within this locally convex region, it will converge to the optimal set-point. However, in a model-free setting, and especially with a large number of tunable set-points, it is difficult to guarantee that the initial condition is sufficiently close to the global optimum. Therefore, in this paper, we propose a method for step-size adaptation based on gradient estimates obtained online by the ESC algorithm, alleviating the need to perform a large number of trials to select the step-size *a priori*. We append this adaptive ESC with a machine learning algorithm that explores the admissible set-point space and constructs a surrogate model which identifies areas that most likely contain the global optimum. The machine learning algorithm warm-starts the adaptive ESC by instantiating the ESC algorithm with an initial set-point that has a high-likelihood of being sufficiently close to the optimal set-points. The combination of step-size adaptation and warm-starting is expected to result in reliable and accelerated convergence.

Additionally, ESC algorithms use dithering and adaptive filtering to construct a gradient estimate in a data-driven manner, but by design do not create a model of the optimization cost/reward function e.g., the map from inputs to power consumption. This is a missed opportunity and a waste of data, since such a ‘surrogate’ model could accelerate convergence, and indeed be reused for other purposes, such as monitoring performance over time. While linear and quadratic surrogate models have been explored in the past, the linear models are not complex enough to capture nuances of the true function, and the performance of quadratic models are adversely affected by noise in the measurements, which results in poor estimation of the Hessian. Conversely, by switching to a robust optimization methodology like Bayesian Optimization to model the surrogate function, the uncertainty and noise is captured

within confidence bounds around predictions of a Gaussian process regression model, and decisions can be made robustly via this uncertainty-capturing model. These opportunities and potential benefits of surrogate modeling to the building industry have been described in detail [22]. In this paper, we propose adapting a popular momentum-based acceleration control law called Adam [23], which is widely used to train non-convex optimization problems arising in deep learning [24], to the ESC problem. We also extend the classical ESC algorithm to non-convex problems by using a warm-start mechanism based on Bayesian optimization [25], a machine learning algorithm that has recently been used for optimizing tubular solar stills [26]. This intrinsically constructs surrogate models, which can speed up the convergence of ESC.

We develop these new methods for set-point optimization of a modern variable-capacity Vapor Compression System (VCS). These systems have many actuators, including fan speeds, the compressor speed, and electronic expansion valve positions. There are often more actuators than there are variables to be regulated. In this case, we can use a subset of the actuators to minimize the system power consumption while the others are used to regulate process variables. The process of empirically determining these optimal set-points from experiments is time-consuming and expensive, making the application of standard ESC methods to this problem cost-prohibitive. These concerns motivate the development of methods that can find optimal set-points in a time-efficient manner. Moreover, the time-varying behavior of these systems during practical operation, such as changes in heat load due to room occupancy or ambient temperature variations, necessitate the development of methods that can quickly converge to new set-points in response to such variations.

The main contributions of this paper are as follows:

- (i) We provide a machine learning mechanism for using data on-line to construct a surrogate model of a power function without requiring an analytic model of the function or of the underlying system dynamics;
- (ii) We provide a general learning-based warm starting method to accelerate convergence of extremum-seeking algorithms;
- (iii) We propose a novel solution to self-optimize non-convex cost functions and concurrently select multiple actuator set-points;
- (iv) We demonstrate that we can learn optimal set-points at multiple operating modes and do not require re-learning of gradients from scratch when the operating mode changes.

The rest of the paper is organized as follows. We provide some VCS background and context in Section 2. A fast version of ESC algorithm that has been considered previously [16] is presented in Section 3, and our specific modifications based on Adam and Bayesian optimization are then presented. Details of our Modelica implementation of the algorithm and results including robustness




Figure 1: Schematic diagram of vapor compression system with capacity control feedback (blue) actuating the compressor speed, and proposed Bayesian optimization warm-started Extremum Seeking Control (ESC, red) actuating the EEV and two fan speeds. Some additional process measurements and feedback are not shown for simplicity.

analysis is presented in Section 4. We also present a study of the proposed algorithm in realistic scenarios with time-varying ambient temperature and heat loads, and demonstrate that our proposed improvements to the ESC approach results in significant power savings. Conclusions are offered in Section 5.

2. VCS Background

A typical air-to-air VCS includes four main components: an evaporating heat exchanger, a compressor, a condensing heat exchanger, and an expansion valve, as shown in Figure 1. Our cycle also includes an accumulator to ensure that vapor refrigerant flows into the compressor. This component does not affect the system dynamics in this work. These components work together to transfer thermal energy from one environment to another via the circulation of a refrigerant under temperature and pressure gradients. The operation of this cycle is well-known, but is briefly described here. Typically, superheated refrigerant vapor at an elevated pressure and temperature leaves the compressor in a state at which the refrigerant saturation temperature at the condensing pressure is higher than the temperature of the surrounding air, causing the refrigerant to condense and transfer thermal energy to the surrounding environment. It exits the heat exchanger as a sub-cooled liquid and expands through the expansion valve and enters the evaporator as cold, two-phase refrigerant. The lower operating pressure of the evaporator results in a commensurately low refrigerant saturation temperature, and the higher temperature of the surrounding air causes the refrigerant to evaporate and transfer thermal energy from the environment into the fluid. The refrigerant then exits this second heat exchanger as a low-pressure superheated vapor, after which it enters the compressor to begin the cycle anew.

Modern high-performance VCS employ variable speed compressors, variable speed fans, and electronic expansion valves (EEVs) in order to improve capacity and energy effi-

ciency across a broad range of operating conditions. These degrees of freedom require a concomitant use of control and optimization strategies to achieve the desired system behavior and energy efficiency in the face of time-varying operating conditions and disturbances. Electrical power is consumed by the fans and compressor. While the fan power consumption can be modeled as a univariate cubic polynomial of fan speed, compressor power consumption is a highly nonlinear function of the inlet and outlet refrigerant states (pressure, flow, temperature) as well as the compressor speed. These boundary conditions are coupled to the dynamic behavior of the evaporators, which in turn are dynamically coupled to the thermodynamics of their environment. As a result, a first principles model of the system that includes compressor power consumption is a very large set of nonlinear differential-algebraic equations, with many parameters that in practice must be identified from data. If such a model is available, then direct gradient feedback can be used to drive the available actuators exponentially fast to their minimum power values [27]. But, in practice, this is usually not the case. On the other hand, the power consumption of both the compressor and fans is typically available from measurement, motivating this work. Furthermore, the steady-state relationship between inputs and total VCS power consumption is known to possess a unique, global minimum, although the map is highly nonlinear and non-convex.

Vapor-compression systems can be applied in a variety of different scenarios. Perhaps the most common application utilizes the VCS to regulate the temperature of an occupied space to a set-point while rejecting an unknown, time-varying heat load disturbance. In this case, the controller modulates the compressor speed to regulate the zone temperature, while the EEV and outdoor fan speed are used to regulate other process variables such as the compressor discharge temperature, and the indoor fan speed may be modulated to minimize system power consumption. In another scenario, typically used for equipment energy efficiency rating in a laboratory [28], the compressor speed is regulated by feedback to produce a constant cooling capacity, while the values of other three actuators are modulated to minimize power consumption subject to enforcing some process variable constraints. This is the situation diagrammed in Figure 1, although we note that the subject of enforcing those constraints is beyond the scope of this paper.

3. Self-optimizing set-points via extremum seeking control algorithms

In this section we begin by briefly summarizing the time-varying ESC (TV-ESC) algorithm presented in [29], used for finding local optima in convex functions in a model-free manner. Subsequently, we describe the novel acceleration methods proposed in this paper. First, adapting the ESC gain using gradient information on-line, and second

by enabling a global search with a Bayesian Optimization warm-start.

3.1. Time-Varying ESC

Consider a partially closed-loop VCS with dynamics

$$x_{t+1} = f(x_t, \nu_t, \theta_t) \quad (1)$$

where $t \in \mathbb{N}$ denotes the time index, $x_t \in \mathbb{R}^n$ denotes the state of the system, $\nu_t \in \mathbb{R}^\nu$ denotes exogenous noise or disturbance inputs, and $\theta_t \in \Theta$ denotes a vector of set-points that will be designed to drive the system to an equilibrium where the measured power output

$$y_t = J(x_t) \quad (2)$$

attains its minimum. For example, θ_t could be the position of an electronically/linear actuated valve, or fan speeds, and/or combinations thereof. Partially closed-loop means that some of the VCS actuators may be driven by feedback, such as in Figure 1, while the remainder comprise θ_t . In this paper, neither f nor J is known to the designer, and x_t is not measured: this is a distinction from a few ESC methods, which assume knowledge of f but not J .

For fixed values of θ_t and ν_t , in the steady-state (i.e., when $x_{t+1} = x_t$), we may write the input-output map from θ_t to y_t as

$$y_t = \bar{J}(\theta_t).$$

We assume that \bar{J} is (at least locally) a strongly convex function of θ_t in some closed set $\Theta_0 \subset \Theta$ where $\Theta \subset \mathbb{R}^p$ denotes a known set of attainable actuator set-points (for example, a range of fan speeds), for constant ν_t . (Formally, \bar{J} is also a function of ν_t , but we ignore this argument to simplify the notation.) Local strong convexity implies that there exists a $\theta^* \in \Theta_0$ such that $\nabla \bar{J}(\theta^*) = 0$; furthermore, there exists a $\chi > 0$ such that $\nabla^2 \bar{J}(\theta) \succeq \chi I$ for every $\theta \in \Theta_0$.

As stated in the previous subsection, we wish to design a control law that drives $\theta_t \rightarrow \theta^*$ and $y_t \rightarrow \bar{J}(\theta^*)$ as $t \rightarrow \infty$. Since by assumption a mathematical representation of $\bar{J}(\theta)$ is not available, we cannot use standard gradient-based optimization methods to derive the optimal set-point $\theta^* = \arg \min_{\theta} \bar{J}(\theta)$. The problem of optimizing θ is further exacerbated by the presence of the nonlinear dynamics f .

A potential solution for this problem is by the use of multivariate ESC. A variant of ESC that has been investigated in prior work and demonstrated good optimization capability for univariate ESC in room air conditioners is the TV-ESC algorithm [29]. In the TV-ESC formulation, a gradient-based controller is used that has the form

$$\theta_{t+1} = \theta_t - k_g g_t + d_t,$$

where g is an estimate of the gradient of \bar{J} , and $k_g \in \mathbb{R}$ is the control gain or step size. In practice, the dither signal d_t is a periodic signal containing p distinct frequencies to provide persistence of excitation.

The gradient estimate g_t is generated in a data-driven manner. Specifically, we compute the increments in the cost function

$$\Delta J_t \triangleq \bar{J}(\theta_t) - \bar{J}(\theta_{t-1})$$

along with the incremental set points

$$\Delta \theta_t \triangleq \theta_t - \theta_{t-1}.$$

The gradient is estimated from the linear equation

$$\begin{bmatrix} \Delta J_{t-N_d+1} \\ \vdots \\ \Delta J_t \end{bmatrix} = \begin{bmatrix} \Delta \theta_{t-N_d+1}^\top \\ \vdots \\ \Delta \theta_t^\top \end{bmatrix} g_t,$$

using a history of N_d data points. As the optimization problem is not static, the most recent data points contribute most strongly to the gradient. Therefore, we implement a recursive filter to estimate the gradient, and include a forgetting factor $\alpha \in (0, 1)$ to assign more importance to recent data points rather than data of the past. This recursive filter is described by

$$e_t = \Delta J_t - \Delta \theta_t^\top g_{t-1}, \quad (3a)$$

$$K_t = \frac{P_{t-1} \Delta \theta_t}{\alpha + \Delta \theta_t^\top P_{t-1} \Delta \theta_t}, \quad (3b)$$

$$P_t = \frac{1}{\alpha} (P_{t-1} - K_t \Delta \theta_t^\top P_{t-1}), \quad (3c)$$

$$g_t = g_{t-1} + K_t e_t, \quad (3d)$$

initialized with $P_{-1} = \alpha^{-1} I$ and $g_{-1} = 0$.

We have observed that during implementation, resetting the matrix P_t to P_{-1} after a user-defined number of iterations prevents ill-conditioning of the matrices in the gradient estimator. Since the recursive filter needs a burn-in period [30], denoted τ , to generate meaningful estimates g , our implementation of the ESC control law incorporates gradient estimates only after the burn-in period, as follows:

$$\theta_{t+1} = \begin{cases} \theta_t + d_t & \text{for } t \leq \tau, \\ \theta_t - k_g g_t + d_t & \text{for } t > \tau. \end{cases} \quad (4)$$

3.2. Acceleration Via Gradient-Based Step-Size Adaptation

Unfortunately, evaluating $\bar{J}(\theta_t)$ requires data obtained by running experiments to steady-state, which is expensive and time-consuming. Therefore, we propose a mechanism to accelerate the ESC algorithm by modifying the gain of the ESC based on gradient statistics. We begin by providing an example where maintaining a constant gain k_g can result in degradation of optimization performance, even for a static problem.

3.2.1. A Motivating Example

We illustrate the importance of step-size selection by implementing ESC on a static cost function

$$\bar{J}(\theta) = \left(\theta - \begin{bmatrix} 1 & 2 & 3 \end{bmatrix}^\top \right)^\top \left(\theta - \begin{bmatrix} 1 & 2 & 3 \end{bmatrix}^\top \right).$$

Keeping initial conditions and all other parameters constant, we run the ESC algorithm with multiple step-sizes $k_g \in \{0.001, 0.025, 0.1\}$ and plot $\|\theta_t - \theta^*\|$ in Figure 2, where $\theta^* = [1, 2, 3]$. While choosing a gain of $k_g = 0.025$ exhibits excellent closed-loop behavior in terms of optimizing the cost and convergence speed, making k_g larger to 0.1 result in significant jumps (around 4000 steps) despite initially good performance. Increasing k_g further leads to diverging trajectories. Conversely, decreasing k_g below 0.01, as evident from the $k_g = 0.005$ trajectory, results in good asymptotics but a slow convergence rate. To reiterate, this simple static example shows that increasing k_g results in oscillatory behaviour of the algorithm and unreliable convergence, and making k_g too small results in sluggish convergence. These effects are worsened when the underlying optimization problem contains dynamics in the map from θ to \mathcal{J} . Since we do not have a model of the cost function, we cannot use model information, such as analytical gradients, to change k_g online. Thus, we need to be able to adapt the step-size k_g based on gradient information obtained online, which we discuss next.

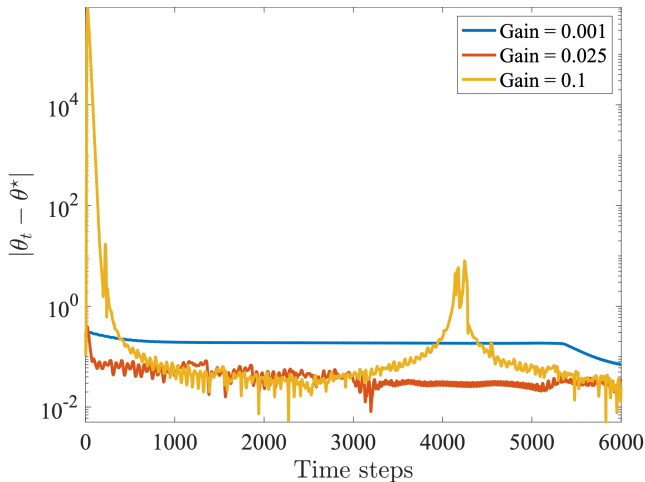


Figure 2: Importance of step-size selection on an exemplar static map (y -axis is in log-space).

3.2.2. Controller Update with Step-Size Adaptation

The step-size can be determined by using the history of the gradient estimates. Let m_t and v_t denote the mean and standard deviation computed using the last N gradient estimates $g_{t-N+1:t}$, respectively. Our proposed controller has the form

$$\theta_{t+1} = \theta_t - k_t \frac{m_t}{\sqrt{v_t} + \epsilon} + d_t \quad (5a)$$

$$m_{t+1} = \beta_1 m_t + (1 - \beta_1) g_t \quad (5b)$$

$$v_{t+1} = \beta_2 v_t + (1 - \beta_2) g_t \odot g_t, \quad (5c)$$

where

$$k_t = k_g \frac{\sqrt{1 - \beta_2^t}}{1 - \beta_1^t}, \quad (5d)$$

where \odot denotes the element-wise product operator, m is the first moment (weighted mean) of the gradients obtained using the past N measurements, v is the second moment, k_g is a user-defined step size parameter, and ϵ is a user-defined parameter that ensures that the denominator in (5a) is always positive, since the $\sqrt{v_t}$ term could be zero at points where the function has been measured and noise was negligible. The first- and second-order moments m and v are estimated via a low-pass filter of the history of moments of the gradients; the filter coefficients are $\beta_1 > 0$ and $\beta_2 > 0$.

The structure of this controller is motivated by moment-based adaptive gradient algorithms such as AdaGrad, RMSProp, and Adam, that are commonly used in stochastic gradient descent methods [24]. While our formulation does not contain stochasticity in the gradient computation in the traditional sense (arising from randomly selecting directions in which to compute the gradient), recent work has demonstrated the effectiveness of these adaptive gradient methods [31, 32] in eliminating the need to hand-tune hyperparameters such as step-sizes. In the ESC setting, this is analogous to *reducing the need to fine-tune the step-size in the controller update* (4), which is especially advantageous in practice, when selecting k_g in a model-free setting requires multiple trials.




Figure 3: Illustration of adaptive step-size selection in adaptive gradient methods such as Adam. The blue surface shows an exemplar nonconvex cost function.

The principle by which the control gain adapts is illustrated in Fig. 3. Concretely, we consider two cases: the top left trajectory of the figure, where the estimated gradients are consistently in the same direction with little variation, therefore, the mean estimated gradient m_t is similar over iterations, and the variance v_t is small. Therefore, the element-wise ratio $m_t / \sqrt{v_t} + \epsilon$ is large, and allows the controller to take a more aggressive stance in the descent direction. Another case is illustrated by the bottom right trajectory, wherein the estimated gradients are noisy, and therefore the variance v_t is large: this ensures that the ratio $m_t / \sqrt{v_t} + \epsilon$ is small, result in a smaller control gain

and cautious movement in the descent direction.

Another advantage of this method is that it performs well when the number of set-points is large, such as operating conditions of a power grid or a cluster of buildings. In such settings, one cannot always change the entire vector of set-points or obtain all the measured power outputs synchronously. Therefore, only a subset of set-points can change, and only a subset of gradients can be estimated. By evaluating the statistics of a history of gradients as in the Adam algorithm, one can get an estimate of the full gradient vector. Thus, adaptive gradients like Adam are well-suited to these large-scale problems since they are primarily used for mini-batch gradient-based updates in training deep neural networks, where random subsets of gradients are used to find search directions that lead to a local optimum.

3.3. Acceleration Via Bayesian Optimization Warm-Start

In practice, the function \bar{J} may be non-convex or convex but with some directions or regions exhibiting small gradients resulting in slow convergence of un-accelerated ESC methods. Note that even though \bar{J} is non-convex, a gradient-based method can still converge to the unique minimum if the function is locally convex and the initial guess is made within this region of local convexity. Finding the optimal set-points in these scenarios involves identifying a good initial guess, ideally in a region where the function is locally convex Θ_0 , then using gradient estimates within that region to identify the optimal solution. Without a good initial guess the TV-ESC may converge very slowly. But, using a machine-learning method to approximate \bar{J} with operational data is very useful because the approximated function can be used as a surrogate for the true (but unknown) \bar{J} . To these ends, we employ a Bayesian optimization (BO) framework to accelerate and warm-start our TV-ESC algorithm. We now detail the BO warm-start procedure.

BO methods work in two phases. First, regression models are used to construct probabilistic approximations of the steady-state power function \bar{J} with respect to the set-points θ . This probabilistic surrogate model can be used to generate a predictive distribution for \bar{J} at each point of the input space Θ . Second, the predictive distribution is used to guide subsequent search directions, with focus given mainly to subregions of Θ where the function is expected to contain the most information about the global solution of the optimization problem $\min_{\Theta} \bar{J}$. Widely used models for constructing these maps include Gaussian processes (GPs) [33], student- t processes [34] or Bayesian neural networks [35].

In this work, we propose the use of GPs to define a prior distribution over functions. The underlying assumption made is that the cost function \bar{J} to be optimized has been generated from such a prior distribution, characterized by a zero mean and a kernelized covariance function $\mathcal{K}(\theta, \theta')$. The covariance function \mathcal{K} is singularly responsible for defining the characteristics of the function such as

smoothness, robustness to additive noise, and so on. While many kernel functions are available, we have found (empirically) that the Matérn 3/2 function provides a good approximation of steady-state power functions without over-smoothing, which is a drawback of squared exponential kernel families. The Matérn 3/2 kernel is defined by

$$\mathcal{K}(\theta, \theta') = \sigma_0^2 \left(1 + \frac{\sqrt{3}\|\theta - \theta'\|}{l} \right) \exp \left(-\frac{\sqrt{3}\|\theta - \theta'\|}{l} \right),$$

where l is a vector known as length-scale, which controls the smoothness of the functions induced by the kernel in each dimension of θ , and σ_0 controls the range of values, i.e., support of the GP distribution.

Assume that we have already evaluated the objective at N_θ input samples. Let this training data be denoted by $\{(\theta_k^D, \bar{J}(\theta_k^D) + \nu_k)\}_{k=1}^{N_\theta}$, where $\nu_k \sim \mathcal{N}(0, \sigma_n^2)$ is additive white noise in the measurement channel with zero-mean and unknown covariance σ_n^2 . Given a set of hyperparameters (length scale l , kernel variance σ_0 , and noise variance σ_n) for a pre-selected kernel, one can compute the matrices

$$K_D(\theta) = [\mathcal{K}(\theta, \theta_1^D) \quad \dots \quad \mathcal{K}(\theta, \theta_{N_\theta}^D)]$$

and

$$K_D = \begin{bmatrix} \mathcal{K}(\theta_1^D, \theta_1^D) & \dots & \mathcal{K}(\theta_1^D, \theta_{N_\theta}^D) \\ \vdots & \ddots & \vdots \\ \mathcal{K}(\theta_{N_\theta}^D, \theta_1^D) & \dots & \mathcal{K}(\theta_{N_\theta}^D, \theta_{N_\theta}^D) \end{bmatrix}.$$

With these matrices, we can compute a posterior distribution characterized by a mean function $\mu(\theta)$ and variance function $\sigma^2(\theta)$ given by

$$\mu(\theta) = K_D(\theta)^\top (\mathcal{K}_D + \sigma_n^2 I)^{-1} J(\theta), \quad (6a)$$

$$\sigma^2(\theta) = \mathcal{K}(\theta, \theta) - K_D(\theta)^\top (\mathcal{K}_D + \sigma_n^2 I)^{-1} K_D(\theta). \quad (6b)$$

Now, 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 l , σ_0 and σ_n . There are a variety of methods to obtain these hyperparameters, but the most common method involves maximizing the log-marginal likelihood function

$$\mathcal{L}(\sigma_0, \sigma_n, l) = -\frac{1}{2} \log |\mathcal{K}_n| - \frac{1}{2} J(\theta)^\top \mathcal{K}_n^{-1} J(\theta) + \frac{p}{2} \log 2\pi \quad (6c)$$

with $\mathcal{K}_n = \mathcal{K}_D + \sigma_n^2 I$. This problem is non-convex but can be solved using quasi-Newton methods or adaptive gradient methods [33].

The exploration-exploitation trade-off in BO methods is performed via an acquisition function $\mathcal{A}(\cdot)$. The acquisition function uses the predictive distribution given by the GP to compute the expected utility of performing an evaluation of the objective at each set-point θ . The next set-point at which the objective has to be evaluated is given by $\theta_{N_\theta+1} := \arg \max \mathcal{A}(\theta)$. As this function only depends on the GP approximated function and not on the actual

objective \bar{J} , the maximization of $\mathcal{A}(\cdot)$ involves computing (6) rather than expensive function evaluations. In this work, we use an expected improvement (EI) acquisition function, given by

$$\mathcal{A}_{\text{EI}}(\theta) = \sigma(\theta) (\gamma(\theta)\Phi(\gamma(\theta)) + \phi(\gamma(\theta))), \quad (7)$$

where $\phi(\cdot)$ is the density function of the zero-mean one-variance normal distribution, $\Phi(\cdot)$ is its cumulative distribution function, and

$$\gamma(\theta) = \frac{\min\{\bar{J}(\theta_k) + \nu_k\}_{k=1}^{N_\theta} - \mu(\theta)}{\sigma(\theta)}. \quad (8)$$

This acquisition function estimates the expected improvement of the steady-state power generated by the next set-point versus the current best solution. An efficient way of computing the maximum of this acquisition function is by generating random samples on Θ , computing \mathcal{A}_{EI} for each sample, and choosing the sample maximum as the next set-point.

After a suitable number of iterations N_θ , the GP regressor is expected to learn the underlying function \bar{J} and the best solution obtained thus far by the acquisition function serves as a warm-start to the TV-ESC algorithm. The selection of N_θ is a design decision: it is informed usually by practical considerations such as the total amount of evaluations of \bar{J} , that is, the total amount of data one can feasibly collect.

While we do not explicitly consider convergence of the BO-ESC method, we can present some intuition into why the BO-ESC algorithm is expected to converge and explain why a rigorous proof will be limited in scope. The BO-ESC algorithm works on the premise that the BO warm-start collects enough data to ensure that a local region of convexity is identified within the domain of the cost function \bar{J} . If the BO enables the ESC algorithm to start in such a domain of local convexity, then the proofs of [29] are valid, and the results on asymptotic convergence and stability of the BO-ESC are identical to that of the TV-ESC. The challenge lies in proving that the BO correctly identifies a solution in a region of local convexity *with finite data*, since BO proofs are typically asymptotic in nature [36], and involve assumptions that are hard to check in practice; for example, that the cost is globally Lipschitz or that the minimum is known to the designer *a priori*. Since we cannot afford to sample infinitely in practical settings, we refrain from a rigorous treatment of the convergence properties of the BO-ESC algorithm.

Pseudocode for the proposed algorithm is provided in Algorithm 1.

4. Results and Discussion

We apply this method to two candidate test problems that reflect a range of applications. The first test problem is adapted from a common scenario in which the VCS

performance is measured to determine compliance with government efficiency standards. Such standards [37] are used to compare energy performance of different models and manufacturers. The methodology for these tests requires the system to regulate the cooling capacity to a specified set-point, which can be achieved by using a feedback controller to modulate the compressor speed. Disturbances are held constant during the test. We use the self-optimization methods to modulate the other three inputs for this system (the expansion valve position and two fan speeds) to minimize the power consumption, as shown in Figure 1. The second test problem is the more common scenario in which the VCS is used to regulate the room temperature of an occupied space in the face of a slowly varying heat load and ambient temperature disturbances. This second problem incorporates two separate feedback loops, one to regulate room temperature and the other to regulate a process variable.

A high-fidelity model of the dynamics of a prototype VCS was constructed using the Modelica language [38]. Geometric and operational parameters, such as heat exchanger geometries, heat transfer coefficients, and frictional pressure drop correlations, were adapted and slightly modified from experimentally-calibrated models for the purposes of generalization. Equation-oriented models of the compressor, expansion valve, accumulator, and both heat exchangers were interconnected to form a cycle model. The dynamic behavior of the system, represented by f in (1), is dominated by that of the heat exchangers. These are represented as finite volume models that spatially discretize the mass, momentum, and energy balance equations, respectively, into a set of differential algebraic equations (DAEs), which also characterized the conjugate heat and mass transfer on both the refrigerant-side and air-side. This discretization was performed along the length of the heat exchanger tubes, while three-dimensional variations on the air-side of the heat exchangers were represented with a large set of equations interrelating the boundary conditions. A fundamental equation-of-state model was also used to describe the algebraic coupling between thermodynamic property variables of the refrigerant, such as temperature, density, pressure, and specific enthalpy. Sixteen volumes were used for the condensing heat exchanger, while twenty-seven volumes were used for the evaporating heat exchanger, owing to its larger number of tubes. Nonlinear algebraic models were used for the compressor and expansion valve. A nonlinear algebraic model is also used for $J(x_t)$ in (2). This resulted in a cycle model with 12,114 equations. Additional information on the modeling approach used for the cycle components in this work can be found in [39].

The model was interfaced to the accelerated ESC code using the Functional Mockup Interface (FMI) standard [40]. The model was first developed in the Dymola [41] environment (see Figure 4), and then exported as a Functional Mockup Unit (FMU). The FMU contains both executable simulation code as well as a DAE

Algorithm 1 PROPOSED ADAM-ACCELERATED ESC WITH BO WARM-START

Require: Bounds on admissible set-points Θ

Require: Time period between set-point changes

Bayesian Optimization Warm-Start

Require: Number of burn-in BO iterations, N_0

▷ default: 5

Require: Number of BO iterations for warm-start, N_θ

▷ default: > 75

Require: Acquisition function \mathcal{A}

▷ default: expectation-improvement

Require: Kernel \mathcal{K} for GP regression

▷ default: Matérn 3/2

Require: Samples Θ' within Θ to evaluate acquisition function

- 1: **for** k in $1 : N_0$ **do**
 - 2: Randomly select set-points θ_k and evaluate cost function $\bar{J}(\theta_k)$ via simulation/experiment
 - 3: **end for**
 - 4: **for** k in $N_0 + 1 : N_\theta + N_0$ **do**
 - 5: Fit GP with kernel \mathcal{K} using data $\{\theta_k, \bar{J}(\theta_k)\}$ using (6)
 - 6: Evaluate acquisition function \mathcal{A} on samples Θ'
 - 7: Compute next best $\theta_{k+1} \leftarrow \arg \max \mathcal{A}$
 - 8: Evaluate $\bar{J}(\theta_{k+1})$ via simulation/experiment
 - 9: **end for**
-

Adam-Accelerated TV-ESC

Require: ESC controller gain, k_g

Require: Data window length for gradient estimation, N_d

Require: Dither signal, d

▷ small sinusoid

Require: Initial gradient filter parameters, P_{-1}, α

▷ default: $10^3 I, 0.7$

Require: Adam parameters, β_1, β_2

▷ default: 0.9, 0.999

10: **for** k in $-1 : N_d$ **do**

11: $g_t \leftarrow$ gradient estimate using (3)

12: $\theta_t \leftarrow$ update set-point using (5)

13: **end for**

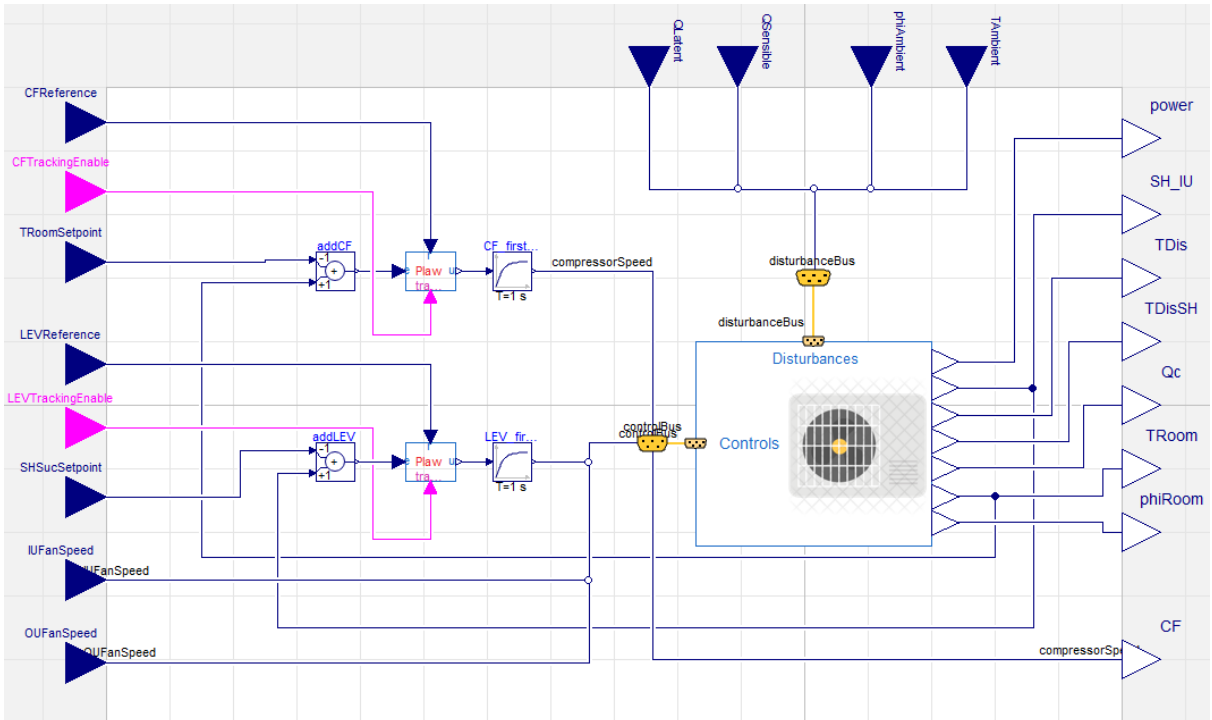


Figure 4: Iconic view of the Modelica implementation of the VCS.

solver, and was then connected to the ESC code written in Python. Such a platform is advantageous because it al-

lows the model to be formulated in Modelica, which is generally superior for numerical simulation of large stiff non-

linear differential equations, and coupled to the Bayesian optimization code, which is built using existing machine learning tools in Python (see the Appendix for a list of critical Python dependences).

4.1. Energy Efficiency Test Scenario

For this study, we employ a proportional-integral-derivative (PID) controller that regulates the cooling capacity of the cycle to 2 kW, and use the extremum seeking algorithm to find optimal values for the expansion valve position and both fan speeds at this operating condition. The PID controller gains are designed and fixed offline and demonstrate good regulatory performance under regular operating conditions.

Note that the extremum-seeking controller (ESC) design is completely decoupled from the PID controller design: in fact, we do not need to know the PID controller gains. The only aspect of the closed-loop system that we are interested in is the time it needs to obtain a good estimate of the steady-state power consumption. This is strongly correlated with the PID controller gains and the closed-loop system characteristics. Knowledge of the time required to reach steady-state (or within a small neighborhood of the steady-state), also referred to as a settling time, informs how often we can update the set-points using the extremum-seeking controller, whose objective is to drive the closed-loop system to the minimal steady-state power consumption. This settling time forms a lower bound on the time period between set-point changes; that is, the set-points should not be altered more frequently than the settling time of the closed-loop system. There is, however, no hard upper bound: the set-point can be updated as slowly as the user desires. The benefit of longer time periods between set-point updates is that one gets a better estimate of the steady-state power consumption, but a disadvantage is that it lowers the convergence rate of ESC algorithms. In these simulations, we update set-points every 5 minutes, which is both practically feasible, and gives us good estimates of the steady-state power consumption.

Figure 5 illustrates the performance of Time-Varying Extremum Seeking Control (TV-ESC) (green), Adam-accelerated TV-ESC (orange), and our proposed algorithm (blue). It is immediately clear from the power plot that the ESC algorithms, even with Adam acceleration, exhibit slower convergence² compared with our Bayesian-Optimization warm-started TV-ESC (hybrid) approach. Indeed, the proposed approach converges to a steady-state cycle power of 0.331 kW within 3 hr, whereas the other algorithms require more than 10 hours to approach the minimum cycle power. Since our goal was to minimize power consumption, we observe that the hybrid approach yields the best results. This is further correlated by Monte-Carlo

²Note that this convergence is the speed at which the closed-loop system performance is optimized by the ESC variants, not the convergence rate at which the closed-loop system is regulated, which is of the order of 5 minutes.

sampling of the admissible set-point space exhaustively with 10^4 samples of combinations of expansion valve positions and fan speeds: the empirical minimum is 0.331 kW and the empirical minimizer is within 5% of the minimizer obtained by our hybrid method. Compared to a baseline simulation obtained by fixing the set-points at their initial values which results in a steady-state power of 0.423 kW, the energy savings of BO-ESC, Adam-ESC, and ESC are 21.75%, 19.62%, and 17.26%, respectively.

Note the TV-ESC algorithm has very slow evolution of the Outdoor Fan Speed (OFS), which the Adam-ESC shows faster convergence. The major difference power consumption is caused by the difference in the optimal EEV positions found by the three algorithms. Since the TV-ESC and Adam-ESC algorithms are dependent on the initial EEV position, they converge to a steady-state EEV position of around 300 counts, whereas the BO algorithm, due to a wider exploration of the EEV space, determines that a steady-state value of 260 counts yields significantly lower power consumption (see the inset plot for comparison around 8-12 hours). The reason for this is that the BO algorithm has an exploration phase, at the end of which, the BO warm-starter yields an excellent local set of tuning parameters. However, there is still local gradient information that can be exploited on-line, which leads to further lowering of power consumption after the BO warm-start. The subplots in the bottom row show that the heat capacity set point is attained throughout the procedure, but the compressor frequency, which is strongly coupled with EEV position, gradually decreases to around 17 Hz, since the compressor is the primary power consumer.




Figure 5: Performance comparison of ESC algorithms discussed in this paper, including TV-ESC (green), Adam-Accelerated TV-ESC (orange), and the proposed BO-ESC (blue). Power consumption, the three inputs (EEV position and indoor/outdoor fan speeds), heat capacity (Q_c), and the compressor frequency (CF) are plotted.

Note that the hybrid BO-ESC algorithm exhibits fluctuations in the exploration phase, whereas the other two algorithms are always smooth. We reiterate this is be-




Figure 6: Zoomed-in version of Bayesian exploration phase in BO-ESC from $t = 30$ min to $t = 60$ min.

cause the BO-ESC explores the entire space of set-points rapidly and therefore set-point changes are typically larger than that exhibited by the other two methods. However, in the zoomed-in Figure 6, the hybrid trajectories in blue are clearly smooth, because these set-points are changed every 5 minutes, which is larger than the settling-time of the closed-loop system, and so the closed-loop system smoothly settles to these new operating conditions within those 5 minutes; that is, the trajectories are actually not non-smooth. This is also reflected in the compressor frequency plot: while it would be of practical concern if the compressor was made to fluctuate non-smoothly during the exploration, this is not the case, as the compressor frequency also settles smoothly to its steady-state within each ESC iteration.

Implementing BO requires the selection of an appropriate acquisition function to balance the exploration-exploitation trade-off. It is not immediately clear which acquisition function is the best choice. To determine this, we test the BO-ESC algorithm with three most commonly used acquisition functions: Expected Improvement (EI) as described in (7), Probability of Improvement (PI), given by $\mathcal{A}_{PI}(\theta) = \Phi(\gamma(\theta))$ and Lower Confidence Bound (LCB), given by $\mathcal{A}_{UCB}(\theta) = \mu(\theta) - 1.96\sigma(\theta)$, where Φ , γ , μ , and σ have been defined previously in (6)–(8). As we can see in Figure 7, all three acquisition functions result in faster convergence than generic TV-ESC and Adam-accelerated TV-ESC (see Figure 5). The EI acquisition function provides slightly lower power consumption after 10 hr. This is not surprising, since it is well-known (empirically) that the expected improvement criterion is better-behaved than PI, and unlike LCB, does not have its own tunable parameter, the selection of which greatly changes the quality of solutions [25]. This motivates the use of EI in the next section, where we investigate the robustness of our proposed method to disturbances.




Figure 7: Performance of three different acquisition functions, with Expected Improvement (blue), Probability of Improvement (orange), and Lower Confidence Bound (green), with the proposed Bayesian-Optimized TV-ESC.

4.2. Robustness Analysis

Next we study the robustness of our proposed algorithm several factors that are important to consider in any practical realization, including rate of actuation, time-varying ambient temperature, closed-loop capacity set-point, variations in controller gain/step-size, actuator quantization, multi-rate or aperiodic actuation, and time-varying loads.

4.2.1. Robustness to Sample Rate

This experiment demonstrates the robustness of our proposed algorithm to the discrete-time sample rate at which sensors are measured and actuators are changed. While higher sample rates might result in faster convergence, this is not always the case since the system is not in steady-state. We conducted simulations with three different sample rates: 60 s, 120 s, and 200 s. We observe from Figure 8 that the TV-ESC algorithm converges more slowly for each of the chosen sample rates, whereas our proposed hybrid algorithm converges within 5 hr, 2.5 hr, and 1.5 hr with data samples available every 200 s, 120 s, and 60 s, respectively. Note the BO-ESC remains in exploration mode for longer with the slower sample rate, which is not surprising. Interestingly, with a 60 s sample rate, the BO-ESC performance actually deteriorates. This is most likely because the system has fast dynamics in the sub-minute time scale, and these are excited, so the system is experiencing more transient behavior than for the longer sample rates. This is most evident from the EEV position, which is the most sensitive actuator.

For subsequent simulations in this section, we fix the sample rate at 120 s.

4.2.2. Robustness to Ambient Temperature

We tested the performance of the proposed algorithm for three different ambient temperatures in the range 30 to

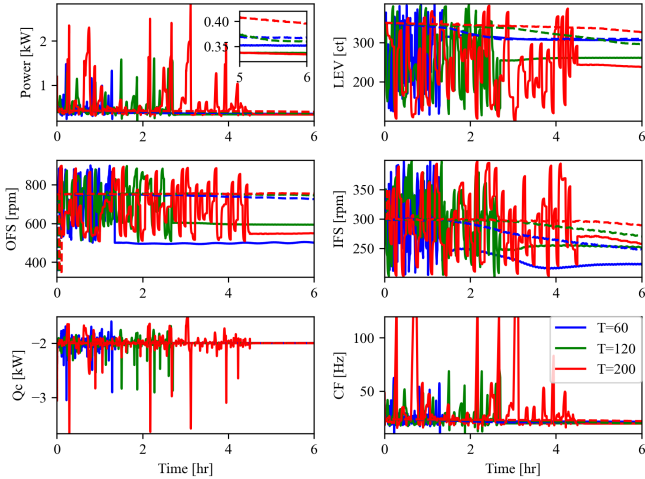


Figure 8: Robustness of proposed BO-ESC to varying sample rate of 60 s (blue), 120 s (green), and 200 s (red). Dashed lines are conventional TV-ESC, while continuous lines are the proposed BO-ESC.

40°C at fixed closed-loop capacity, shown in Figure 9. For each simulation, the ambient temperature was held constant. As expected, the steady-state power increases with increasing ambient temperature, as seen from the power subplot. The CF and capacity feedback loops are regulated to the same values for each case. It is interesting to note that the most sensitive actuator is the EEV, whereas the least sensitive actuator is the Outdoor Fan Speed (OFS). In all cases, the TV-ESC algorithm converges very slowly without BO warm-start, as evident by the dashed lines.

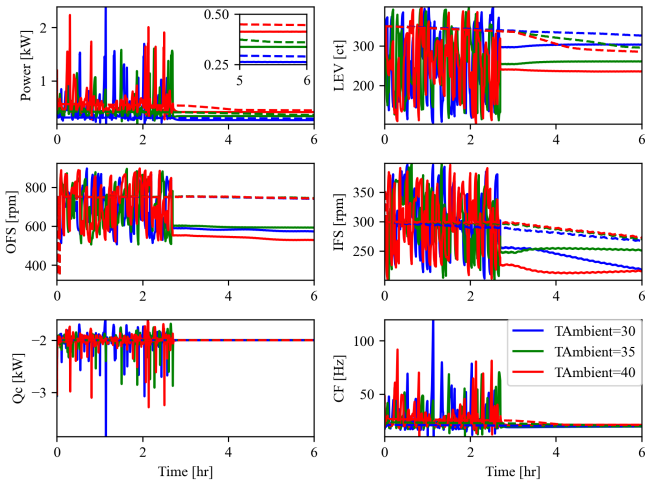


Figure 9: Robustness of proposed BO-ESC to different ambient temperatures (legend entities are in °C).

4.2.3. Robustness to Capacity Set-Point

A major advantage of our proposed approach is that it is agnostic to the control inputs regulating HVAC sub-components. To test this, we change the set-point to the capacity-regulating PI controller from 2 kW in increments

of 0.5 kW. The plots in Figure 10 demonstrate that despite changes in control inputs associated with different capacity set points, our proposed algorithm converges within 2.5 hr to the respective minima. As expected, the steady-state cycle power at lower capacity is lower, and the two most sensitive actuators exhibit the same trend: Higher EEV counts and indoor unit fan speeds for higher capacities. Interestingly, the compressor speed CF converges to the different values for the different capacities.




Figure 10: Robustness of proposed BO-ESC to different capacity set-points (legend entities are in Watt).

4.2.4. Robustness to Sensor and Actuator Quantization

Another cogent problem in implementing extremum-seeking algorithms is the fact that actuators and sensors are quantized. For example, expansion valves actuate in terms of counts, which are integers. We test our proposed algorithm to include quantization on the sensors and actuators. The results of our simulations are shown in Figure 11. We observe that there is no major change in the two sets of power trajectories, although with quantization there are small oscillations in the EEV position. The steady-state power is lower for quantized BO-ESC compared with quantized or continuous TV-ESC.

4.2.5. Robustness to Aperiodic Tuning

Most papers in the literature consider optimizing set-points in a periodic/synchronous manner, that is, the set-points are all updated at the same time, an assumption that may not always be true in engineered systems. To understand how well our proposed algorithm works on systems where the actuators are asynchronous, we ran simulations where the set of actuators that can be altered at any given time is extracted randomly using a Bernoulli random processes. An example of such an asynchronous tuning schedule is shown in Figure 12, where ones indicate at which iterations of the proposed algorithm the set-point could be altered, and zeros when they remained fixed. This




Figure 11: Robustness of proposed BO-ESC to quantization.

mode of implementation is particularly useful when there are combinations of inputs (usually equipment-specific) that should be avoided as per design-level constraints. In such cases, we can incorporate logic for switching off certain components of the inputs, and as seen in this subsection, will not significantly affect performance. That is,




Figure 12: Exemplar asynchronous tuning profile generated from a random process.

at every time, we generate a random vector with three elements, where each element is one or zero depending on which set-point component is alterable. The results of aperiodic/asynchronous tuning is illustrated in Figure 13. In particular, we note the effectiveness of Adam as an accelerating mechanism in this scenario. Since Adam is based on stochastic gradients and is designed to adapt step-sizes in accordance with stochastic properties of the gradient history (as discussed in Section 3.2), the Adam-ESC algorithm outperforms the TV-ESC algorithm even when starting from the same initial set-point. This results in

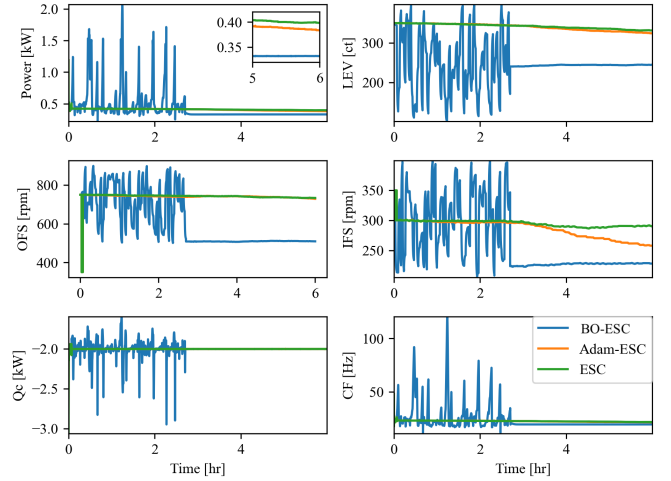


Figure 13: Robustness of proposed BO-ESC to aperiodic tuning.

slight improvement of the efficiency, which is further improved by the BO warm-start. Thus, in scenarios where all the components are not simultaneously tunable, the Adam mechanism is helpful to immunize against component-wise stochasticity.

4.2.6. Robustness to Varying Heat Loads and Environmental Conditions

As a more practical example of the application of this algorithm to vapor compression cycle-based equipment, we developed a more complex model in which the cycle regulated the room temperature in an occupied space with continuous oscillations in both the heat load and the ambient temperature. Standard extremum-seeking algorithms may not function well in the face of large and continual deviations in the operating conditions because the disturbances violate assumed time-scale separation.

This system model used the same cycle model as described in the previous section, but incorporated two feedback loops. The first regulates the room temperature to a set-point via the compressor speed, and the second regulates the evaporator superheat temperature to a set-point of 5°C via the control of the EEV. The regulation of the evaporator superheat temperature is important to ensure the safe long-term operation of the equipment. This closed-loop system was connected to a room model with a volume of 27 m^3 and a 20 cm thick envelope with the thermal properties of wood and external and internal and outer heat transfer coefficients of $1.5\text{ W/m}^2\text{K}$ and $3.5\text{ W/m}^2\text{K}$, respectively. This system was operated according to a standard commercial office building that is occupied during the day and unoccupied at night. As a result, the sensible and latent heat loads in the space were set to 3200 W and 800 W between the hours of 8:00am and 6:00pm, and were set to 2000 W and 100 W otherwise. The ambient temperature was assumed to vary sinusoidally between 24°C and 40°C over a 24 hour period, with a peak

temperature at 2:00pm. We refer the reader to Fig. 14 for an illustration of the sensible and latent heat variations, as well as the ambient temperature and humidity oscillations over a period of 7 days. This study used the TMY3 weather dataset for Atlanta³ corresponding to the Hartsfield-Jackson airport, which is located in the southeastern US climate zone 3A, for the week of August 13-21. One week was determined to be sufficient since the slowest time constant in the model is approximately one day.




Figure 14: Ambient temperature, relative humidity, and internal heat loads (sensitive + latent).

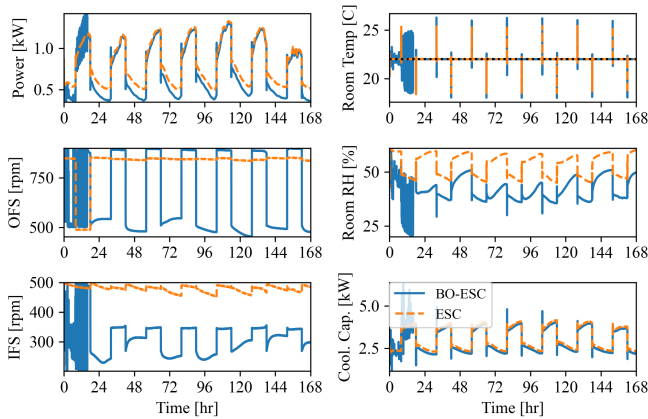


Figure 15: Robustness to varying ambient conditions and head loads. Power, set point, room conditions and room temperature set point (RTSP), and cooling capacity variation.

We compare our proposed BO-ESC method with the TV-ESC algorithm for optimizing two inputs: the indoor and outdoor unit fan speeds (IFS and OFS, respectively). Both set points are assumed to be within a known bounded range: $IFS \in [200, 500]$ rpm, and $OFS \in [500, 1000]$ rpm.

Neither BO-ESC nor TV-ESC are provided with information about the ambient condition variations or heat load profiles *a priori*, although both algorithms assume that the loads can be measured at the current time instant. Numerical estimates of the loads would not be necessary in practice, and an estimate of building occupancy from sensors could provide equivalent information to the algorithms. The standard TV-ESC approach relies on frequent covariance resets because sharp changes in power measurements results in poor gradient estimates that limit ESC performance. For BO-ESC, we use the first 18 hr to collect data for each of the load-modes, i.e., when the heat load is high, and when it is low. For each mode, the BO algorithm learns a surrogate model of the energy function and a corresponding IFS and OFS pair that minimizes the energy based on this surrogate model. After 18 hr, the Adam-ESC algorithm takes over and gradient-based updates for both fan speeds are used.

A comparison of the performance of the two methods are provided in Fig. 15. We note that the ESC algorithm exhibits limited flexibility: both the OFS and IFS set points drift gradually without significantly decreasing the energy consumption over the 7 days. In fact, because the energy function is quickly varying with head load and ambient variations, the ESC gradients are not always accurate, and incremental changes in the fan speeds do not always result in decreasing energy gradients. This is a limitation of ESC: it does not perform well in scenarios where multiple conditions are concurrently changing since it chases ‘incorrect’ local gradients. Conversely, the BO-ESC effectively switches between surrogate models at different load conditions to quickly reset initial guesses for Adam-ESC. This is very clear from the IFS and OFS plots - the sudden changes correspond closely to the heat load changes in Fig. 14. This results in marked improvements of energy efficiency, as one can infer from the power (top-left) subplot in Fig. 15. Also note that our proposed method leads to significant reduction of power consumption both at nighttime and in the morning for most days. Since BO-ESC depends on warm-starting within a given range of IFS and OFS values, it is safer to use for long time-horizons, since it does not exhibit ESC’s drifting behavior: any time the set point starts to drift, a load change will bring it back to the safe range. This is another beneficial feature of the BO-ESC.

Additional insights may be gained by considering the room temperatures, cooling capacities, and relative humidities under both experiments (see Fig. 15). While some deviation between the room temperature and its setpoint is visible during load changes, this is expected because limitations in the vapor-compression cycle model due to the complexity of the accompanying fluid transients in these conditions prevent it from exhibiting on/off behavior. As a result, the cooling capacity when the compressor is at its minimum frequency is still greater than that which is required for the space at night, which causes the room temperature to fall below the setpoint. Given these limi-

³<https://www.nrel.gov/docs/fy08osti/43156.pdf>

tations, however, the room temperature control is mostly excellent for both BO-ESC and ESC, which suggests that the BO-ESC method will not adversely affect the occupant’s thermal comfort. The cooling capacities of the system under both types of control are also quite similar, which correlates well with the consistent room temperature behavior. It is perhaps most notable that the room’s relative humidity under BO-ESC is much lower than under ESC. This represents a beneficial side-effect of minimizing the fan power consumption; because the fan speeds are reduced to achieve lower electrical power consumption at night, the longer residence time of the air in the evaporating heat exchanger causes the air to be dehumidified more effectively and reduce the relative humidity in the room. In this case, the BO-ESC method is able to achieve two benefits for the price of one: by lowering the fan speed, it can both reduce the energy consumption of the system and improve occupant comfort.




Figure 16: Comparison of total energy in kW-hr and energy savings using TV-ESC and our proposed BO-ESC approach.

Since ESC is not equipped with the information that there are high load and low load periods of the day, it does considerably worse in terms of energy savings than BO-ESC as we see in Fig. 15. To enable a more fair comparison, we conduct an additional set of simulations in which we supply the times of the heat load mode changes to ESC, and refer to this algorithm as Switched-ESC. Subsequently, we compare the total energy expenditure of the building self-optimized by ESC, switched-ESC, and BO-ESC against no self-optimization, i.e., a baseline where the indoor unit and outdoor unit fan speeds were kept constant at 450 rpm and 850 rpm, respectively. The energy expenditure in kW-hr is plotted in the left subplot of Figure 16 and the right subplot of Figure 16 illustrates the energy savings of each competing algorithm over the baseline. It is clear that the Switched ESC shows somewhat improved performance with the newly supplied heat load information. Though there is improvement in ESC performance, it is nevertheless clear from the figure that the BO-ESC outperforms both variants of the ESC. This leads us to reaffirm that in time-varying environments like the one

simulated in this experiment, ESC algorithms need non-trivial modifications such as BO warm-starting to yield competitive performance.

5. Conclusions

In this paper, we proposed a self-optimization control algorithm for set-point optimization of HVAC vapor compression systems. The method uses data to construct a probabilistic surrogate model that maps the set-points to the power consumed by the cycle. This probabilistic map subsequently informs set-point selection in a Bayesian Optimization (BO) framework. With few data points, the BO method computes an initial set of set-points that accelerate the convergence of extremum seeking methods, that have previously exhibited excellent self-tuning performance in similar problems. A combination of BO warm-starting and Adam acceleration (a well-known optimization method in the machine learning literature) results in an extremum seeking solution that can handle highly nonlinear objective functions, is robust to measurement noise and variation in environmental conditions, is agnostic to model knowledge and control algorithms present in the vapor compression cycle, and in our experience consistently converges faster than state-of-the-art extremum seeking methods.

While our method is generalizable to any closed-loop system whose settling time is measurable, one open question is how the BO exploration can take place when there are operational constraints. Although this is not within the scope of this paper, we intend to explore constrained versions of ESC and BO that can be synergized for safe operation applicable to a wide class of thermal engineering and energy systems in future work.

References

- [1] Álvaro Roberto Gardenghi, J. F. Lacerda, C. B. Tibiriçá, L. Cabezas-Gómez, Numerical and experimental study of the transient behavior of a domestic vapor compression refrigeration system - influence of refrigerant charge and ambient temperature, *Applied Thermal Engineering* (2021) 116728.
- [2] Y. Zhang, J. Wang, F. Hu, Y. Wang, Comparison of evaluation standards for green building in China, Britain, United States (2017).
- [3] L. Dong, Y. Li, J. M. House, T. I. Salsbury, Model-free control and staging for real-time energy efficient operation of a variable refrigerant flow system with multiple outdoor units, *Applied Thermal Engineering* 180 (2020) 115787.
- [4] H. Huang, W. I. Binti Wan Mohd Nazi, Y. Yu, Y. Wang, Energy performance of a high-rise residential building retrofitted to passive building standard – a case study, *Applied Thermal Engineering* 181 (2020) 115902.
- [5] Z. Zhao, Y. Li, T. I. Salsbury, J. M. House, C. F. Alcalá, Local self-optimizing control based on extremum seeking control, *Control Eng. Pract.* 99 (August 2019) (2020) 104394.
- [6] J. Song, W. Tian, X. Xu, Y. Wang, Z. Li, Thermal performance of a novel ultrasonic evaporator based on machine learning algorithms, *Applied Thermal Engineering* 148 (2019) 438–446.
- [7] Y. Wang, W. Li, Z. Zhang, J. Shi, J. Chen, Performance evaluation and prediction for electric vehicle heat pump using machine learning method, *Applied Thermal Engineering* 159 (2019) 113901.

- [8] D. E. Marasco, C. E. Kontokosta, Applications of machine learning methods to identifying and predicting building retrofit opportunities, *Energy Build.* (2016).
- [9] K. Amasyali, N. M. El-Gohary, A review of data-driven building energy consumption prediction studies, *Renewable and Sustainable Energy Reviews* 81 (2018) 1192–1205.
- [10] J. Drgoňa, D. Picard, M. Kvasnica, L. Helsen, Approximate model predictive building control via machine learning, *Applied Energy* 218 (2018) 199 – 216.
- [11] D. Picard, J. Drgoňa, M. Kvasnica, L. Helsen, Impact of the controller model complexity on model predictive control performance for buildings, *Energy and Buildings* 152 (2017) 739 – 751.
- [12] K. B. Ariyur, M. Krstic, *Real-time optimization by extremum-seeking control*, John Wiley & Sons, 2003.
- [13] M. Benosman, Learning-based adaptive control: An extremum seeking approach—theory and applications, Butterworth-Heinemann, 2016.
- [14] M. Guay, D. J. Burns, A comparison of extremum seeking algorithms applied to vapor compression system optimization, in: 2014 American Control Conference, IEEE, 2014, pp. 1076–1081.
- [15] D. J. Burns, W. K. Weiss, M. Guay, Realtime set-point optimization with time-varying extremum seeking for vapor compression systems, in: 2015 American Control Conference (ACC), IEEE, 2015, pp. 974–979.
- [16] M. Guay, D. J. Burns, A proportional integral extremum-seeking control approach for discrete-time nonlinear systems, *International Journal of Control* 90 (8) (2017) 1543–1554.
- [17] S. Sharma, A. G. Alleyne, Extremum seeking control of battery powered vapor compression systems for commercial vehicles, *International Journal of Refrigeration* (2020).
- [18] P. Li, Y. Li, J. E. Seem, Efficient operation of air-side economizer using extremum seeking control, *J. Dyn. Syst. Meas. Control. Trans. ASME* (2010).
- [19] B. Mu, Y. Li, J. M. House, T. I. Salsbury, Real-time optimization of a chilled water plant with parallel chillers based on extremum seeking control, *Appl. Energy* (2017).
- [20] L. Dong, Y. Li, J. M. House, T. I. Salsbury, Model-free control and staging for real-time energy efficient operation of a variable refrigerant flow system with multiple outdoor units, *Appl. Therm. Eng.* 180 (2020) 115787.
- [21] W. Wang, Y. Li, B. Hu, Real-time efficiency optimization of a cascade heat pump system via multivariable extremum seeking, *Appl. Therm. Eng.* 176 (May) (2020) 115399.
- [22] P. Westermann, R. Evins, Surrogate modelling for sustainable building design – A review, *Energy Build.* 198 (2019) 170–186.
- [23] D. P. Kingma, J. L. Ba, Adam: A method for stochastic optimization, 3rd Int. Conf. Learn. Represent. (2015) 1–15 arXiv:1412.6980.
- [24] I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, The MIT Press, 2016.
- [25] J. Snoek, H. Larochelle, R. P. Adams, Practical bayesian optimization of machine learning algorithms, in: *Advances in Neural Information Processing Systems*, 2012, pp. 2951–2959.
- [26] Y. Wang, et al., Prediction of tubular solar still performance by machine learning integrated with bayesian optimization algorithm, *Applied Thermal Engineering* 184 (2021) 116233.
- [27] S. A. Bortoff, D. J. Burns, C. R. Laughman, H. Qiao, C. Danielson, A. Goldsmith, S. Di Cairano, Power optimizing control of multi-zone heat pumps, in: *IEEE Conference on Control Technology and Applications*, 2018, pp. 826–833.
- [28] Air-Conditioning, Heating & Refrigeration Institute (AHRI), 2111 Wilson Blvd, Suite 500, Standard for Performance Rating of Unitary Air-conditioning and Air-source Heat Pump Equipment (2017).
- [29] M. Guay, A time-varying extremum-seeking control approach for discrete-time systems, *J. Process Control* 24 (3) (2014) 98–112.
- [30] G. C. Goodwin, K. S. Sin, *Adaptive Filtering Prediction And Control*, Courier Corporation, 2014.
- [31] S. J. Reddi, S. Kale, S. Kumar, On the convergence of adam and beyond, arXiv preprint arXiv:1904.09237 (2019).
- [32] R. Ward, X. Wu, L. Bottou, Adagrad stepsizes: sharp convergence over nonconvex landscapes, in: *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9–15 June 2019, Long Beach, California, USA, Vol. 97, 2019*, pp. 6677–6686.
- [33] C. K. Williams, C. E. Rasmussen, *Gaussian Processes For Machine Learning*, Vol. 2, MIT press Cambridge, MA, 2006.
- [34] A. Shah, A. Wilson, Z. Ghahramani, Student-*t* processes as alternatives to gaussian processes, in: *Artificial Intelligence and Statistics*, 2014, pp. 877–885.
- [35] J. T. Springenberg, A. Klein, S. Falkner, F. Hutter, Bayesian optimization with robust Bayesian neural networks, in: *Advances in Neural Information Processing Systems*, 2016, pp. 4134–4142.
- [36] K. Kawaguchi, L. P. Kaelbling, T. Lozano-Pérez, Bayesian optimization with exponential convergence (2015).
- [37] Standard 210/240: Performance rating of unitary air-conditioning & air-source heat pump equipment (2017).
- [38] Modelica Association, *Modelica specification, Version 3.4* (2017). URL www.modelica.org
- [39] H. Qiao, V. Aute, R. Radermacher, Transient modeling of a flash tank vapor injection heat pump system - Part I: Model development, *Int. J. Refrigeration* 49 (2015) 169–182.
- [40] Modelica Association, *Functional Mockup Interface for Model Exchange and Co-Simulation, Version 2.0.1* (2019). URL www.fmi-standard.org
- [41] Dassault Systemes, *Dymola* 2020 (2019).

Appendix

Acronyms

VCS	vapor compression systems	BO	Bayesian optimization
ESC	extremum seeking control	GP	Gaussian processes
EEV/LEV	electronic expansion valve	TV-ESC	time-varying extremum seeking control
IFS	indoor unit fan speed	OFS	outdoor unit fan speed
DAE	differential algebraic equation	FMU/FMI	functional mockup unit/interface
Adam-ESC	Adam-acclerated ESC	BO-ESC	Bayesian optimization warm-started ESC
EI/PI	expected/probability of improvement	UCB	upper confidence bound

Symbols

\mathbb{N}	natural numbers	\mathbb{R}	real numbers
f	VCS dynamical equations	t	time index
x	VCS system state	ν	VCS exogenous disturbance/noise
θ	VCS set points	Θ	set of admissible set points
y	measured power output	J	power cost function from state
\bar{J}	power cost function from set point	∇	gradient
∇^2	Hessian	θ^*	optimal set point
g	estimated gradient	d	dither signal
k_g	step size or ESC control gain	p	number of set points
ΔJ	cost function increment	$\Delta\theta$	set point change
N_d	size of data history for gradient filter	α	forgetting factor for gradient estimator/filter
e	gradient filter error	K	gradient filter gain
P	gradient filter covariance	m_t	Adam first moment
v_t	Adam second moment	ϵ	Adam user-defined parameter
\odot	element-wise product operator	N	number of past measurements in Adam
k_t	Adam control gain/step size	β	Adam filter coefficients
\mathcal{K}	kernel function	l	GP length-scale
σ_0	GP variance	σ_n	GP additive noise
\mathcal{N}	normal density function	μ	GP predicted mean
σ^2	GP predicted variance	\mathcal{L}	GP training loss
K_D	GP kernel vector	\mathcal{K}_D	GP kernel matrix
\mathcal{A}	GP acquisition function	ϕ	zero-mean unit-variance normal density
Φ	cdf of ϕ	γ	expected improvement component
N_0	initial number of burn-in iterations	N_θ	number of BO iterations for warm-starting

Python 3.7/3.8 Toolboxes with Version

FMPy	0.2.27	GPFLOW	2.2.1
Tensorflow	2.3.1	Numpy	1.19.2
Scipy	1.5.3	Pandas	1.1.3