Realtime Optimization of MPC Setpoints using Time-Varying Extremum Seeking Control for Vapor Compression Machines

Weiss, W.; Burns, D.J.; Guay, M.

TR2014-072 July 2014

Abstract

Recently, model predictive control (MPC) has received increased attention in the HVAC community, largely due to its ability to systematically manage constraints while optimally regulating signals of interest to setpoints. For example, in a common formulation of an MPC control problem for variable compressor speed vapor compression machines, the setpoints often include the zone temperature and the evaporator superheat temperature. However, the energy consumption of vapor compression systems has been shown to be sensitive to these setpoints. Further, while superheat temperature is often preferred because it can be easily correlated to heat exchanger efficiency (and therefore cycle efficiency), direct measurement of superheat is not always available. Therefore, identifying alternate signals in the control of vapor compression machines that correlate to efficiency is desired.

In this paper, we consider a model-free extremum seeking algorithm that adjusts setpoints provided to a model predictive controller. While perturbation-based extremum seeking methods have been known for some time, they suffer from slow convergence rates–a problem emphasized in application by the long time constants associated with thermal systems. Our method uses a new algorithm (time-varying extremum seeking), which has dramatically faster and more reliable convergence properties. In particular, we regulate the compressor discharge temperature using an MPC controller with setpoints selected from a model-free time-varying extremum seeking algorithm. We show that the relationship between compressor discharge temperature and power consumption is convex (a requirement for this class of realtime optimization), and use time-varying extremum seeking to drive these setpoints to values that minimize power. The results are compared to the traditional perturbation-based extremum seeking approach. Further, because the required cooling capacity (and therefore compressor speed) is a function of measured and unmeasured disturbances, the optimal compressor discharge temperature setpoint must vary according to these conditions. We show that the energy optimal discharge temperature is tracked with the time-varying extremum seeking algorithm in the presence of disturbances.

International Refrigeration and Air Conditioning Conference at Purdue, 2014

Realtime Optimization of MPC Setpoints using Time-Varying Extremum Seeking Control for Vapor Compression Machines

Walter K. WEISS[1] Daniel J. BURNS[2*] Martin GUAY[1]

[1]Queen's University
Department of Chemical Engineering
Kingston, Ontario
walter.weiss@queensu.ca, martin.guay@chee.queensu.ca

[2]Mitsubishi Electric Research Laboratories
Cambridge, MA, USA
burns@merl.com
*corresponding author

ABSTRACT

Recently, model predictive control (MPC) has received increased attention in the HVAC community, largely due to its ability to systematically manage constraints while optimally regulating signals of interest to setpoints. For example, in a common formulation of an MPC control problem for variable compressor speed vapor compression machines, the setpoints often include the zone temperature and the evaporator superheat temperature. However, the energy consumption of vapor compression systems has been shown to be sensitive to these setpoints. Further, while superheat temperature is often preferred because it can be easily correlated to heat exchanger efficiency (and therefore cycle efficiency), direct measurement of superheat is not always available. Therefore, identifying alternate signals in the control of vapor compression machines that correlate to efficiency is desired.

In this paper, we consider a model-free extremum seeking algorithm that adjusts setpoints provided to a model predictive controller. While perturbation-based extremum seeking methods have been known for some time, they suffer from slow convergence rates—a problem emphasized in application by the long time constants associated with thermal systems. Our method uses a new algorithm (time-varying extremum seeking), which has dramatically faster and more reliable convergence properties. In particular, we regulate the compressor discharge temperature using an MPC controller with setpoints selected from a model-free time-varying extremum seeking algorithm. We show that the relationship between compressor discharge temperature and power consumption is convex (a requirement for this class of realtime optimization), and use time-varying extremum seeking to drive these setpoints to values that minimize power. The results are compared to the traditional perturbation-based extremum seeking approach. Further, because the required cooling capacity (and therefore compressor speed) is a function of measured and unmeasured disturbances, the optimal compressor discharge temperature setpoint must vary according to these conditions. We show that the energy optimal discharge temperature is tracked with the time-varying extremum seeking algorithm in the presence of disturbances.

1 INTRODUCTION

Vapor compression systems (VCS), such as heat pumps, refrigeration and air-conditioning systems, are widely used in industrial and residential applications. The introduction of variable speed compressors, electronically-positioned valves, and variable speed fans to the vapor compression cycle has greatly improved the flexibility of the operation of such systems (Fig. 1A). This increased actuator flexibility, along with increasing onboard computational power, enables more sophisticated control schemes than traditional on-off or decentralized PI control. For example, Model Predictive Control (MPC) of vapor compression systems (Fig. 1B) offers a flexible and rigorous design process in which the constraints are enforced during transients and can be modified as the design evolves; and the resulting controller can be computed and analyzed immediately, providing rigorous guarantees on feasibility, optimality, convergence, transient performance and stability. Further, process control variables such as zone temperature and superheat temperature can be regulated to their setpoints in steady state. Previous work has shown that the energy efficiency of these systems is strongly dependent on these setpoints (Burns and Laughman, 2012), however, determining appropriate setpoints is not always straightforward.

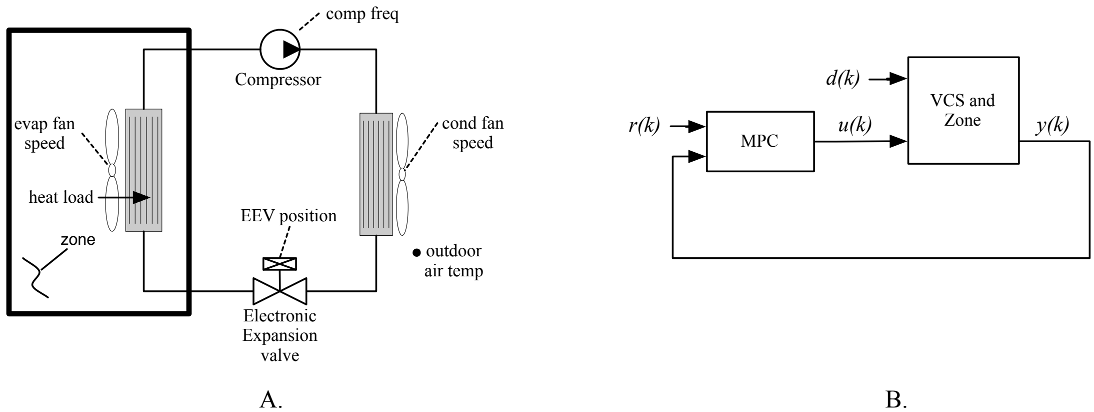

Figure 1: A. The vapor compression system under study consists of a variable speed compressor, condensing heat exchanger, electronically controlled expansion valve, and evaporating heating exchanger. The inputs to the VCS that are manipulated by the control system include (i) the compressor frequency, (ii) the condenser fan speed, (iii) the EEV position, and (iv) the evaporator fan speed. B. An MPC controller is nominally configured to use measurements $y(k)$ to drive regulated variables of a vapor compression system and zone to their setpoints $r(k)$ in the presence of disturbances $d(k)$ such as changes in outdoor air temperature and heat load.

In a typical air conditioning system, setpoints to the controller[1] may include (1) zone temperatures selected by the user and (2) internal machine signals, the regulation of which is required for delivering the required cooling capacity in the presence of given thermodynamic boundary conditions such as heat load and outdoor air temperature. Assuming there exist flexibility with actual zone temperatures, the optimization of setpoints of type (1) have been extensively investigated, especially in the context of a model predictive controller where disturbances $d(k)$ such as ambient temperature and occupancy may be predicted known for some horizon into the future. The interested reader may refer to (Ma *et al.*, 2012, Oldewurtel *et al.*, 2012, Zhang *et al.*, 2013) for more information on problems of this type. However, in this paper, we consider the optimization of setpoints of type (2); that is, internal machine process variables whose steady state values determine the energy consumption of the vapor compression machine.

Often, setpoints of type (2) are simply given as a constant evaporator superheat temperature. In this case, it is assumed that the superheat temperature is a good surrogate for overall cycle efficiency, and by regulating the cycle such that all the refrigerant passing through the evaporator becomes saturated vapor upon exiting, it is assumed that the overall process is performed efficiently. Additionally, it is often stated that because compressors may be damaged when ingesting two-phase refrigerant, regulating the superheat temperature to a positive value ensures only saturated vapor is ingested. However, strict measurement of superheat requires at least one temperature and one pressure measurement (and perhaps more sensors are required depending on the assumptions one makes regarding pressure losses in the evaporator), and these sensors are often too expensive to be included in commercial systems. Additionally, for systems with multiple evaporators, requiring independent regulation of both superheat temperature and zone temperature may not even be possible with the typical set of actuators, because the number of regulated variables may exceed the number of controls. Therefore, alternatives to superheat setpoints for regulating cycle capacity and efficiency are desired.

In this paper, we select the compressor discharge temperature as a signal to be regulated by an MPC controller. The discharge temperature is often measured for equipment protection making it a commonly available signal, and because

[1] In this paper, we consider the 'controller' as the low-level algorithm that determines actuator commands and not, as in other applications, as a user interface to the vapor compression machine for programming zone setpoints. However, we assume that zone setpoints entered by the user in the case of the latter are available to the former.

the refrigerant state at this location in the cycle is always superheated, this signal is a one-to-one function of the disturbances over the full range of expected operating points. (Contrast this with evaporator superheat temperature, which is not defined for values less than zero and produces no change in sensible temperature measurement when two-phase refrigerant exits the evaporator. One of the main challenges of superheat regulation is that low superheat temperature, which is good for efficiency, is easily perturbed to zero in the presence of disturbances, causing the loss of signal information and therefore of feedback control.) Because discharge temperature changes with heat loads and outdoor air temperatures, its setpoint cannot be regulated to a constant, but instead must vary with these conditions. It is the aim of this paper to automate the generation of such setpoints in order to maximize energy efficiency.

One common way to optimize the performance of a vapor compression system is to use a mathematical model of the governing physics. However, models that attempt to describe the influence of steady state operating points on thermodynamic behavior and power consumption are often low in fidelity, and while they may have useful predictive capabilities over the conditions in which they were calibrated, the environments into which these systems are deployed are so diverse as to render comprehensive calibration and model tuning intractable. Therefore, relying on model-based strategies for realtime (online) selection of optimal setpoints is tenuous.

Recently, model-free methods that operate in realtime and aim to optimize a cost have received increased attention and have demonstrated improvements in the optimization of vapor compression systems and other HVAC applications (Burns and Laughman, 2012, Li *et al.*, 2010, Sane *et al.*, 2006, Tyagi *et al.*, 2006) To date, the dominant extremum seeking algorithm that appears in the HVAC research literature is the traditional perturbation-based algorithm first developed in the 1920s (Leblanc, 1922) and re-popularized in the late 1990s by an elegant proof of convergence for a general class of nonlinear systems (Krstic, 2000).

While all extremum seeking techniques optimize a performance metric by estimating its gradient and driving inputs such that the metric is optimized, the way in which the gradient is estimated has a strong influence on its convergence properties. In the traditional perturbation-based method, a sinusoidal term is added to the input at a slower frequency than the natural plant dynamics, inducing a sinusoidal response in the performance metric (Tan *et al.*, 2010). The extremum seeking controller then filters and averages this signal to obtain an estimate of the gradient. Averaging the perturbation introduces yet another (and slower) time scale in the optimization process. For thermal systems such as vapor compression machines where the dynamics are already on the order of tens of minutes, the slow convergence properties of perturbation-based extremum seeking become impediments to wide-scale deployment.

However, new extremum seeking approaches have been developed that estimate the gradient of the performance metric in a way that does not introduce two time scales. Time-varying extremum seeking uses adaptive filtering techniques to estimate the parameters of the gradient function from measured data, eliminating averaging in the controller (Guay *et al.*, 2013). In this paper, we apply time-varying extremum seeking to the problem of obtaining setpoints that optimize energy efficiency in a vapor compression system.

The rest of the paper is organized as follows: A brief problem description is given in Section 2, followed by an explanation of the time-varying extremum seeking control (TV-ESC) routine in Section 3. In Section 5, simulation results for the application of the TV-ESC on an air conditioning system are presented. A comparison of TV-ESC and perturbation based ESC is shown in Section 4 using a simplified example. Lastly, concluding remarks are offered in section 6.

2 PROBLEM DESCRIPTION

Consider the following nonlinear system to represent the dynamics of the vapor compression system

$$x_{k+1} = x_k + f(x_k, r_k) \tag{1a}$$

$$z_k = h(x_k) \tag{1b}$$

where $x_k \in \mathbb{R}^n$ is the vector of state variables at the k^{th} time step, $r_k \in \mathbb{R}$ is the reference signal for discharge temperature, and $z_k \in \mathbb{R}$ is the measured variable to be minimized, power consumption. The goal is to drive the system

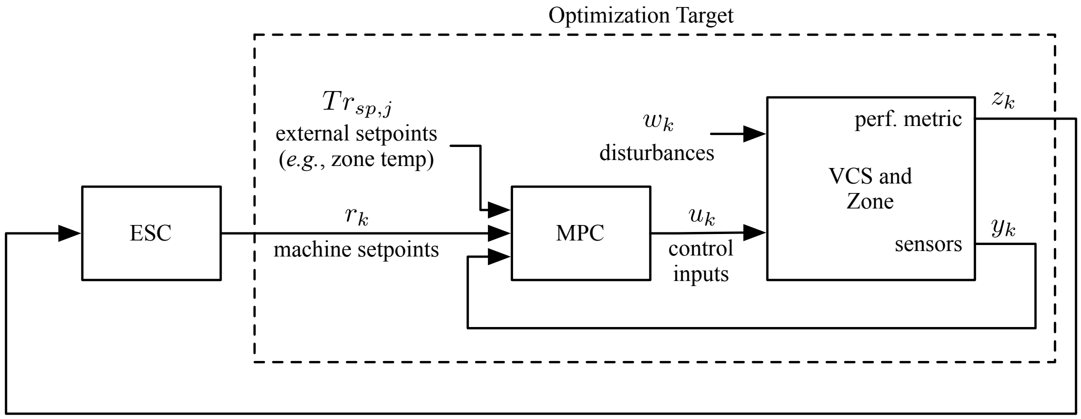

Figure 2: Block diagram of system and controllers.

to the equilibrium x^* and r^* so that z_k is minimized. The optimal input r^* is the setpoint for discharge temperature that minimizes power consumption. The MPC is designed for zone temperature and discharge temperature tracking, as well as to enforce constraints on the system inputs and outputs.

The MPC uses a linear model to approximate the nonlinear dynamics of the VCS. In compact form, this model is given by

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + B_d w_k \\ y_k &= Cx_k + Du_k \end{aligned} \tag{2}$$

where $A \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times p}$, $B_d \in \mathbb{R}^{n\times m}$, $C \in \mathbb{R}^{q\times n}$, $D \in \mathbb{R}^{q\times p}$, $u_k \in \mathbb{R}^p$ is the vector of control inputs to the VCS, $w_k \in \mathbb{R}^m$ is the disturbance vector, and $y_k \in \mathbb{R}^q$ is the vector of system outputs.

Let $T_{r_{sp}} \in \mathbb{R}$ represent the setpoint for zone temperature which is known a priori. Further, $y_{T_d}(k)$ and $y_{T_r}(k) \in y_k$ are the measured discharge temperature and zone temperature respectively. We can then define the performance metric $v(k)$ as follows:

$$v(k) = \begin{pmatrix} y_{T_d}(k) - r(k) \\ y_{T_r}(k) - T_{r_{sp}} \end{pmatrix} \tag{3}$$

The MPC optimization problem is then formulated as

$$\begin{aligned} \min_{u(0),\ldots,u(k)} \quad & x(N)^T Px(N) + \sum_{k=0}^{N-1} v(k)^T Qv(k) + u(k)^T Ru(k) \\ \text{s.t.} \quad & y_{min} \leq y(k) \leq y_{max}, \quad k = 0, \ldots, N \\ & u_{min} \leq u(k) \leq u_{max}, \quad k = 0, \ldots, N \end{aligned} \tag{4}$$

where N is the MPC prediction horizon. The matrix P is obtained from the solution to the algebraic Riccati equation (which is solved to determine the terminal cost for the finite horizon control problem). The matrices Q and R are tuned to appropriately weight the objectives in the cost function.

The MPC optimization problem is solved online at each time step, providing a set of control moves from the current time step to the a time N steps in the future. The first control move is applied to the system, new measurements are obtained, and the optimization problem is recalculated. The goal is to improve the efficiency of the vapor compression system by supplying the MPC with a discharge temperature set-point that will minimize power consumption while tracking a given zone temperature. The proposed realtime optimizer is a time-varying extremum seeking control (TV-ESC) scheme.

3 EXTREMUM SEEKING CONTROLLER

The ESC is tasked with supplying the MPC with the setpoint for discharge temperature r that will minimize power consumption z. We follow the discrete-time ESC update law outlined in Guay (2014).

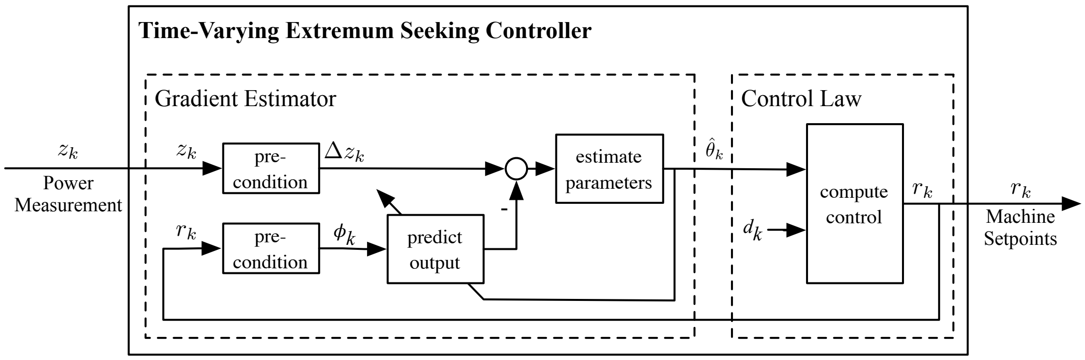

Figure 3: Overview of the TV-ESC algorithm.

The equilibrium cost $z = \ell(r^*)$ satisfies the following optimality conditions:

$$\frac{\partial \ell(r^*)}{\partial r} = 0 \tag{5}$$

$$\frac{\partial^2 \ell(r^*)}{\partial r \partial r^T} > \beta I \quad \forall r \in R \tag{6}$$

where β is a strictly positive constant.

Let $\phi_k = \Delta r_k$. The dynamics of the cost function can be parametrized as:

$$\Delta z_k = \theta_k^T \Delta r_k = \phi_k^T \theta_k \tag{7}$$

Let the estimator for (7) be

$$\Delta \hat{z}_k = \hat{\theta}_k^T \Delta r_k = \phi_k^T \hat{\theta}_k \tag{8}$$

where $\hat{\theta}_k$ is the vector of parameter estimates. The output prediction error is defined as $e_k = \Delta z_k - \Delta \hat{z}_k$.

The dynamical system operates at the faster time-scale with sampling time $\epsilon \Delta t$ while the steady-state optimization operates at the slow time scale with sampling time Δt, where ϵ is a time-scale separation parameter. The parameter estimate update approach is as follows:

$$\Sigma_{k+1}^{-1} = \Sigma_k^{-1} + \epsilon \left(\frac{1}{\alpha} - 1 \right) \Sigma_k^{-1} - \frac{\epsilon}{\alpha^2} \Sigma_k^{-1} \phi_k (1 + \frac{1}{\alpha} \phi_k^T \Sigma_k^{-1} \phi_k)^{-1} \phi_k^T \Sigma_k^{-1} \tag{9}$$

$$\bar{\hat{\theta}}_{k+1} = Proj \left[\hat{\theta}_k + \frac{\epsilon}{\alpha} \Sigma_k^{-1} \phi_k (1 + \frac{1}{\alpha} \phi_k^T \Sigma_k^{-1} \phi_k)^{-1} (e_k), \Theta_0 \right] \ . \tag{10}$$

Where $\Sigma \in \mathbb{R}^{n_\theta \times n_\theta}$ is the covariance matrix and $Proj$ is an orthogonal projection operator. For a more detailed discussion on this operator see Guay (2014) and Goodwin and Sin (2013).

The gradient descent controller is given by:

$$r_{k+1} = r_k - \epsilon k_g \hat{\theta}_k + \epsilon d_k \tag{11}$$

where d_k is a bounded dither signal and k_g is the optimization gain .

Together, the iterative extremum seeking routine is given by:

$$r_{k+1} = r_k - \epsilon k_g \hat{\theta}_k + \epsilon d_k \tag{12a}$$

$$\phi_k = \Delta r_k = r_{k+1} - r_k \tag{12b}$$

$$\Delta \hat{z}_k = \phi_k^T \hat{\theta}_k \tag{12c}$$

$$\Sigma_{k+1}^{-1} = \Sigma_k^{-1} + \epsilon(\frac{1}{\alpha} - 1)\Sigma_k^{-1} - \frac{\epsilon}{\alpha^2}\Sigma_k^{-1}\phi_k(1 + \frac{1}{\alpha}\phi_k^T\Sigma_k^{-1}\phi_k)^{-1}\phi_k^T\Sigma_k^{-1} \tag{12d}$$

$$\bar{\hat{\theta}}_{k+1} = Proj\left[\hat{\theta}_k + \frac{\epsilon}{\alpha}\Sigma_k^{-1}\phi_k(1 + \frac{1}{\alpha}\phi_k^T\Sigma_k^{-1}\phi_k)^{-1}(e_k), \Theta_0\right] \quad . \tag{12e}$$

As shown in Figure 3, at the k^{th} iteration step, the ESC algorithm uses the difference between current r_k and next input r_{k+1}, and the difference between measured Δz_k and predicted $\Delta \hat{z}_k$ change in power consumption for the gradient estimation. The estimated gradient will be used to parameterize the unknown but measured cost function describing power consumption. The gradient is estimated by employing a recursive least squares filter with forgetting factor α. Further, the estimated gradient is used to compute the gradient descent controller which will reduce power consumption. The new setpoint is applied by the MPC, and the ESC algorithm is repeated at the next ESC sample time.

Note that the time-varying extremum seeking controller does not require averaging the effect of the perturbation as in the case of the traditional perturbation-based extremum seeking controller. For this reason, time-varying extremum seeking converges substantially faster, as demonstrated in an example in the following section.

4 COMPARISON OF TIME-VARYING AND PERTURBATION EXTREMUM SEEKING CONTROL

To illustrate the differences in convergence rate between perturbation-based ESC and TV-ESC, these two methods are used to optimize a Hammerstein system consisting of first-order linear difference equation and a static output nonlinearity (see Figure 4A). The equations for this system are given by

$$x_{k+1} = 0.8x_k + r_k \tag{13}$$

$$z_k = (x_k - 3)^2 + 1 \tag{14}$$

which has a single optimum point at

$$r^* = 0.6 \tag{15}$$

$$z^* = 1. \tag{16}$$

Note that the pole location in the difference equation component establishes a dominant timescale and therefore sets a fundamental limit for the convergence rate.

In order to illustrate the difference in convergence rates, a discrete-time perturbation-based extremum seeking controller (perturb-ESC) and a time-varying extremum seeking controller (TV-ESC) is applied to the problem of finding the input r that minimizes the output z, without a model of the process or any explicit knowledge of the nature of the optimum. Reasonable effort is made to obtain algorithm parameters for both ESC methods that achieve the best possible convergence rates. The parameters for the perturbation based ESC used for simulation are

$$d_k = 0.2\sin(0.1k) \tag{17}$$

$$\omega_{LP} = 0.03 \tag{18}$$

$$K = -0.005 \tag{19}$$

Where d_k is the sinusoidal perturbation, ω_{LP} is the cutoff frequency for a first-order low-pass averaging filter, and K is the adaptation gain. Note that the high-pass washout filter was not used as convergence rate was improved without it. For details of a discrete-time perturb-ESC formulation, see Killingsworth and Krstic (2006).

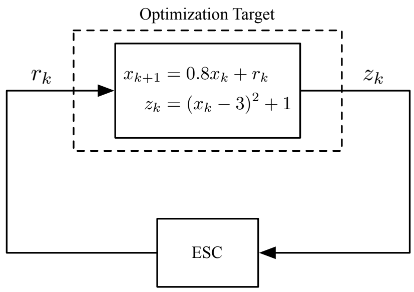

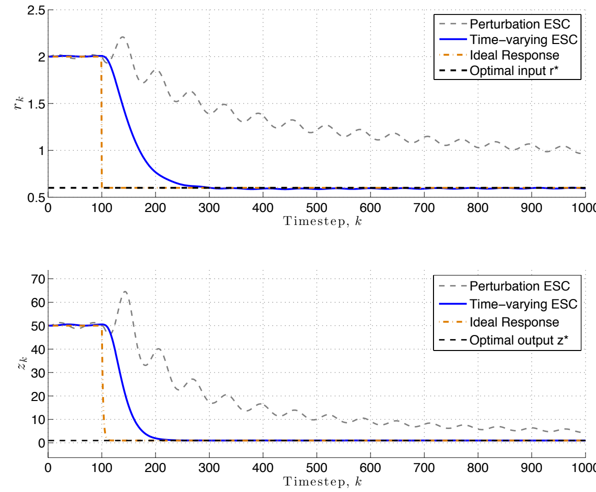

A. B.

Figure 4: Comparing TV-ESC with perturbation ESC. For this application, TV-ESC converges considerably faster to the optimum.

The parameters used for the TV-ESC are

$$d_k = 0.001 \sin(0.1k) \tag{20}$$
$$k_g = 0.001 \tag{21}$$
$$\alpha = 0.1 \tag{22}$$
$$\epsilon = 0.4 \tag{23}$$
$$\tag{24}$$

Where k_g is the adaptation gain, α is the forgetting factor, and ϵ is the timescale separation factor. No projection algorithm was needed for this example.

Simulations are performed starting from an initial input value of $r = 2$ and the ESC methods are turned on after 100 steps. The resulting simulations are shown in Figure 4B. The perturb-ESC method converges to a neighborhood around the optimum in about 4000 steps (not shown in the figure), while the TV-ESC method converges in about 250 steps.

The fast convergence characteristic of TV-ESC is well suited to the optimization of thermal systems with their associated long time constants. In the next section, we apply the TV-ESC algorithm to the problem of selecting setpoints for a MPC controller of a vapor compression machine and present simulation results.

5 SIMULATION RESULTS

The extremum seeking algorithm is used to determine the optimal discharge temperature setpoint that will minimize power consumption at steady state for the air conditioning system. The setpoints for discharge temperature r_k determined by the ESC are passed to the MPC which ensures that discharge and zone temperature setpoints are met, while maintaining operating constraints. As previously illustrated in Figure 2, the inputs u_k to the VCS are compressor frequency, EEV position, and condenser fan speed. Heat load changes are modeled as a disturbance w_k. The performance metric z_k is power consumption measured in Watts. Simulations are performed on a model of vapor compression systems that has been developed based on the Thermosys toolbox for MATLAB/Simulink (Alleyne, A., *et.*

al., 2012). This model captures pertinent dynamics through a moving-boundary approximation to the heat exchanger dynamics. The parameters used in this model have been calibrated to data obtained from a 2.6 kW single-zone room air conditioner operating in cooling mode.

The following simulations were performed for a fixed zone temperature Tr_{sp}. The ESC tuning parameters are as follows:

$$\begin{aligned} k_g &= 0.06, \\ \alpha &= 0.05, \\ \epsilon &= 0.8, \\ A &= 0.2, \\ \omega &= 20k, \end{aligned}$$

where A is the dither signal amplitude, ω is the dither frequency, and k is the current iteration step. The ESC routine was sampled every 100 s. The MPC was executed every 15 s with a prediction horizon of 64 steps. No projection algorithm was used for the ESC, although this could be added to ensure that the discharge temperature set-point lies within a specified set.

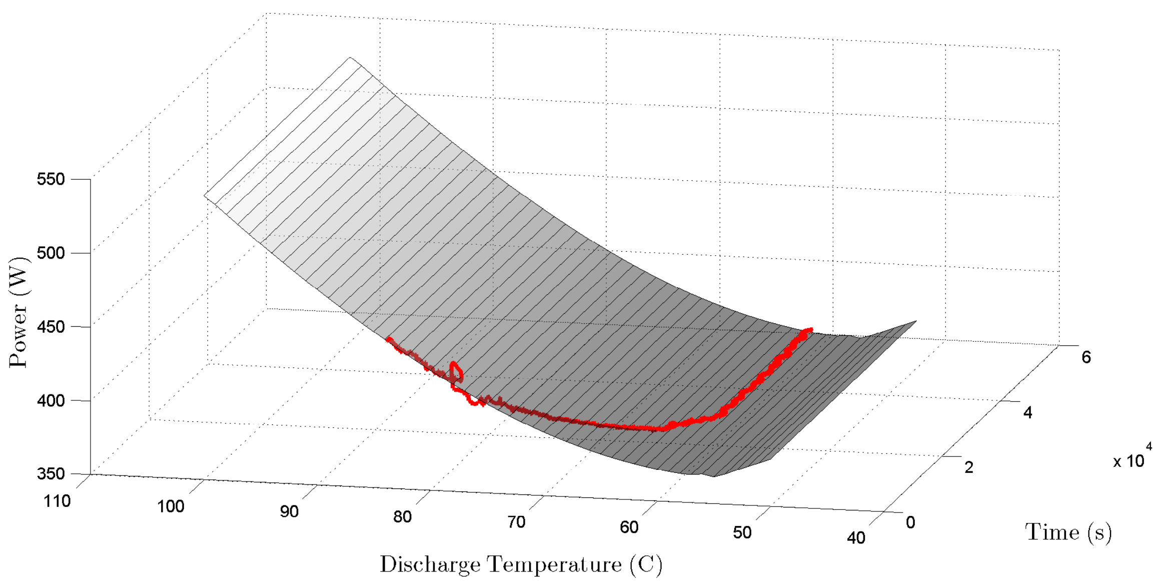

Figure 5: The trajectory of the ESC is shown to go near the minimum of the power surface while maintaining a constant zone temperature for a heat load of 1900 W. The extremum seeker begins at a discharge temperature of 83.8 °C and power consumption of 449.6 W, and converges to 59.3 °C and 372.5 W respectively.

As shown in Figure 5, the function describing power consumption is convex at steady state and can thus be minimized by changing the setpoint for discharge temperature. The ESC routine steers the system toward its optimum. The trajectory generated by the ESC is superimposed on the power surface. Note that the system is initialized at steady state (and achieves the zone temperature setpoint) at the start of the trajectory. Therefore, the significant improvement in power consumption occurs only after the ESC routine is applied.

The trajectory starts at a discharge temperature of 83.8 °C and power consumption of 449.6 W, and converges to 59.3 °C and 372.5 W respectively. The cost surface indicates that the true optimum lies at 55 °C with a power consumption of 367.8 W. The ESC converges to a region near the optimum within five simulated hours and improves power consumption by 77 W or 17%. The convergence rate can be improved by employing a higher optimization gain and a faster ESC sample rate, but system stability may suffer if the ESC is too aggressive.

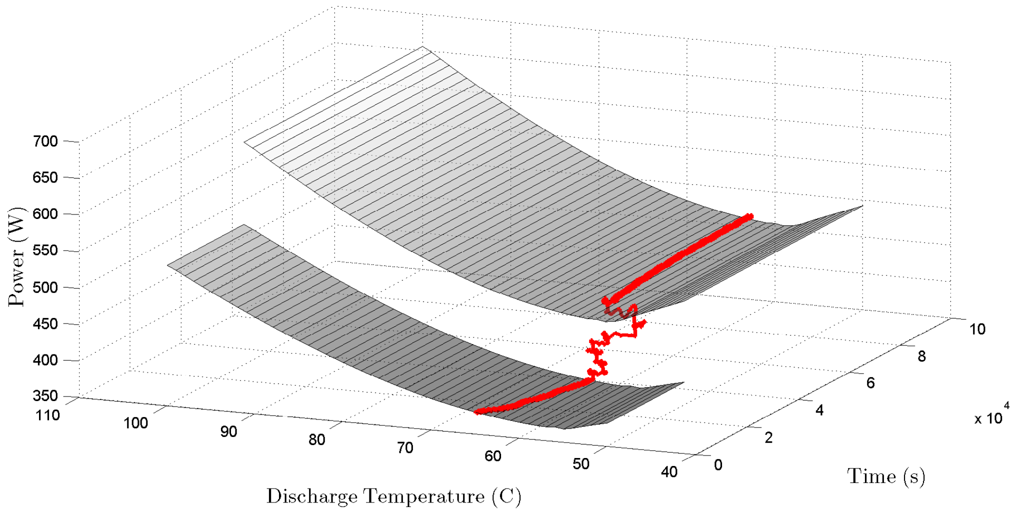

Figure 6: A heat load increase of 400 W is modeled at $30,000$ s. As the cost surface changes due to the new heat load of 2300 W, the time-varying extremum seeking algorithm converges to 62.5 °C and 463.6 W, a region near the true optimum at 58 °C and 456 W.

To demonstrate that the time-varying extremum seeking algorithm can reject disturbances, an increase in heat load of 400 W is applied at $30,000$ s. The extremum seeker appropriately changes the discharge temperature setpoint as shown in Figure 6 in order to obtain the minimum energy operating point. The trajectory shown begins where the previous simulation ended, as the extremum seeker reaches the near-optimum discharge temperature for a heat load of 1900 W. The extremum seeker converges to 62.5 °C and 463.6 W, a region near the true optimum at 58 °C and 456 W. The size of the region near the true optimum that the extremum seeker converges to can be adjusted by varying the optimization gain and dither signal amplitude.

Note that using perturbation-based extremum seeking and its associated convergence characteristics would have prohibitively slow performance. Time-varying extremum seeking allows disturbances to be rejected on time scales necessary for online implementation.

6 CONCLUSION

In this paper, we show that time-varying extremum seeking control is a viable and effective approach to steady state real-time optimization of vapor compression systems. TV-ESC does not rely on averaging singular perturbation in order to obtain an estimate of the gradient of the power consumption, and as a result, converges to optimal setpoints substantially faster than traditional perturbation-based extremum seeking methods.

REFERENCES

Alleyne, A., *et. al.* THERMOSYS 4 Toolbox. University of Illinois at Urbana-Champaign, `http://arg.mechse.illinois.edu/thermosys` (2012).

Burns, D. and Laughman, C. Extremum seeking control for energy optimization of vapor compression systems. In *International Refrigeration and Air Conditioning Conference* (2012).

Goodwin, G. and Sin, K. *Adaptive Filtering Prediction and Control*. Dover Books on Electrical Engineering Series. Dover Publications, Incorporated (2013). ISBN 9780486137728.

Guay, M., Dochain, D., and Dhaliwal, S. A time-varying extremum seeking control technique. In *Proc. 2013 ACC, Washington, D.C.* (2013).

Guay, M. A time-varying extremum-seeking control approach for discrete-time systems. *Journal of Process Control*, 24(3):98 – 112 (2014). ISSN 0959-1524. doi:http://dx.doi.org/10.1016/j.jprocont.2013.11.014.

Killingsworth, N.J. and Krstic, M. PID tuning using extremum seeking: online, model-free performance optimization. *IEEE Control Systems Magazine*, 26(1):70–79 (2006). ISSN 0272-1708. doi:10.1109/MCS.2006.1580155.

Krstic, M. Performance Improvement and Limitations in Extremum Seeking Control. *Systems & Control Letters*, 39(5):313–326 (2000).

Leblanc, M. Sur l'*é*lectrification des chemins de fer au moyen de courants alternatifs de fr*é*quence *é*lev*é*e. Revue G*é*n*é*rale de l'Electricit*é* (1922).

Li, P., Li, Y., and Seem, J.E. Efficient Operation of Air-Side Economizer Using Extremum Seeking Control. *Journal of Dynamic Systems, Measurement, and Control*, 132(3) (2010). ISSN 0022-0434. doi:10.1115/1.4001216.

Ma, Y., Borrelli, F., Hencey, B., Coffey, B., Bengea, S., and Haves, P. Model Predictive Control for the Operation of Building Cooling Systems. *IEEE TRANSACTIONS ON CONTROL SYSTEMS TECHNOLOGY*, 20(3):796–803 (2012). ISSN 1063-6536. doi:10.1109/TCST.2011.2124461.

Oldewurtel, F., Parisio, A., Jones, C., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., and Morari, M. Use of Model Predictive Control and Weather Forecasts for Energy Efficient Building Climate Control. *Energy and Buildings*, 45:15–27 (2012).

Sane, H., Haugstetter, C., and Bortoff, S. Building hvac control systems - role of controls and optimization. In *Proceedings of the 2006 American Controls Conference* (2006).

Tan, Y., Moase, W., Manzie, C., Nesic and, D., and Mareels, I. Extremum seeking from 1922 to 2010. In *29th Chinese Control Conference (CCC)*, pages 14 –26 (2010).

Tyagi, V., Sane, H., and Darbha, S. An extremum seeking algorithm for determining the set point temperature for condensed water in a cooling tower. In *American Control Conference, 2006*, pages 5 pp.– (2006). doi: 10.1109/ACC.2006.1656368.

Zhang, X., Schildbach, G., Sturzenegger, D., and Morari, M. Scenario-Based MPC for Energy-Efficient Building Climate Control under Weather and Occupancy Uncertainty. In *European Control Conference*, pages 1029–1034. Zurich, Switzerland (2013).