Integration of Extremum Seeking and Model Predictive Control for Discrete Time Systems

by

Walter Weiss de la Vega

A thesis submitted to the
Department of Chemical Engineering
in conformity with the requirements for
the degree of Master of Applied Science

Queen's University
Kingston, Ontario, Canada
July 2014

Copyright © Walter Weiss de la Vega, 2014

QUEEN'S UNIVERSITY

Abstract

Department of Chemical Engineering

Master of Applied Science

Integration of Extremum Seeking and Model Predictive Control for Discrete Time Systems

by Walter Weiss de la Vega

This thesis considers a time-varying extremum seeking control algorithm that adjusts set-points provided to a model predictive controller for a vapour compression system. 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. The proposed method uses time-varying extremum seeking, which has faster and more reliable convergence properties for this application. In particular, we regulate the compressor discharge temperature using a model predictive controller with set-points 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 real-time optimization), and use discrete-time extremum seeking control to drive these set-points 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 set-point 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.

Acknowledgements

I would like to express my deepest gratitude and appreciation to Dr. Martin Guay, my supervisor, for his guidance and patience, for all the time and care he dedicated to helping me succeed. My grateful thanks are also extended to Dr. Dan Burns, my gracious host in Boston, for helping me write my first conference paper and ensuring my stay at Mitsubishi Electric Research Labs was pleasant and productive. Special thanks to my colleagues Ehsan Moskhsar and Sean Dougherty for their help and support.

I would also like to thank my caring family for their unconditional support and encouragement. Finally, my heartfelt thanks to Kaitlin Bradley for her love and support when I need it most.

Contents

List of Figures

Abbreviations

ESC	**E**xtremum **S**eeking **C**ontrol(ler)
TV	**T**ime **V**arying
PERB	**P**erturbation-**B**ased
MPC	**M**odel **P**redictive **C**ontrol(ler)
RTO	**R**eal **T**ime **O**timization
RHC	**R**eceding **H**orizon **C**ontrol
LEV	**L**inear **E**xpansion **V**alve
EEV	**E**lectronic **E**xpansion **V**alve
CF	**C**ompressor **F**requency
ODF	**O**utdoor **F**an
VCS	**V**apour **C**ompression **S**ystem
HVAC	**H**eating, **V**entilation, and **A**ir **C**onditioning

Symbols

A, B, C, D	Matrices used in state-space system definition
u	Input vector
x	State vector
y	Output vector
$\hat{y}$	Predicted output
J	Cost function
N	Prediction horizon
Q, R	Positive definite weighting matrices
P	Matrix used for terminal cost
r	Set-point generated by ESC
k_g	ESC optimization gain
z	Variable to be minimized by ESC
ℓ	ESC cost function
e	Estimation error
d	ESC dither signal
ω	Dither frequency
Σ	Information matrix (or summation operator)
α	Estimation gain
ϵ	Time-scale separation parameter
θ	Parameter vector (gradient of ESC cost function)

*F*or my parents, their distant nagging and their ever-present love.

Chapter 1

Introduction

Model predictive control (MPC) has received increased attention in the HVAC (Heating, Ventilation, and Air Conditioning) community over the last years. This interest is largely due to MPC's ability to systematically manage constraints while optimally regulating signals of interest to set-points. For example, in a common formulation of an MPC control problem for variable compressor speed vapour compression machines, the set-points often include the zone temperature and the evaporator superheat temperature. However, the energy consumption of vapour compression systems has been shown to be sensitive to these set-points. 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 vapour compression machines that correlate to efficiency is desired.

Conventionally, methods to improve a vapour compressions system's energy efficiency have relied on the use of models relating commanded inputs, the system's thermodynamic behaviour, and the resulting power consumption. However, these models are highly complex and often unreliable, as they need to be calibrated for a diverse range of environmental factors. Extremum seeking control (ESC), a mechanism by which a variable of interest can

be driven to its optimum, does not rely on a model of the system dynamics and is thus not subject to the unreliability associated with it. As a result, ESC is well suited to optimize a vapour compression system.

In this thesis, we consider ESC and MPC in a hierarchical formulation where the ESC computes set-points which are then provided to the MPC. Among the possible methods for the ESC component, a time-varying ESC (TV-ESC) routine was chosen because of its fast and reliable convergence properties for this application. While perturbation-based ESC, an approach that was also considered, is well known and has been studied extensively, it can suffer from slow convergence rates, a problem that is exacerbated by applications with long time constants such as those associated with thermal systems. The proposed ESC & MPC scheme is applied to a vapour compression system. Specifically, TV-ESC is used to compute set-points for compressor discharge temperature in a vapour compression system, set-points which the MPC then enforces while also maintaining a user-set room temperature. The relationship between compressor discharge temperature and the vapour compression system's power consumption is shown to be convex, thus allowing the ESC to compute set-points that minimize power consumption. Further, because the required cooling capacity (and therefore compressor speed) is affected by measured and unmeasured disturbances, the optimal compressor discharge temperature set-point must vary according to these conditions. The TV-ESC algorithm is shown to provide energy optimal discharge temperature set-points in the presence of disturbances.

1.1 Objectives and Contributions

Most of the research on integrated real-time optimization (RTO) and MPC has focused on a model based approach for the RTO component. The main contribution of this work is to provide a framework for an integrated ESC and MPC approach to real-time optimization that makes minimal assumptions about the underlying cost function. Further, this approach

is applied to a vapour compression system, which in practice usually relies on rudimentary control methods. This particular application is shown to greatly benefit from the MPC's constraint management and the ESC's optimization routine. Further, the use of time-varying ESC, as opposed to traditional ESC methods, positively impacts the rate of convergence of the extremum seeking controller which further improves the performance of the system.

In this thesis, we select the compressor discharge temperature as a signal to be regulated by an MPC controller. Discharge temperature is often measured for equipment protection making it a commonly available signal, and because 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 set-point cannot be regulated to a constant, but instead must vary with these conditions. It is the aim of this thesis to automate the generation of such set-points to maximize energy efficiency.

1.2 Overview of the Thesis

The rest of the thesis is organized into three chapters.

In Chapter 2, a literature review is given. First, model predictive control is discussed including topics such as strategy, stability, and varying set-point. Second, some background on extremum seeking control is presented. Third, the use of real-time optimization in

conjunction with model predictive is briefly studied. Finally, set-point generation for vapour compressions systems is covered.

Next, in Chapter 3, a formal problem statement for the thesis is given. The time-varying extremum seeking control (TV-ESC) routine is presented and compared to standard perturbation-based ESC. Then, some simulation examples are provided to demonstrate how ESC can be used in tandem with MPC or to replace it altogether–while also discussing the limitations and advantages of the suggested control schemes. Finally, TV-ESC is applied as a real-time optimizer to a vapour compression system using MPC.

Lastly, concluding remarks are offered in Chapter 4 including a summary of the simulation results for the vapour compression system. Some thoughts on what future research on this topic could entail are also provided. Auxiliary results are presented in Appendix A.

Chapter 2

Literature Review

This chapter is organized as follows: First, an overview of model predictive control (MPC) and extremum seeking control (ESC) is given. Then, some background on using ESC in conjunction with MPC is provided, followed by a review of the use of ESC and MPC on vapour compression systems.

2.1 Model Predictive Control

MPC is a control framework that originated in the late seventies and that obtains the control input by solving a discrete-time constrained optimal control problem over a defined control horizon. MPC, also known as receding or moving horizon control, makes use of a prediction model to assess the impact of controller inputs on the system. After the optimal control problem is solved, and the first control input in the found optimal control sequence is applied, the optimal control problem is solved again and a new optimal control is computed at the next sampling instant using updated measurements, as shown in Figure 2.1. To summarize, MPC:

- uses a model to predict future process output over a given time-frame, the prediction horizon;
- calculates a control sequence to minimize an objective function;
- implements the first control action in the calculated sequence, and repeats the procedure at the next sampling step, thus resulting in a moving horizon strategy.

As outlined in Camacho and Bordons [2004], MPC has several advantages over other methods:

- tuning is simple and the concept is intuitive;
- the multi-variable case is easily dealt with;
- it naturally introduces feedforward control to compensate for measurable disturbances;
- it handles process constraints with ease and these can be included during the controller design state.

Some of the disadvantages to MPC include potentially high computational load when the controller cannot be calculated beforehand(as in adaptive control), and, more importantly, its reliance on an appropriate model for the process. The performance of the MPC is directly linked to the accuracy of the process model as the control decisions are based on the predictions of the model. To summarize, the effectiveness of MPC is dependent on a model of acceptable accuracy and the availability of sufficiently fast computational resources [Venkat et al., 2008].

MPC has had much success in industry because it provides a framework that handles multi-variable processes and accounts for control objectives and system constraints. Numerous successful MPC applications have been reported over the last two decades [Zeng et al., 2008] including applications for chemical and petro-chemical processes. MPC is a well-developed field, where issues such as online optimization, stability, and performance have been addressed [Bemporad and Morari, 1999].

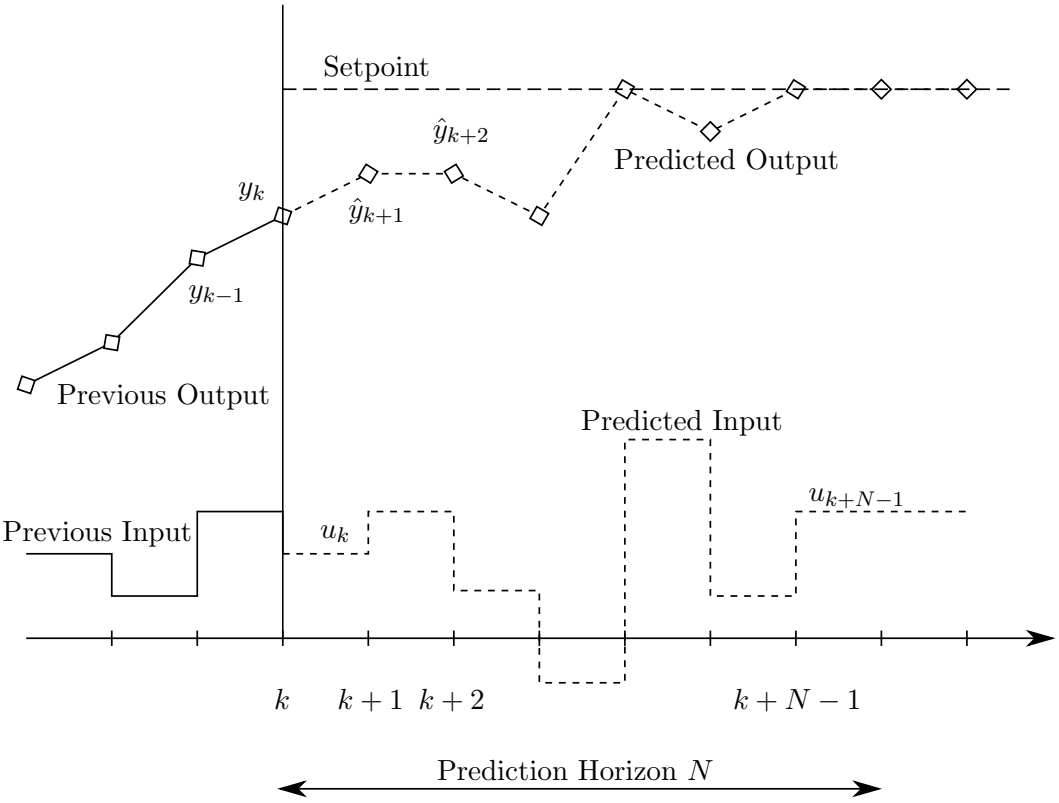

FIGURE 2.1: Model predictive control (MPC) is a control framework that obtains the control input by solving a constrained optimal control problem over a defined control horizon. MPC, or receding horizon control, makes use of a prediction model to assess the impact of controller inputs on the system. After the optimal control problem is solved, and the first control input in the found optimal control sequence is applied, the optimal control problem is solved again and a new optimal control is computed at the next sampling instant using updated measurements.

The general linear MPC problem for discrete-time systems can be formulated as [Goodwin et al., 2001]:

$$\begin{aligned} \min_{u(0),\ldots,u(k)} \quad & \sum_{k=0}^{N-1} \ell(x(k), u(k)) \\ \text{subject to} \quad & u(k) \in \mathbb{U}, x(k) \in \mathbb{X} \text{ for } k = 0, \ldots, N \\ & x(k+1) = Ax(k) + Bu(k) \text{ for } k = 0, \ldots, N-1 \end{aligned} \tag{2.1}$$

where ℓ is a user-defined objective function to be minimized while meeting input and state constraints. The prediction horizon N is the time interval over which the optimization is performed. An optimal sequence of inputs $u(0), \ldots, u(N-1)$ is generated over this horizon, but only the first input is implemented, after which the optimization problem is repeated at the next sampling instant. Stability is not guaranteed for a finite horizon, but can be achieved by adding a terminal cost and/or terminal constraints. Though many applications are best described by nonlinear models, linear approximations are often considered and can be adequate for the control of such systems.

2.1.1 Strategy

Using the process model, future outputs for a predetermined prediction horizon N are calculated. At time k, predicted outputs $\hat{y}_{k+1}$,$\hat{y}_{k+2}$, $\ldots$, $\hat{y}_{k+N}$ are computed. These predicted outputs will depend on the known output values up to that sampling instant and on the predicted inputs.

Future control signals are calculated by solving an optimization problem of the form (2.1) to keep the systems trajectory as close as possible to a reference trajectory, which in some cases is the set-point itself or a trajectory that exponentially approaches the set-point. The optimization problem usually consists of a quadratic costs function that captures the difference between predicted output and reference trajectory. Usually the control effort is also included in the cost function to prevent the control input sequence from varying too much.

The first control signal (in the sequence of optimal control signals calculated at the previous step) is applied to the process. At the next sampling step, the process is repeated, and new information about the process output and thus the system's trajectory relative to the reference is obtained.

2.1.2 Stability

Even when the MPC is able to find a solution that optimizes the process, closed-loop stability is not guaranteed. Indeed, optimality alone does not result in stability. However, the use of terminal costs and/or terminal constraints, among other methods, can guarantee close-loop stability.

In [Mayne et al., 2000], stability and optimality results are summarized for both linear and nonlinear MPC. Here, key findings from this study are summarized.

Let us generalize Equation (2.1) by expanding it to include nonlinear systems and a terminal cost. The system to be controlled is of the form

$$x(k+1) = f(x(k), u(k)), \tag{2.2}$$

$$y(k) = h(x(k)), \tag{2.3}$$

where $f(\cdot)$ is defined by the systems dynamics and has an equilibrium point at the origin. The control and state sequences must satisfy $u(k) \in \mathbb{U}$ and $x(k) \in \mathbb{X}$ respectively, where $\mathbb{U}$ is a convex, compact subset of $\mathbb{R}^p$, and $\mathbb{X}$ is a convex closed subset of $\mathbb{R}^n$, both containing the origin in their interior. Let us define the cost function as follows:

$$V_N(x, \boldsymbol{u}) = \sum_{i=0}^{N-1} \ell(x(i), u(i)) + F(x(N)), \tag{2.4}$$

where $\boldsymbol{u} = \{u(0), u(1), \ldots, u(N-1)\}$ is the control sequence over the prediction horizon N. Let us define a terminal constraint, which is sometimes used for stability purposes, as $x(N) \in X_f \subset \mathbb{X}$. Note there is no need to define the cost function as dependent on the current time-step k because $f(\cdot)$ and $\ell(\cdot)$ are time invariant. We can now formulate the optimal MPC problem as

$$\mathcal{P}_N(x) : V_N^0(x) = \min_{\boldsymbol{u}} \left\{ V_N(x, \boldsymbol{u}) | \boldsymbol{u} \in \mathcal{U}_N(x) \right\}, \tag{2.5}$$

where $x(i)$ denotes the state trajectory resulting from the control sequence $\boldsymbol{u}$ and $\mathcal{U}_N(x)$ is the set of feasible control sequences that satisfy the control, state, and terminal constraints.

In the 1990s, significant research effort was invested in ensuring stability for MPC. One of the proposed methods to modify the optimal MPC problem to ensure closed-loop stability was the addition of a terminal cost $F(\cdot)$ (without a a requirement for a terminal constraint). An alternative method was to impose a terminal equality constraint, where the terminal cost and the terminal constraint would need to satisfy $F(x) \equiv 0$ and $X_f = \{0\}$ respectively. In yet another approach, the terminal constraint set, X_f is required to be a subset of $\mathbb{R}^n$ containing a neighbourhood of the origin where $F(x) \equiv 0$, but a terminal cost is not explicitly used in the problem formulation. Most existing model predictive controller design approaches employ the terminal cost and constraint set approach, where $F(\cdot)$ is picked such that it is approximately equal to the infinite horizon cost function in a suitable neighbourhood of the origin.

2.1.3 Varying Setpoint

Model predictive control can be shown to satisfy constraints and be asymptotically stable when a suitable penalization of the terminal state and an additional constrained are enforced [Mayne et al., 2000]. However, this is suitable for a given operating set-point, but may not be enough to guarantee stability if the set-point changes. In [Ferramosca et al., 2009] it is shown that convergence to the set-point within an admissible set can be guaranteed by adding an artificial offset term to the MPC cost function.

2.2 Extremum Seeking Control

Generally, adaptive control schemes are used to regulate linear and nonlinear systems with known set points or reference trajectories. However, in some applications such as anti-lock

braking system control, the control objective may be to optimize an objective function of unknown parameters or to keep an unknown objective function at its extremum. Extremum seeking control (ESC) can be applied to solve such problems and find the set-points that optimize the unknown or uncertain objective function.

2.2.1 Background

Extremum seeking control dates back to the 1922 application to electric railway systems by Leblanc [1922] and was also popular in the 1940s-1960s. With the rise of model adaptive control, it fell out of favour until the 1990s, as applications such as fluid flow, combustion, and biomedical systems were described by increasingly complex and unreliable models [Ariyur and Krstic, 2003] that lent themselves to the simplified modelling requirements of ESC. The resurgence of ESC took hold following the publication of Krstić and Wang [2000] on the stability of perturbation based ESC for a very general class of dynamical systems.

In [Tan et al., 2006], non-local stability results are provided for a variety of extremum seeking schemes, thus expanding the local stability results in [Krstić and Wang, 2000]. This is achieved by showing semi-global practical stability of the closed loop with respect to the design parameters. In other words, given an arbitrarily large set of initial conditions, it is possible to tune the controller so that all solutions starting in this set eventually converge to a neighbourhood near the optimum. Moreover, it is shown that under certain conditions the ESC parameters can be made arbitrarily small while enlarging the domain of attraction, thus posing an interesting trade-off for the design parameters.

A variety of ESC techniques have been developed over the years, but one of the most popular methods remains the so-called perturbation-based ESC [Ariyur and Krstic, 2003]. This method involves probing the system using a sinusoidal perturbation (or dither signal) in a clever fashion to obtain a gradient estimate.

Work on extremum seeking control prior to the ESC schemes proposed in [Ariyur and Krstic, 2003] generally limited the plant to be a nonlinear static map possibly cascaded with a linear dynamic block. The work in [Krstić and Wang, 2000] generalized the approach to the solution of ESC problems for a general class of nonlinear dynamical systems. The analysis was based on averaging and singular perturbation results and guaranteed local stability. The performance of this approach is highly sensitive to the choice of dither signal. A detailed analysis of how the shape of the dither signal affects the performance of the ESC was carried out in [Tan et al., 2008]. As a result of the difficulties of properly tuning the dither in practice, other adaptive ESC approaches that are less impacted by the choice of dither signal, such as the one in [Guay et al., 2013], can be more attractive.

ESC can not only be used for tuning a set-point to optimize a particular output, but also for tuning parameters of a feedback law. For example, in [Killingsworth and Krstic, 2006], perturbation-based ESC is used to tune proportional-integral-derivative (PID) controllers, a usually cumbersome manual task. As opposed to other automatic tuning methods, ESC does not require knowledge of the plant. The proposed closed-loop ESC scheme makes use of the multiple parameter perturbation-based algorithm proposed in [Ariyur and Krstic, 2003] to optimize the three key parameters of a PID controller: gain, integral time, and derivative time. The performance of the proposed extremum seeking based tuning method was shown to be comparable to that of other popular model-based tuning methods.

In [Guay and Zhang, 2003], an extremum seeking problem for nonlinear systems with parametric uncertainties was proposed. Unlike the usual ESC schemes, the objective function cannot be measured directly, but the explicit structure of the objective function is assumed to be known (which depends on unknown plant parameters). The work is extended to state constrained nonlinear systems in [DeHaan and Guay, 2005]. Studies on persistently exciting signals used to guarantee parameter convergence are shown in [Adetola and Guay, 2007].

A global extremum seeking scheme is presented in [Tan et al., 2009]. The proposed scheme can identify the global optimal value even in the presence of local extrema. Also, when sufficient conditions are met, the ESC can converge to an arbitrarily small neighbourhood of the optimum from an arbitrarily large set of initial conditions.

2.2.2 Strategy

To provide some background on the concept of extremum seeking we refer to [Ariyur and Krstic, 2003], where extremum seeking is presented in great analytical rigour. Figure 2.2 shows the basic ESC scheme described here. Suppose we are interested in driving the static map $f(\theta)$ to its minimum:

$$f(\theta) = f^* + \frac{f''}{2}(\theta - \theta^*)^2 \tag{2.6}$$

where $f(\theta)$ is of class C^2 and $f'' > 0$. Note that any C^2 function can be approximated locally by (2.6) by applying Taylor's Theorem. The aim is to minimize $\theta - \theta^*$ thus driving $f(\theta)$ to its minimum. To obtain gradient information of the static map, a perturbation signal (or dither signal) $a\sin(\omega t)$ is fed into the plant. The high pass filter serves to remove the constant component f^* from the output y. Then, by multiplication with the signal $\sin(\omega t)$, the signal is demodulated. The resulting frequency contains some unwanted higher frequency components that are then attenuated by the integrator. Note that the second sinusoidal signal is additive because the nonlinearity in $f(\theta)$ already provides a multiplication effect.

Let $\tilde{\theta} = \theta^* - \hat{\theta}$ where $\hat{\theta}$ denotes the estimate of the unknown optimal input θ^* in Figure 2.2. Then, by the procedure outlined above, it can be shown that

$$\dot{\tilde{\theta}} \approx -\frac{kaf''}{2}\tilde{\theta} \ . \tag{2.7}$$

With $k > 0$ (note $k < 0$ for maximization problems), equation (2.7) indicates that this is a stable system since $kf'' > 0$. Further, $\tilde{\theta} \to 0$ and thus our system converges to a

small distance of the optimum θ^*. As noted in [Ariyur and Krstic, 2003], given that the perturbation frequency ω is sufficiently large, the output error $y - f^*$ achieves exponential convergence to an $\mathcal{O}(a^2 + 1/\omega^2)$ neighbourhood of the origin.

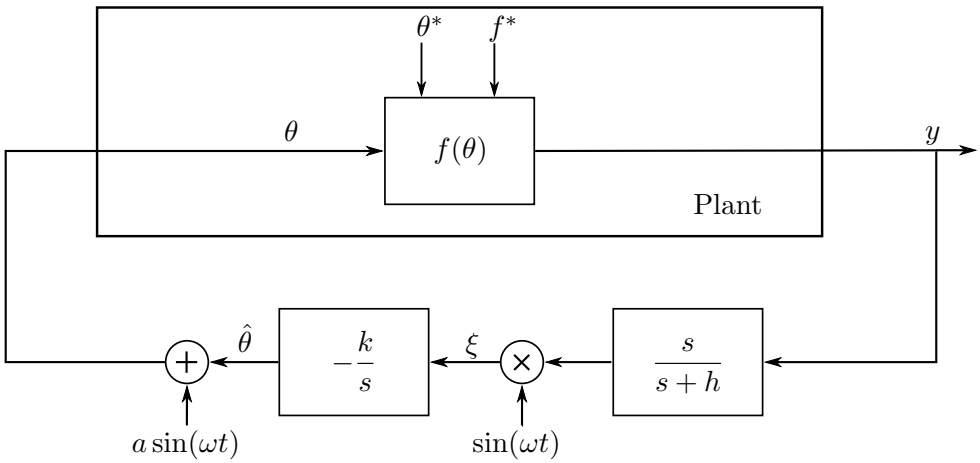

FIGURE 2.2: Extremum seeking control scheme for static map [Ariyur and Krstic, 2003].

2.3 Real-Time Optimization and Model Predictive Control

Real-time optimization aims to optimize an economic objective while the system is online. In conjunction with MPC, the goal generally is to optimize the economic objective while letting the MPC handle system constraints and inputs. In [De Souza et al., 2010] a method to integrate real-time optimization and model predictive control in one layer is suggested. Here, the gradient of the objective is included in the cost function of the MPC, thus allowing for simultaneous computation of control and optimization, simplifying the problem to a QP. However, knowledge of the gradient is needed and tuning is more complex due to the added term in the cost function.

In [Adetola and Guay, 2006], a control algorithm that incorporates real-time optimization and receding horizon control is applied to an output feedback extremum seeking control problem for a linear unknown system, assuming one can provide a suitable functional expression for

the cost function. Similar to the suggested approach in this thesis, the optimum set-point that minimizes a given performance function is computed and then the control input which will drive the system to the computed optimum is obtained. Further, in [Adetola and Guay, 2010] this work is expanded to control of constrained uncertain nonlinear systems. The performance function of interest is assumed to be a known functional of the systems states and parameterized by unknown parameters.

An approach to use extremum seeking to solve finite-horizon LQ control problems, an unconstrained variant of MPC, for unknown discrete time systems was proposed in [Frihauf et al., 2013]. The convergence to the open-loop optimal control sequence that minimizes the cost function is shown using a perturbation and a Newton based extremum seeking controller.

2.4 Set-Point Generation for Vapour Compression Systems to Improve Efficiency

Vapour 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 vapour compression cycle has greatly improved the flexibility of the operation of such systems (Fig. 2.3A). 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 of vapour compression systems (Fig. 2.3B) 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 analysed 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 set-points in steady-state. Previous work has shown that the energy efficiency of these systems is strongly dependent on these set-points [Burns and Laughman, 2012], however, determining appropriate set-points is not always straightforward.

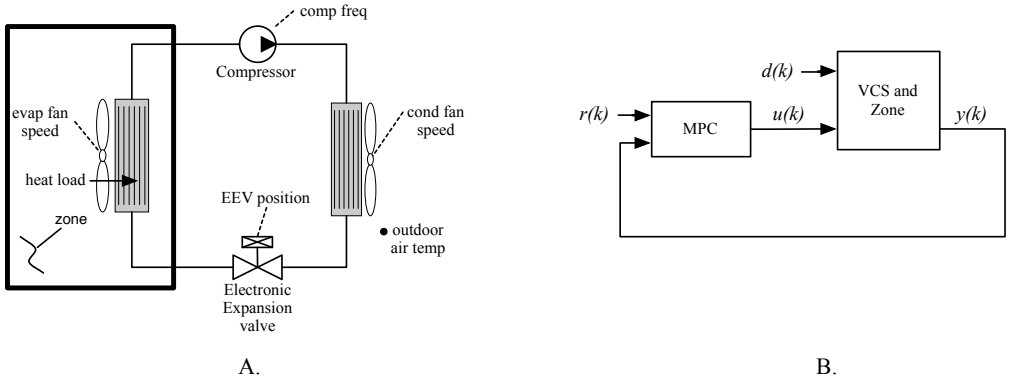

FIGURE 2.3: A. The vapour compression system under study consists of a variable speed compressor, condensing heat exchanger, electronically controlled expansion valve, and evaporating heat exchanger. The inputs to the VCS that are manipulated by the control system include (i) the compressor frequency, (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 vapour compression system and zone to their set-points $r(k)$ in the presence of disturbances $d(k)$ such as changes in outdoor air temperature and heat load.

In a typical air conditioning system, set-points to the controller 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 set-points 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 [Oldewurtel et al., 2012, Zhang et al., 2013, Ma et al., 2012] for more information on problems of this type. However, in this thesis, we consider the optimization of set-points of type (2); that is, internal machine

process variables whose steady-state values determine the energy consumption of the vapour compression machine.

Often, set-points 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 vapour 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 vapour 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 set-points for regulating cycle capacity and efficiency are desired.

One common way to optimize the performance of a vapour 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 behaviour 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 real-time (online) selection of optimal set-points is tenuous.

Recently, model-free methods that operate in real-time and aim to optimize a cost have

received increased attention and have demonstrated improvements in the optimization of vapour compression systems and other HVAC applications [Burns and Laughman, 2012, Sane et al., 2006, Li et al., 2010, 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 [Krstić and Wang, 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 vapour 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 thesis, we apply time-varying extremum seeking to the problem of obtaining set-points that optimize energy efficiency in a vapour compression system.

Chapter 3

Real-Time Extremum Seeking Optimization for MPC

In this chapter, we first establish a framework for ESC and MPC in a hierarchical structure where the ESC is operating as the real-time optimizer. We then apply the proposed method to a vapour compression system. Extremum seeking is used to provide set-points to the MPC in order to optimize a measured variable of interest for the VCS. The objective is to optimize a Refrigeration & Air-Conditioning (RAC) unit's efficiency at steady-state by reducing power consumption.

The proposed model predictive controller is tracking two variables: room temperature (set-point chosen by the user) and compressor discharge temperature. Using data collected from simulations for the RAC, it is shown that there is a convex relationship between discharge temperature and power consumption at steady-state for a given room temperature. Therefore, the extremum seeking algorithm can find a set-point for discharge temperature that minimizes power consumption while maintaining the room temperature set-point. Since an explicit model relating discharge temperature to power consumption may be difficult

to obtain, due to the complexity of the system, ESC lends itself to this problem–as power consumption can be measured.

3.1 Problem Description

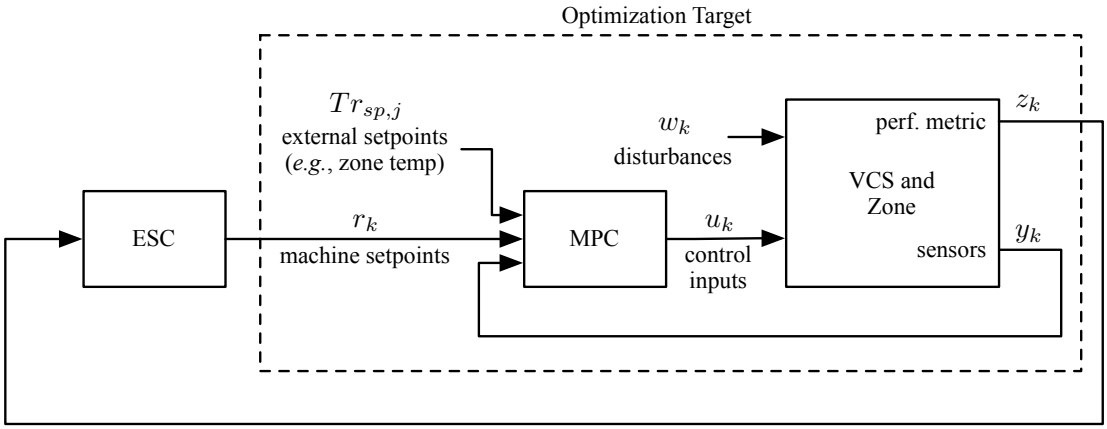

FIGURE 3.1: Block diagram of system and controllers.

Consider the dynamics of the closed-loop system to be given by

$$x_{k+1} = x_k + f(x_k, r_k) \tag{3.1a}$$

$$y_k = g(x_k) \tag{3.1b}$$

$$z_k = h(x_k) \tag{3.1c}$$

where $x_k \in \mathbb{R}^n$ is the vector of state variables at the k^{th} time step, $r_k \in \mathbb{R}^p$ is the vector of reference signals sent to the MPC, and $z_k \in \mathbb{R}$ is the measured variable to be minimized. The goal is to drive the system to the equilibrium x^* and r^* so that z_k is minimized. The optimal input r^* is the set-point that minimizes the variable of interest z_k.

A model predictive controller then tracks these reference signals and controls the plant while enforcing input and output constraints. The MPC optimization problem (3.44) is solved

online at each time step, providing a set of control moves from the current time step k to a time N steps in the future, i.e. over the control horizon. The first control move u_k is applied to the system, new output measurements y_k are obtained, and the optimization problem is recalculated.

In the application considered in this thesis, the goal is to improve the efficiency of a vapour compression system by supplying the MPC with a discharge temperature set-point that minimizes power consumption, while tracking a given zone temperature. Here, r_k is a scalar value for discharge temperature set-point (the optimizer), z_k is power consumption (the measured variable to be minimized), and zone temperature is set externally. The MPC is designed for zone temperature and discharge temperature tracking, as well as to enforce constraints on the system inputs and outputs. More detail on the vapour compression system and on the MPC for this particular application is given in Section 3.5.

Generally, MPC is used for set-point tracking. Here we consider a set-point that is generated by a real-time optimization algorithm. The proposed real-time optimizer is a time-varying extremum seeking control (TV-ESC) scheme [Guay et al., 2014]. In this study, we provide some evidence that the TV-ESC approach provides faster convergence relative to standard perturbation based ESC. The exact structure of the MPC is not significant in this application as long as it can implement the set-points generated by the RTO. However, for the case study provided at the end of the chapter, a formulation for a linear MPC is shown.

3.2 Extremum Seeking Controller

The ESC is tasked with supplying the MPC with the set-points for the optimizer r that minimize the variable of interest z; in the context of the VCS, these signals are discharge temperature and power consumption respectively. We follow the discrete-time ESC update

law outlined in Guay [2014]. First, we show the derivation for a static map, followed by the ESC update law for a system with dynamics such as the vapour compression system.

As shown in Figure 3.2, 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 output (e.g. change in power consumption) for the gradient estimation. The estimated gradient $\hat{\theta}_k$ will be used to parameterize the unknown but measured cost function describing the variable to be minimized. 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 minimize the variable of interest z_k (e.g. power consumption). The new set-point is applied by the MPC, and the ESC algorithm is repeated at the next ESC sample time.

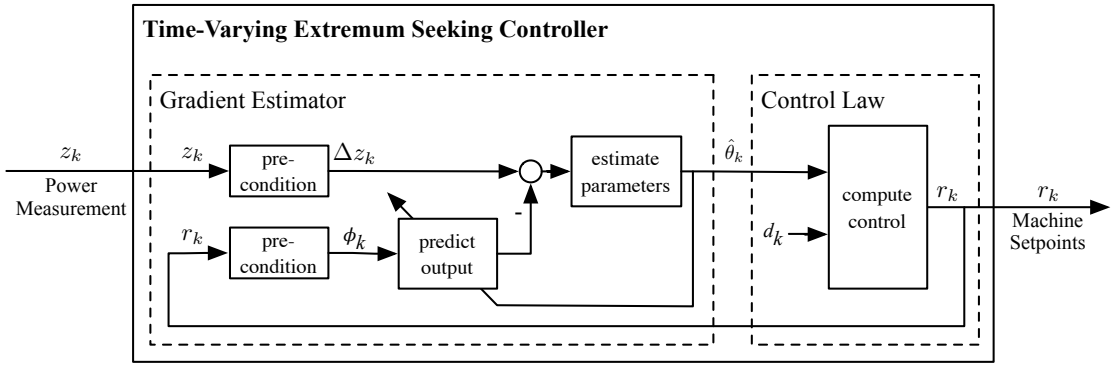

FIGURE 3.2: Overview of the TV-ESC algorithm.

3.2.1 Static Map

Recall the nonlinear system

$$x_{k+1} = x_k + f(x_k, r_k) \tag{3.2a}$$

$$z_k = h(x_k) \ . \tag{3.2b}$$

The goal is to drive the system to the equilibrium x^* and r^* so that z_k is minimized. The steady-state map $\pi(r) \in \mathbb{R}^n$ is such that $x_{k+1} = x_k$ and therefore:

$$f(\pi(r), r) = 0 \tag{3.3}$$

The steady-state cost function is then given by:

$$z = h(\pi(r)) = \ell(r) \tag{3.4}$$

Therefore, at steady-state the problem is reduced to finding the minimizer r^* of $z = \ell(r^*)$. The equilibrium cost $z = \ell(r^*)$ should satisfy the following optimality conditions:

Assumption 3.1. *The equilibrium cost (3.4) is such that:*

$$\frac{\partial \ell(r^*)}{\partial r} = 0 \tag{3.5}$$

$$\frac{\partial^2 \ell(r^*)}{\partial r \partial r^T} \geq \beta I \quad \forall r \in \mathcal{R} \tag{3.6}$$

where β is a strictly positive constant. Thus, ℓ is strictly convex.

Assumption 3.2. *The static-map ℓ is such that*

$$\|z\| \leq Z \tag{3.7}$$

$$\left\| \frac{\partial \ell}{\partial r} \right\| \leq L_1 \tag{3.8}$$

$$\left\| \frac{\partial^2 \ell}{\partial r \partial r^T} \right\| \leq L_2 \tag{3.9}$$

$\forall r \in \mathcal{R}$ with positive constants $Z > 0, L_1 > 0$ and $L_2 > 0$. In other words, the cost function and both its first and second derivatives are bounded.

The input r is assumed to be a time-varying signal, then $z_k = \ell(r_k)$. Further, let $\Delta z_k = z_{k+1} - z_k$. Now we can express the change in the cost function z in terms of the input r:

$$\Delta z_k = \ell(r_{k+1}) - \ell(r_k) \ . \tag{3.10}$$

Given that the function $\ell(r)$ is of class C^1, a linear transform can be applied to Δz_k to parametrize as follows:

$$\Delta z_k = \int_0^1 \ell'(\lambda r_{k+1} + (1-\lambda) r_k) d\lambda \Delta r_k \tag{3.11}$$

where $\ell'(r) = \frac{\partial \ell}{\partial r}$, $\Delta r_k = r_{k+1} - r_k$.

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

$$\Delta z_k = \theta_k^T \Delta r_k = \phi_k^T \theta_k \tag{3.12}$$

where θ_k is a time-varying parameter used to describe the quasi steady-state dynamics of the system during real-time optimization and is defined as

$$\theta_k = \int_0^1 \ell'(\lambda r_{k+1} + (1-\lambda) r_k) d\lambda \ . \tag{3.13}$$

Now that the cost function has been parametrized using a time-varying parameter, an approach to estimate θ_k is needed. First, assume that r is bounded and lies in a set of admissible inputs R, which can be ensured via a projection algorithm, then

Assumption 3.3.

$$r_k \in \mathcal{R} \ \ \forall k \geq 0 \ .$$

Let the estimator for (3.12) be

$$\Delta\hat{z}_k = \hat{\theta}_k^T \Delta r_k = \phi_k^T \hat{\theta}_k \tag{3.14}$$

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

Using a recursive least squares filter with forgetting factor $\alpha \geq 0$ and information matrix $\Sigma \in \mathbb{R}^{n_\theta \times n_\theta}$, the resulting parameter estimation update law is given by:

$$\Sigma_{k+1} = \alpha\Sigma_k + \phi_k\phi_k^T, \quad \Sigma_0 = qI > 0 \tag{3.15}$$

$$\bar{\hat{\theta}}_{k+1} = Proj\left[\hat{\theta}_k + \frac{1}{\alpha}\Sigma_k^{-1}\phi_k(1 + \frac{1}{\alpha}\phi_k^T\Sigma_k^{-1}\phi_k)^{-1}(e_k), \Theta_0\right] \tag{3.16}$$

where q is a strictly positive constant and $Proj$ is an orthogonal projection operator. For a more detailed discussion on this operator see Guay [2014] and Goodwin and Sin [2009].

Applying the Woodbury matrix identity to (3.15), we can update the covariance matrix Σ^{-1} directly using the following expression:

$$\Sigma_{k+1}^{-1} = \frac{1}{\alpha}\Sigma_k^{-1} - \frac{1}{\alpha^2}\Sigma_k^{-1}\phi_k(1 + \frac{1}{\alpha}\phi_k^T\Sigma_k^{-1}\phi_k)^{-1}\phi_k^T\Sigma_k^{-1} \quad . \tag{3.17}$$

The trajectories of the system are assumed to meet the following condition.

Assumption 3.4 (Goodwin and Sin [2009])**.** *Persistency of excitation condition. There exist constants $\beta_T > 0$ and $T > 0$ such that*

$$\frac{1}{T}\sum_{i=k}^{k+T-1} \phi_i\phi_i^T > \beta_T I, \forall k > T.$$

For the univariate case, it can be shown that a sinusoidal perturbation signal sufficed to meet this condition. Generally, for the multivariate case a perturbation signal with a different frequency for each input is needed to satisfy the persistency of excitation condition.

Finally, we use a gradient descent controller to reach the system's extremum. The controller is defined as follows:

$$r_{k+1} = r_k - k_g \hat{\theta}_k + d_k \tag{3.18}$$

where d_k is a bounded dither signal and k_g is the optimization gain. Assume $|d_k| \leq D \ \forall k \geq 0$ where D is a positive constant (can be shown to be the amplitude of the signal if the dither is sinusoidal).

Theorem 3.5 (Guay [2014]). *The extremum seeking controller is such that the system converges to an $\mathcal{O}(D)$ neighbourhood of the minimizer r^* of the static cost z. The size of this neighbourhood can be adjusted by setting the gains α and k_g .*

For a proof of Theorem 3.5, or more detail on the derivation of the update law, see Guay [2014]. To summarize, the update law for a static map is as follows:

$$r_{k+1} = r_k - k_g \bar{\hat{\theta}}_k + d_k \tag{3.19a}$$

$$\phi_k = \Delta r_k = r_{k+1} - r_k \tag{3.19b}$$

$$\Delta \hat{z}_k = \phi_k^T \bar{\hat{\theta}}_k \tag{3.19c}$$

$$\Sigma_{k+1}^{-1} = \frac{1}{\alpha}\Sigma_k^{-1} - \frac{1}{\alpha^2}\Sigma_k^{-1}\phi_k(1 + \frac{1}{\alpha}\phi_k^T\Sigma_k^{-1}\phi_k)^{-1}\phi_k^T\Sigma_k^{-1} \tag{3.19d}$$

$$\bar{\hat{\theta}}_{k+1} = Proj\left[\bar{\hat{\theta}}_k + \frac{1}{\alpha}\Sigma_k^{-1}\phi_k(1 + \frac{1}{\alpha}\phi_k^T\Sigma_k^{-1}\phi_k)^{-1}(e_k), \Theta_0\right] \tag{3.19e}$$

3.2.2 Implementation of Update Law for Systems with Dynamics

The time-varying parameter estimation procedure for a system with dynamics, such as the vapour compression system, is given below, followed by a gradient descent controller to steer

the system to its extremum.

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:

$$\Sigma_{k+1}^{-1} = \Sigma_k^{-1} + \epsilon\left(\frac{1}{\alpha} - 1\right)\Sigma_k^{-1} - \frac{\epsilon}{\alpha^2}\Sigma_k^{-1}\phi_k(1 + \frac{1}{\alpha}\phi_k^T\Sigma_k^{-1}\phi_k)^{-1}\phi_k^T\Sigma_k^{-1} \tag{3.20}$$

$$\bar{\hat{\theta}}_{k+1} = Proj\left[\hat{\theta}_k + \frac{\epsilon}{\alpha}\Sigma_k^{-1}\phi_k(1 + \frac{1}{\alpha}\phi_k^T\Sigma_k^{-1}\phi_k)^{-1}(e_k), \Theta_0\right] \tag{3.21}$$

where $\Sigma^{-1} \in \mathbb{R}^{n_\theta \times n_\theta}$ 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 [2009].

The gradient descent controller is given by:

$$r_{k+1} = r_k - \epsilon k_g \hat{\theta}_k + \epsilon d_k \tag{3.22}$$

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

Together, the iterative extremum seeking routine is given by:

$$r_{k+1} = r_k - \epsilon k_g \bar{\hat{\theta}}_k + \epsilon d_k \tag{3.23a}$$

$$\phi_k = \Delta r_k = r_{k+1} - r_k \tag{3.23b}$$

$$\Delta\hat{z}_k = \phi_k^T \bar{\hat{\theta}}_k \tag{3.23c}$$

$$\Sigma_{k+1}^{-1} = \Sigma_k^{-1} + \epsilon(\frac{1}{\alpha} - 1)\Sigma_k^{-1} - \frac{\epsilon}{\alpha^2}\Sigma_k^{-1}\phi_k(1 + \frac{1}{\alpha}\phi_k^T\Sigma_k^{-1}\phi_k)^{-1}\phi_k^T\Sigma_k^{-1} \tag{3.23d}$$

$$\bar{\hat{\theta}}_{k+1} = Proj\left[\bar{\hat{\theta}}_k + \frac{\epsilon}{\alpha}\Sigma_k^{-1}\phi_k(1 + \frac{1}{\alpha}\phi_k^T\Sigma_k^{-1}\phi_k)^{-1}(e_k), \Theta_0\right] \quad . \tag{3.23e}$$

Note that the time-varying extremum seeking controller does not require averaging the effect of the perturbation as in traditional perturbation-based extremum seeking controller

design techniques. As a result, time-varying extremum seeking converges to the optimum substantially faster, as demonstrated in the following section.

3.3 Comparison of Perturbation-Based and Time-Varying 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 3.3A). This example is taken from Weiss, Burns, and Guay [2014]. The equations for this system are given by

$$x_{k+1} = 0.8x_k + r_k \tag{3.24}$$

$$z_k = (x_k - 3)^2 + 1 \tag{3.25}$$

which has a single optimum point at

$$r^* = 0.6 \tag{3.26}$$

$$z^* = 1. \tag{3.27}$$

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

To illustrate the difference in convergence rates, a discrete-time perturbation-based extremum seeking controller (PERB-ESC) and a time-varying extremum seeking controller (TV-ESC) are 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

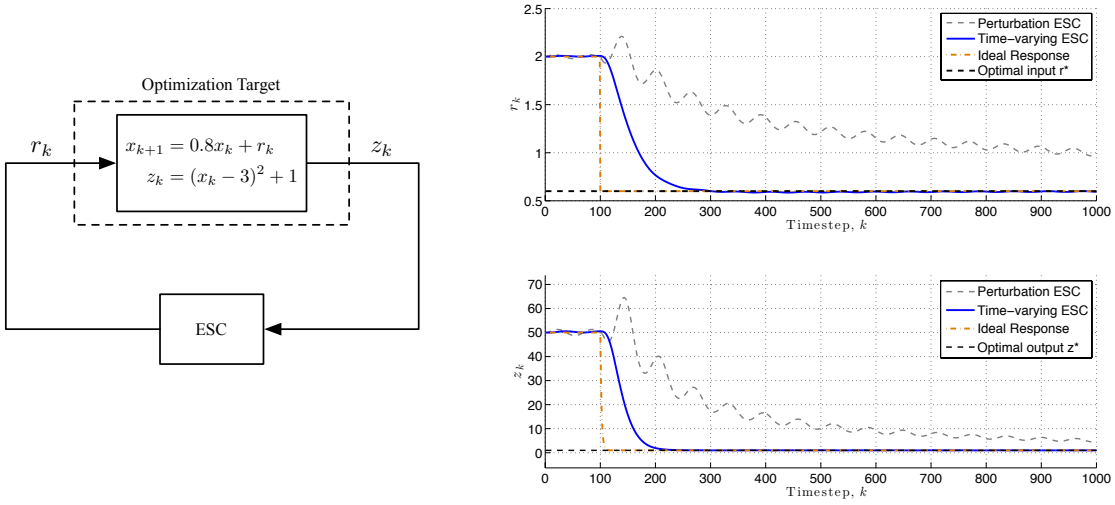

A. B.

FIGURE 3.3: Comparing TV-ESC with perturbation ESC. For this application, TV-ESC converges considerably faster to the optimum.

simulation are

$$d_k = 0.2\sin(0.1k) \tag{3.28}$$

$$\omega_{LP} = 0.03 \tag{3.29}$$

$$K = -0.005 \tag{3.30}$$

where d_k is the sinusoidal perturbation, ω_{LP} is the cut-off 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 perturbation-based ESC formulation, see Killingsworth and Krstic [2006].

The parameters used for the TV-ESC are

$$d_k = 0.001\sin(0.1k) \tag{3.31}$$

$$k_g = 0.001 \tag{3.32}$$

$$\alpha = 0.1 \tag{3.33}$$

$$\epsilon = 0.4 \tag{3.34}$$

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 3.3B. The perturbation-based ESC method converges to a neighbourhood 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 Section 3.5.2, we apply the TV-ESC algorithm to the problem of selecting set-points for a MPC controller of a vapour compression machine and present simulation results.

3.4 Simulation Examples

In this section, we present three examples to demonstrate different applications of extremum seeking control to the problem of optimal set-point generation in MPC. First, we use ESC and MPC in a hierarchical formulation where the ESC is used to compute set-points that are optimal for an unknown but measured cost function. The MPC then implements the set-points computed by the ESC. Second, we use ESC to emulate the behaviour of a MPC so that it simultaneously finds the optimal set-point and the inputs necessary to achieve it. Third, we assume that we can generate values for a performance metric over a prediction horizon and use the ESC to drive the system to its extremum. The latter formulation is akin to economic MPC approaches [Rawlings et al., 2012]. The main difference here is that

the exact form of the prediction model has no impact on the ESC approach. The objective is to allow the use of large unstructured process simulations with MPC without taking into account any mathematical formulation of the process model.

Example 3.1. In the following example, we simplify the problem described in section 3.1 to demonstrate the expected behaviour of the integrated ESC and MPC controller. The extremum seeking controller computes set-points for the MPC, which in turn drives the plant to the desired set-point that minimizes the ESC's cost function.

Consider the quadratic cost function to be minimized by the ESC:

$$z_k = 10 + 50(y_k - 5)^2 \tag{3.35}$$

where y_k is the plant's output and is described by

$$y_k = 0.8531y_{k-1} + 0.1u_k^2 + 0.2713u_k + b_k \tag{3.36}$$

where b_k is a disturbance modelled as a normally distributed pseudorandom number with 0 mean and 0.05 standard deviation, and u_k is the input from the MPC to the plant.

A perfect match, aside from the disturbance b_k, between the plant and the prediction model is assumed. The stage cost of the MPC is defined as follows:

$$J = \sum_{k}^{k+N-1} R(y_i - r_k)^2 + Q(u_i - u_{i-1})^2 \tag{3.37}$$

where $R = Q = 10$ and $N = 15$.

The extremum seeking controller was implemented using the following parameters:

$$k_g = 0.006$$

$$\alpha = 0.1$$

$$\epsilon = 0.2$$
$$d_k = 10^{-3}\sin(20k) \ .$$

A projection algorithm was used to enforce $\|\hat{\theta}_k\| \leq 300$.

$$Proj(\tau) = \begin{cases} \tau & \text{if } \|\hat{\theta}_k\| < 300 \text{ or } 2\hat{\theta}_k^T\tau \leq 0 \\ \left(I - cI\frac{\hat{\theta}_k\hat{\theta}_k^T}{\hat{\theta}_k^T\hat{\theta}_k}\right)\tau & \text{otherwise.} \end{cases}$$

The simulation was initialized with $u_0 = 1$, $r_0 = 0$, and $y_0 = 3$. As shown in Figure 3.5, the ESC starts generating set-points for the MPC after its activation at 200 s which the latter then tracks. As the input converges to $r_k^* = 5$, we observe the unknown but measured performance metric z_k reach its minimum. Note that the ESC converges in less than 150 sampling steps.

This example is similar to the problem posed in Section 3.1. It mainly differs in our consideration of a simplified nonlinear recursive model to play the role of the VCS process model. Here, it is also assumed that there is no model mismatch. As a result of the system's much faster dynamics (compared to the VCS), there is no need to sample the ESC at a rate different from the MPC's; instead we rely on ϵ only to provide the required time-scale separation. In this case, the MPC is only tasked with tracking the trajectory of one output signal. The MPC for the VCS in Section 3.5.2 is tasked with maintaining one signal at a reference value while tracking a second signal using ESC computed set-points. For the vapour compression system, these factors are of importance and lead to lower ESC sample rates as the system dynamics require time to unfold.

Example 3.2. In this example, we assume a known plant model, but unknown cost function that we wish to minimize. We could, as demonstrated in the previous example, approach this using both extremum seeking and model predictive control. However, extremum seeking

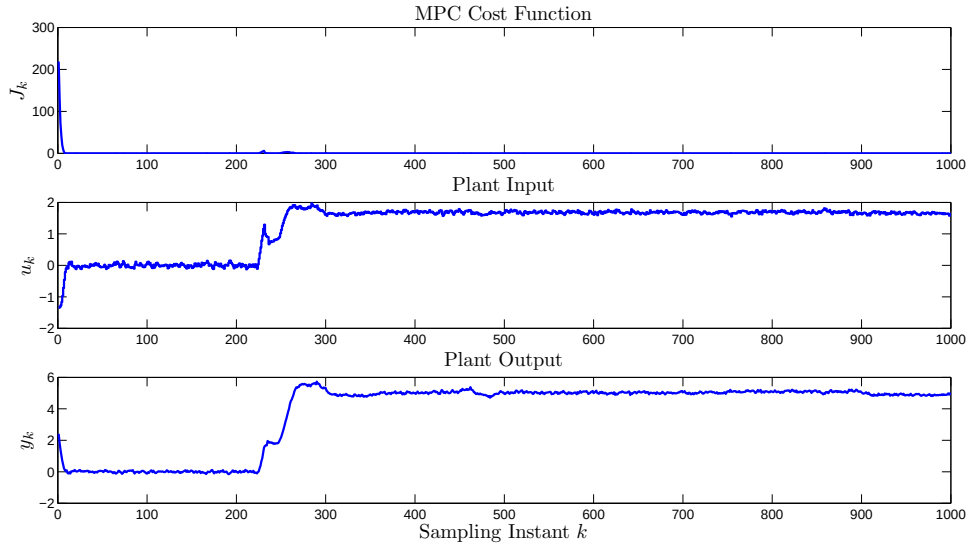

FIGURE 3.4: MPC cost function, generated input, and plant response for Example 3.1 are depicted here. Note that the MPC ensures that the plant output y_k tracks the set-points r_k generated by the ESC.

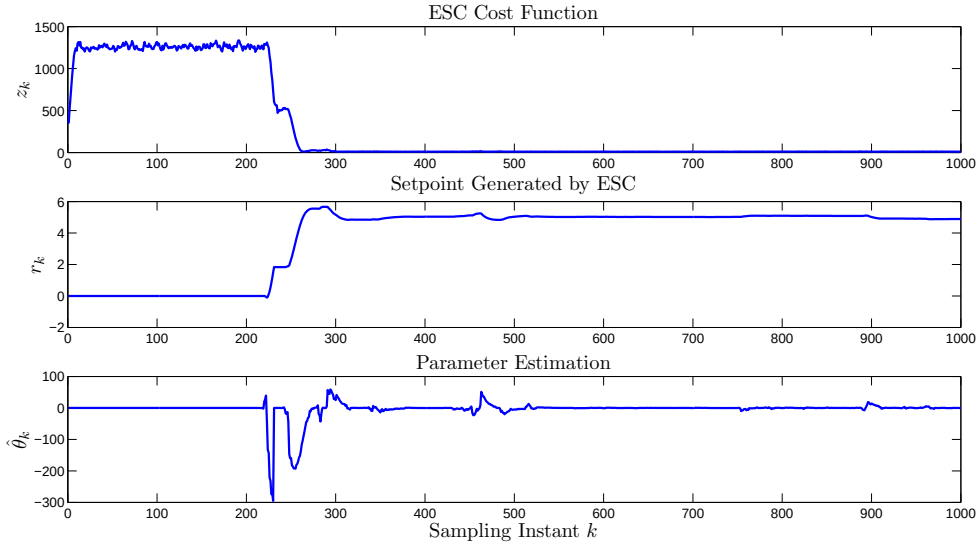

FIGURE 3.5: ESC cost function, generated set-point (which is fed to the MPC), and estimated parameter $\hat{\theta}_k$ for Example 3.1 are shown in this figure. After the ESC routine is activated at $t = 200\,\mathrm{s}$, we see the cost function z_k being driven to its minimum as the set-point correctly converges to $r^* = 5$. Also, the estimated parameter $\hat{\theta}_k$ converges to zero as expected.

control can also be used as a real-time optimization routine to compute input trajectories similar to a MPC's. The plant model is used to predict the plant's response over the control horizon, thus resulting in an emulation of a MPC. The advantage of this approach over the previous example is ease of implementation, as we are now dealing with just one controller instead of two in series. Also, though the convergence rate will be slower in most cases, the computational load should be much lower than before, as the ESC relies on a purely algebraic update law. On the other hand, this approach may not benefit from MPC's stability and convergence guarantees, though this element is still under investigation: evidence shows that ESC can be used for MPC even for stabilization purposes [Guay et al., 2014]. Further, this example does not involve any constraints, which MPC is adept at handling.

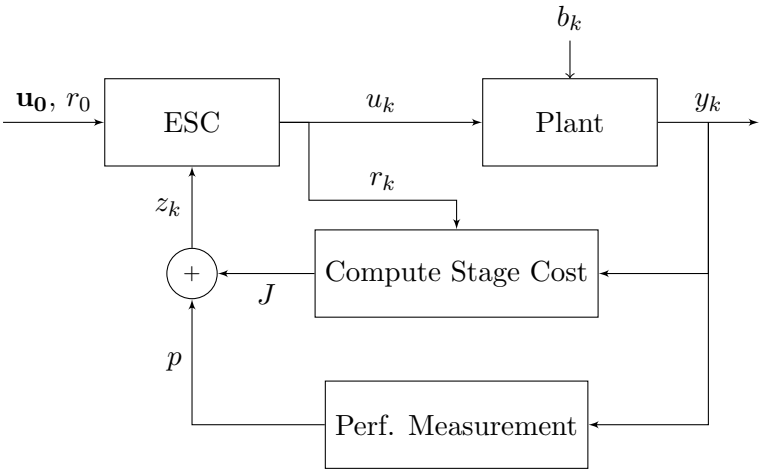

FIGURE 3.6: Block diagram for Example 3.2.

The cost function to be minimized consists of the cost function for the MPC in the previous example, in addition to the unknown performance metric. The ESC will attempt to find the inputs u_k^* that minimize the cost function over the control horizon, in addition to the set-point r_k^* that minimizes the performance metric.

$$z_k = \lambda J(u_k, u_{k+1}, \ldots, u_{k+N-1}) + (1-\lambda)p(r_k) \tag{3.38}$$

where $0 < \lambda < 1$ is a weighting parameter and $p(r_k)$ is the unknown but measured

performance metric. For the following simulation, we will use the same performance metric and plant model as in the previous example. Then, the cost function is given by

$$z_k = \sum_{k}^{k+N-1} R(y_i - r_k)^2 + Q(u_i - u_{i-1})^2 + 10 + 50(y_k - 5)^2 \tag{3.39}$$

where $R = Q = 10$ and $N = 10$.

The following parameters were used for the extremum seeking controller:

$$k_g = 0.003$$

$$\alpha = 0.01$$

$$\epsilon = 0.7$$

$$d_{i,k} = 10^{-2} \sin(5^i k) \quad \forall i = 1, \ldots, n$$

where the dither signal $d_{i,k}$ is defined such that it has a different frequency for each input signal to meet the persistency of excitation condition. For this simulation, we have $n = N + 1$ inputs to be generated by the ESC, which correspond to the plant inputs u_k over the control horizon in the cost function, as well as the performance metric optimizer r_k. It is important to note that, as in the case of MPC, only the first input in the sequence of inputs computed by the ESC over the control horizon is applied to the plant. A projection algorithm, as in the previous example, is used to enforce $\|\hat{\theta}_k\| \leq 100$. The simulation was initialized with $u_0 = 1$, $r_0 = 0$, and $y_0 = 0$. The extremum seeking routine starts at $t = 200\,\text{s}$.

As shown in Figures 3.7 and 3.8, the ESC drives the plant to the expected set-point $r_k^* = 5$. This method is substantially slower to converge than the previous approach. However, it also offers some advantages such as ease of implementation–one controller as opposed to two–and lower computational load (as the ESC update law is purely algebraic). It is important to note that ESC, unlike MPC, does not guarantee that the generated sequence of inputs $u_k, \ldots, u_{k+N-1}$ over the control horizon is optimal at every step. However, as

the ESC routine is repeated, the generated inputs should converge to the optimal sequence of inputs. As a result, the ESC is able to make use of the plant model and the predicted output information over the control horizon, and is thus more effective than a greedy control scheme with one step lookahead.

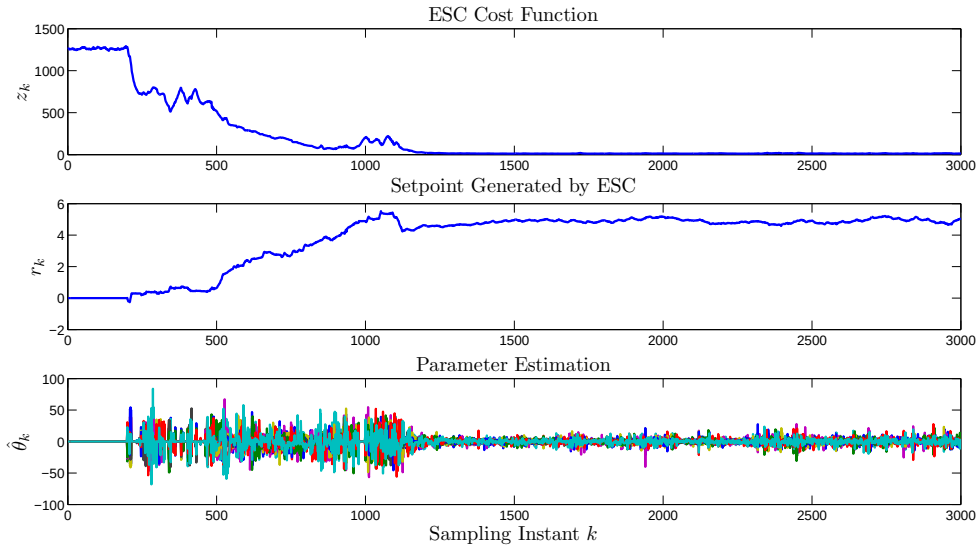

FIGURE 3.7: Above, the ESC cost function, estimated plant set-point, and estimation parameters for Example 3.2 are shown. The extremum seeking controller is generating both the optimal set-point and the inputs necessary to achieve it.

Example 3.3. Now suppose that we can produce values for the performance metric, that we are interested in minimizing, over the prediction horizon and that we have access to the plant model. Note that we need not have an exact mathematical formulation for the performance metric. Such situations arise when the MPC model considered is generated from a third party simulation or from a very complex process model.

Let the cost function be given by:

$$z_k = \sum_{k}^{k+N-1} \lambda p(y_i) + Q(u_i - u_{i-1})^2 \tag{3.40}$$

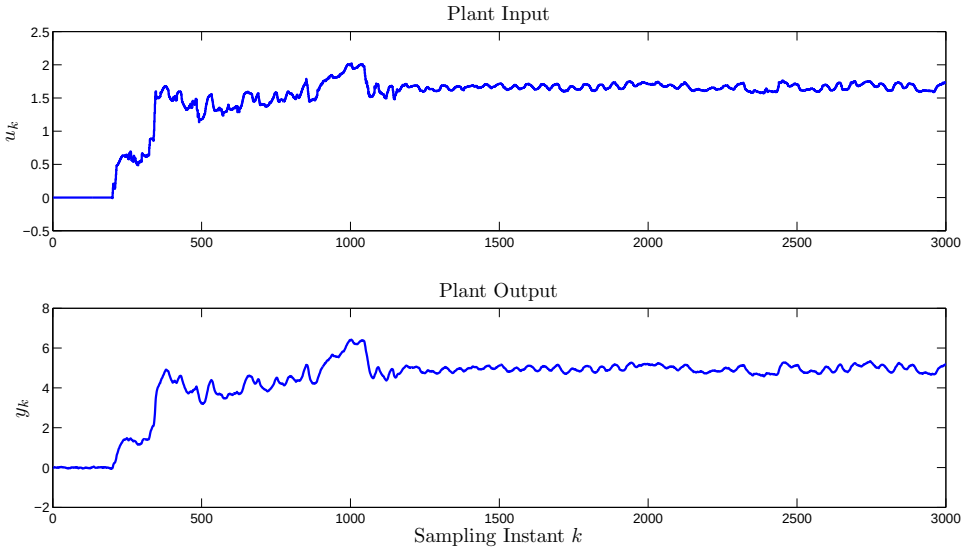

FIGURE 3.8: Plant input and output for Example 3.2. Note that the inputs in this case are generated by the ESC instead of a MPC.

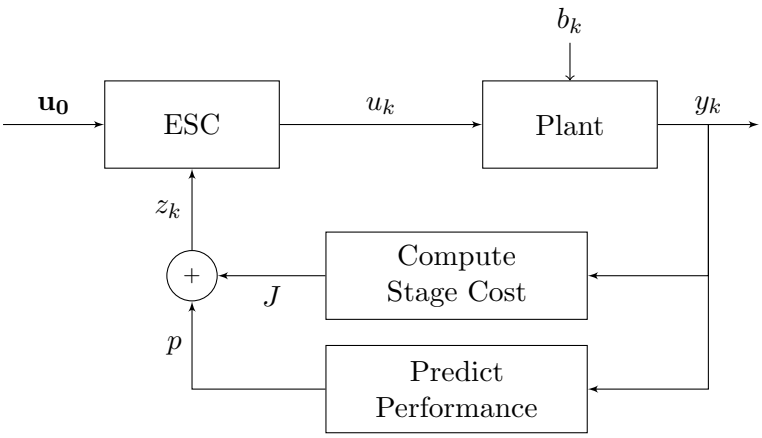

FIGURE 3.9: Block diagram for Example 3.3.

where λ and Q are a weighting constant and a weighting matrix respectively. For this example, we use the same plant model as before and the following performance function:

$$p_k(y_k) = 10 + 50(y_k - 5)^2 \ . \tag{3.41}$$

The following parameter values were used for the ESC:

$$k_g = 0.005$$
$$\alpha = 0.2$$
$$\epsilon = 0.9$$
$$d_{i,k} = 10^{-2}\sin(10^i k) \quad \forall i = 1, \ldots, n$$

where we have n input signals generated by the ESC–though, as before, we only apply the first in the sequence to the plant–corresponding to each time step in the control horizon $N = 10$. The remaining parameters are defined as $\lambda = 1$, $Q = 10$, $y_0 = 0$ and $u_{i,0} = 0$ for all $i = 1, \ldots, n$.

Note that though the performance metric in this example results in a cost function similar to that of a traditional MPC's, it is not generally limited to this form but rather by the assumptions stipulated in Section 3.2. Figures 3.10 and 3.11 show that the convergence rate is much faster than in Example 3.2. This behaviour is expected as we now assume that we can sample the performance metric over the prediction horizon and thus have access to additional information about the system. Solving this type of problem using ESC could be particularly helpful for uncertain performance metrics or stochastic systems where a definite mathematical formulation may not be readily available.

3.5 Vapour Compression System

In the following section we apply the proposed integrated ESC and MPC method to a vapour compression system; specifically, we apply it to an air-conditioner in cooling mode. A basic overview of a VCS was provided in Section 2.4 but is recapped here.

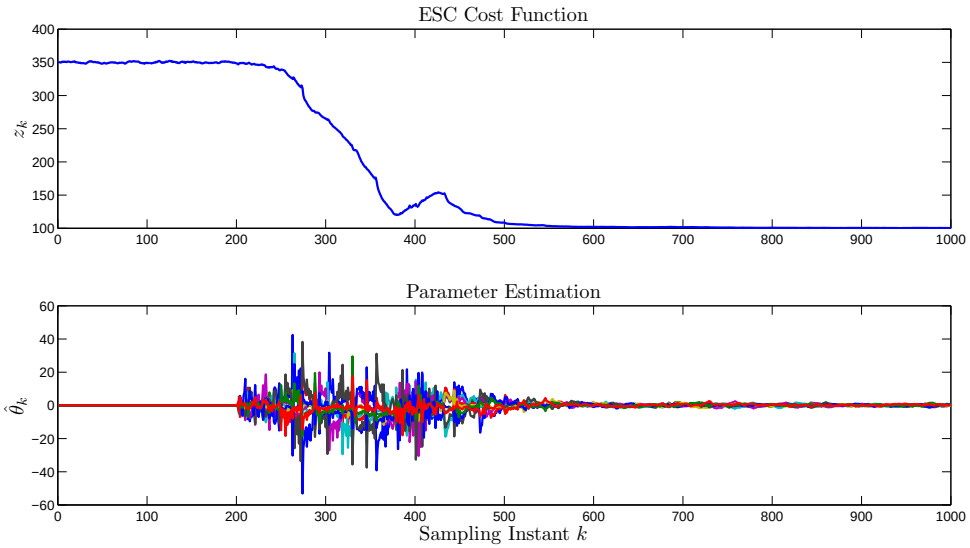

FIGURE 3.10: Above, the ESC cost function and the estimation parameters for Example 3.3 are shown. The ESC routine is started at $t = 200\ s$. The extremum seeking controller is steering the plant so that the performance metric is minimized over the prediction horizon.

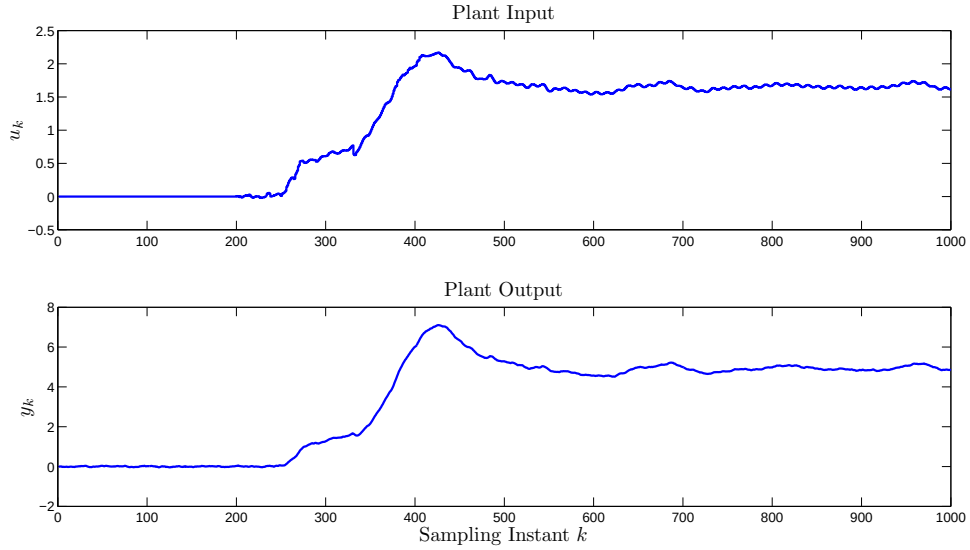

FIGURE 3.11: Plant input and output for Example 3.3. Note the plant converges to the optimal set-point much faster than in Example 3.2. This is expected as we assume knowledge of the cost function.

As shown in Figure 3.12, the air-conditioner consists of a variable speed compressor, condensing heat exchanger, electronic expansion valve or linear expansion valve (EEV and LEV respectively), and evaporating heat exchanger. The (ideal) refrigeration cycle can be described as follows:

1. The refrigerant enters the compressor as a low pressure vapour, where it is compressed and pumped to the condenser.
2. The now high pressure vapour enters the condenser, where it condenses to a liquid while releasing its heat to the outdoors.
3. The liquid now moves through the expansion valve, where its pressure rapidly drops causing flash evaporation and cooling.
4. The refrigerant, a mixture of liquid and vapour at this point, then moves through the evaporator, where it absorbs the heat from the air (which is directed through the evaporator coils by an indoor fan) and changes to a completely gaseous phase.
5. Finally, the hot low-pressure vapour is fed to the compressor, where the cycle is repeated.

The inputs and outputs of the vapour compression system are as follows:

Inputs:

- compressor frequency (Hz),
- condenser fan speed (rpm),
- expansion valve aperture (counts),
- evaporator fan speed (rpm).

Outputs:

- y_1 discharge temp. (°C),
- y_2 discharge temp. superheat (°C),
- y_3 evaporator saturation temp. (°C),
- y_4 condenser subcooling (°C),
- y_5 room air temp. (°C),
- y_6 evaporator superheat (°C).

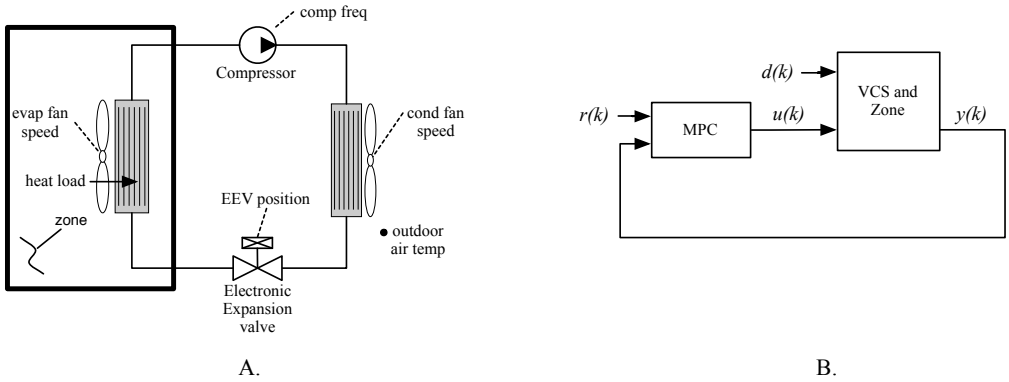

FIGURE 3.12: A. The vapour compression system consists of a variable speed compressor, condensing heat exchanger, electronically controlled expansion valve, and evaporating heat exchanger. B. An MPC controller is nominally configured to use measurements $y(k)$ to drive regulated variables of a vapour compression system and zone to their set-points $r(k)$ in the presence of disturbances $d(k)$ such as changes in outdoor air temperature and heat load.

We are particularly interested in room air temperature, compressor discharge temperature, and evaporator superheat. Room air temperature is of importance as the MPC should track the user-defined set-point for this signal. Evaporator superheat, i.e. the temperature of the refrigerant at the evaporator outlet minus its saturation temperature, has typically been used in order to optimize cycle efficiency. However, calculating evaporator superheat requires at least a pressure and temperature sensor at the evaporator outlet–sensors that are not typically available in commercial systems due to the associated cost. Compressor discharge temperature, i.e. the temperature of the refrigerant as it exits the compressor, is a commonly available measurement and has been suggested as an alternative signal to regulate power consumption [Weiss et al., 2014].

TV-ESC is used to obtain energy optimal set-points for compressor discharge temperature, which are then tracked by the MPC. This approach is well suited to this application since a reliable model relating power consumption and discharge temperature is not available.

3.5.1 Model Predictive Controller

The MPC is controlling the nonlinear vapour compression system using a linearised model. In compact form, this model is given by

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + B_d w_k \\ y_k &= Cx_k + Du_k \end{aligned} \tag{3.42}$$

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. The linearised model was obtained by perturbing the nonlinear system about its nominal operating conditions using binary pseudo-random numbers to obtain single-input single-output (SISO) step responses. Transfer functions were fitted to each of the open-loop step responses, and were then combined into a single transfer function matrix with 56 states. Finally, a reduced order model was obtained by applying an optimal Hankel norm approximation, resulting in a 6th order linear model.

Let $T_{r_{sp}} \in \mathbb{R}$ represent the set-point 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 and zone (room) temperature respectively. We can then define the performance metric $v(k)$ as follows:

$$v(k) = \begin{pmatrix} y_{T_d}(k) - r(k) \\ y_{T_r}(k) - T_{r_{sp}} \end{pmatrix} \tag{3.43}$$

The MPC optimization problem is then formulated as

$$\begin{aligned}
\min_{u(0),\ldots,u(k)} \quad & x(N)^T P x(N) + \sum_{k=0}^{N-1} v(k)^T Q v(k) + \Delta u(k)^T R \Delta u(k) \\
\text{s.t.} \quad & y_{min} \leq y(k) \leq y_{max}, && k = 0, \ldots, N \\
& u_{min} \leq u(k) \leq u_{max}, && k = 0, \ldots, N \\
& x_{k+1} = Ax_k + Bu_k + B_d w_k, && k = 0, \ldots, N-1 \\
& y_k \quad = Cx_k + Du_k, && k = 0, \ldots, N-1.
\end{aligned} \tag{3.44}$$

where N is the MPC prediction horizon and $\Delta u(k) = u(k) - u(k-1)$, i.e. the change in the control input signal relative to the previous time step. 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 positive definite matrices $Q \in \mathbb{R}^{n_v \times n_v}$ and $R \in \mathbb{R}^{p \times p}$ are tuned to appropriately weight the objectives in the cost function (where n_v is the size of the vector v). Further, Q and R are diagonal matrices that penalize output regulation error and control effort respectively. The states of the nonlinear model describe the properties of the refrigerant as it traverses the VCS. However, the states of the linear model do not have a physical meaning and can thus not be measured. To achieve full-state feedback, which is necessary for the implementation of the MPC, we use a Luenberger observer which relies on a similar linear model.

3.5.2 Simulation Results

The extremum seeking control algorithm is used to determine the optimal discharge temperature set-point that minimizes power consumption at steady-state for the air conditioning system. The set-points for compressor discharge temperature r_k determined by the ESC are passed to the MPC which ensures that discharge and zone temperature set-points are met, while maintaining operating constraints. As previously illustrated in Figure 3.1, the

inputs u_k to the VCS are compressor frequency, EEV position, and condenser fan speed (evaporator fan speed was kept constant for these simulations). Heat load changes are modelled as a disturbance w_k. The performance metric z_k is power consumption measured in Watts. Simulations are performed on a model of vapour compression systems that has been developed based on the Thermosys toolbox for MATLAB/Simulink [Alleyne Research Group, 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. For the following simulations, condenser fan speed was kept at its nominal value for clarity purposes so that the ESC trajectory could align with the cost function surface. However, successful simulations have been obtained when fan speed was allowed to be manipulated by the MPC, as is the case in practice.

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

$$k_g = 0.06,$$
$$\alpha = 0.05,$$
$$\epsilon = 0.8,$$
$$A = 0.2,$$
$$\omega = 20k,$$

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.

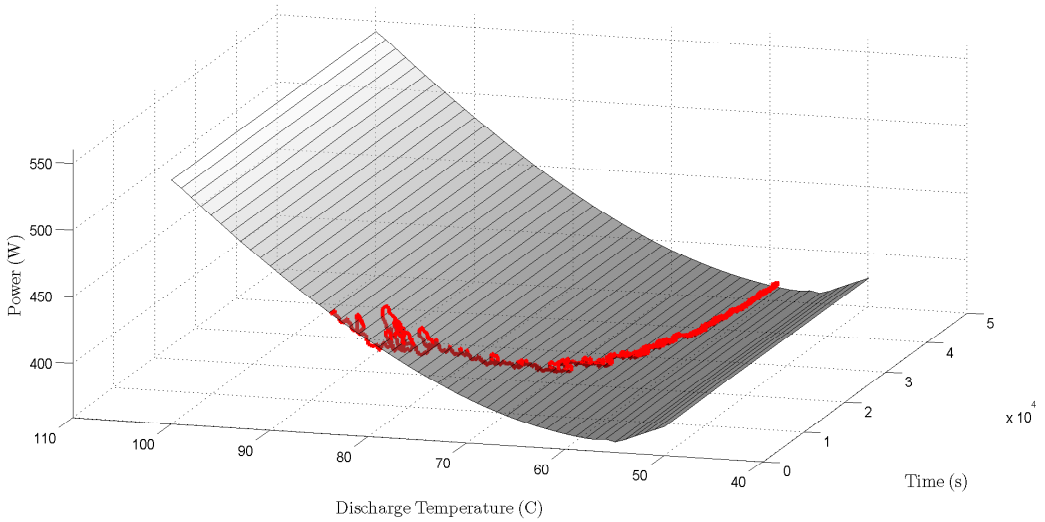

FIGURE 3.13: The trajectory of the ESC is shown to go near the minimum of the power surface while maintaining a constant zone temperature for a heat load of 1900 W. The extremum seeking controller begins at a discharge temperature of 83.8 °C and power consumption of 449.6 W, and converges to 59.3 °C and 372.5 W respectively.

The surface shown in Figure 3.13 describes the function relating power consumption to discharge temperature at steady-state. This cost surface was obtained from simulation results using two PI (proportional-integral) controllers to regulate room and discharge temperature at a given set-point. The steady-state power consumption was then measured and plotted. This process was repeated for different discharge temperature set-points thus resulting in the graph depicted. Outdoor and indoor fan speed was kept constant while collecting this data. As a result, the simulations involving the ESC and MPC also have to be conducted at constant fan speeds in order to align with the cost surface. This is acceptable for the purposes of demonstrating that the proposed method is working as intended, but in practice we would allow the MPC to alter outdoor and indoor fan speed.

The convex shape of the function indicates that the ESC can minimize power consumption at steady state by finding the optimal set-point for compressor 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 set-point) 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.

Figures 3.14, 3.15, and 3.16 show that the MPC objectives and constraints are being met during the optimization of the system.

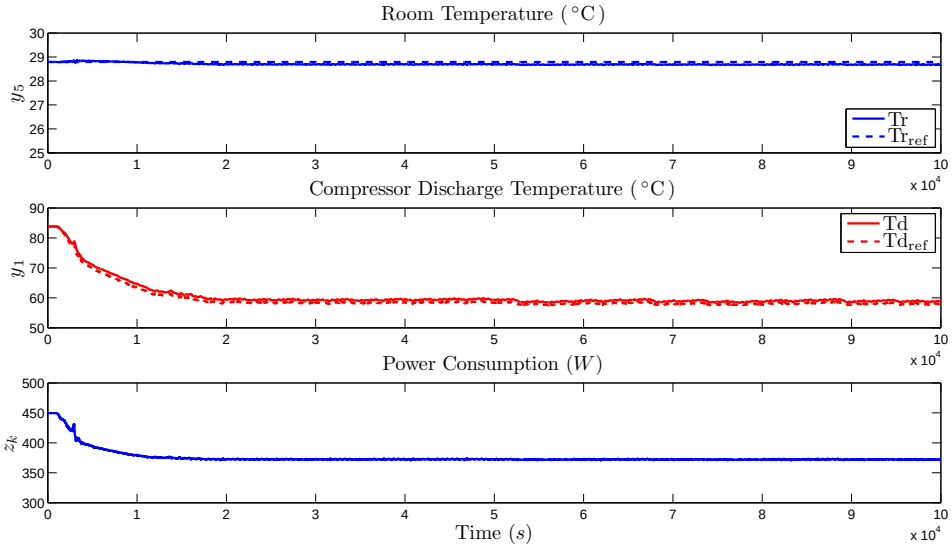

FIGURE 3.14: The MPC ensures that the performance variables, room and compressor discharge temperature, follow the given set-points. Power consumption is minimized as the ESC optimizes the discharge temperature set-points.

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 seeking controller appropriately changes the discharge temperature set-point as shown in Figure 3.17 to obtain

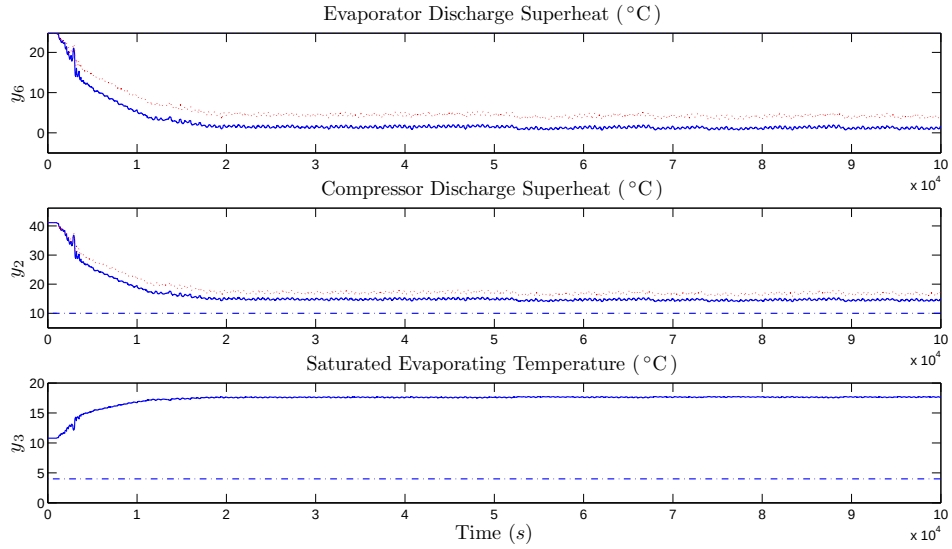

FIGURE 3.15: Trajectories shown in solid blue and dotted red correspond to the outputs as computed by the observer and Thermosys respectively. Note that the outputs are kept within their constraints shown by the dashed blue line.

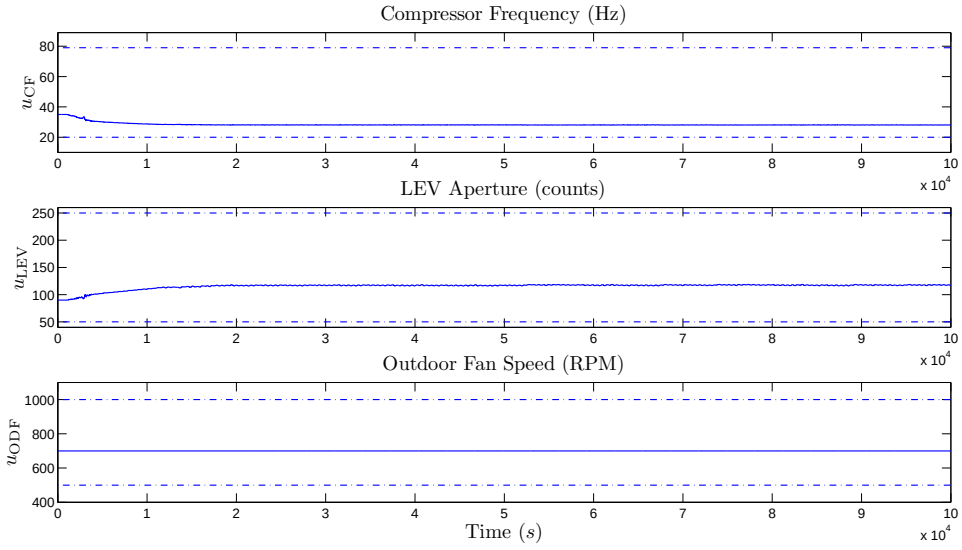

FIGURE 3.16: As the LEV opens to increase cooling capacity, the compressor can reduce its frequency, which lowers power consumption. The inputs are kept within their constraints due to the MPC.

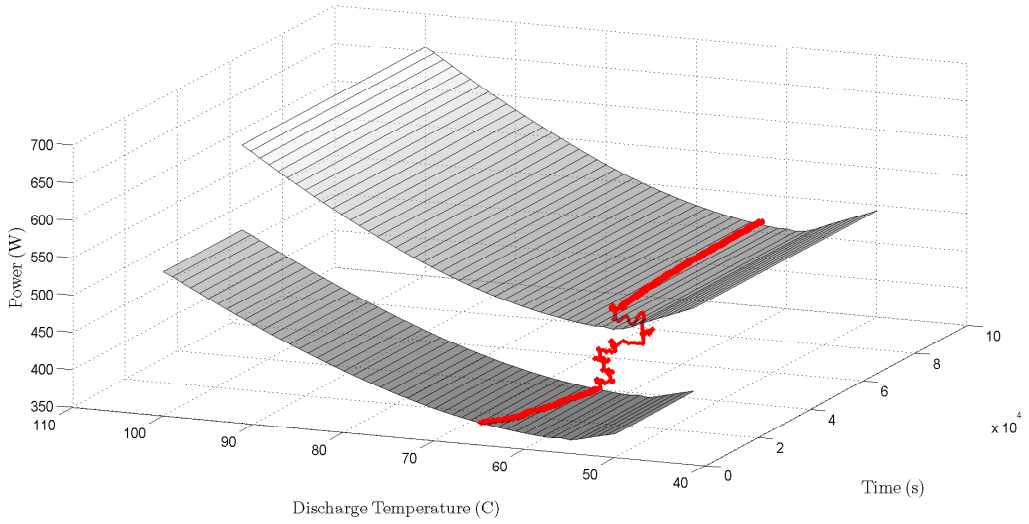

FIGURE 3.17: A heat load increase of 400 W is modelled at 30 000 s. As the cost surface changes due to the new heat load of 2300 W, the time-varying extremum seeking algorithm converges to 62.5 °C and 463.6 W, a region near the true optimum at 58 °C and 456 W.

the minimum energy operating point. The trajectory shown begins where the previous simulation ended, as the extremum seeking controller reaches the near-optimum discharge temperature for a heat load of 1900 W. The ESC 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 ESC converges to can be adjusted by varying the optimization gain and dither signal amplitude.

For this application, time-varying extremum seeking converges and allows disturbances to be rejected on time scales necessary for online implementation. As shown in the simulations above, TV-ESC convergence for the VCS can take multiple hours. As a result, other ESC methods with slower convergence properties would likely not be viable in practice for this application.

3.6 Summary

In this chapter, we show that time-varying extremum seeking control is a viable and effective approach to steady-state real-time optimization of vapour compression systems. TV-ESC does not rely on averaging and singular perturbation to obtain an estimate of the gradient of the power consumption. As a result, it converges to optimal set-points substantially faster than conventional extremum seeking methods.

This approach does not rely on dither signal frequency as the sole tuning parameter, leading to more freedom in tuning and thus making the tuning process easier. The rate at which the ESC is sampled directly impacts the convergence time to the optimum set-point. This rate cannot be set arbitrarily fast, however, as the dynamics of the system and the speed at which the MPC is able to achieve new set-points have to be considered. Further, previous work has shown that the energy efficiency of vapour compression systems is strongly dependent on set-points for superheat temperature, though these set-points have not always been easy to determine. Here, we showed that discharge temperature set-points can be generated via a model-free ESC approach and that the integration of ESC and MPC resulted in substantial energy efficiency improvements while also maintaining operating constraints. Future work will focus on obtaining experimental results for this approach.

Chapter 4

Conclusion

In this thesis, time-varying extremum seeking control was shown to be a viable and effective approach to steady-state real-time optimization of vapour compression systems. The proposed approach could be extended to other systems that use MPC and require optimization, e.g. petro-chemical systems that operate within set constraints and require optimization to maximize yield. Real-time optimization of processes using MPC is of great interest, because of the inherent capacity of MPC to manage a wide array of processes and constraints. As accurate models become harder to develop for complex systems, or, perhaps, are computationally inefficient to use, it is beneficial to have a non-model based solution for the optimization of these systems in lieu of the model predictive controller.

TV-ESC does not rely on averaging and singular perturbation to obtain an estimate of the gradient of the power consumption. As a result, it converges to optimal set-points substantially faster than conventional extremum seeking methods. Since this approach does not rely on dither signal frequency as the sole time-scale separation parameter, there is more freedom in tuning and thus the tuning process is generally easier. The rate at which the ESC is sampled directly impacts the convergence time to the optimum set-point. This rate cannot be set arbitrarily fast, however, as the dynamics of the system and the speed at which

the MPC is able to achieve new set-points have to be considered. Further, previous work has shown that the energy efficiency of vapour compression systems is strongly dependent on set-points for superheat temperature, though these set-points have not always been easy to determine. Here, we showed that discharge temperature set-points can be generated via a model-free ESC approach and that the integration of ESC and MPC resulted in substantial energy efficiency improvements while also maintaining operating constraints.

Future research will focus on obtaining experimental results for the approach proposed in this thesis. It is also of interest to investigate the stability and robustness properties of the integration of time varying extremum seeking and MPC. Further, the non-local stability properties of TV-ESC for discrete time systems are of significant interest. Also, further exploring the class of stability that can be achieved when emulating MPC using ESC would be beneficial. Finally, it would also be of interest to look into other potential uses of ESC on MPC controlled systems, such as optimizing the tuning of the model predictive controller for improved performance.

Appendix A

Auxiliary Results

Figures A.1, A.2, A.3, and A.4 show the results for the ESC-MPC routine applied to the VCS at high load, i.e. 2300 W. Here, the outdoor fan was not restricted to its nominal operating value. The convergence rate could be improved by using a higher gain for the ESC.

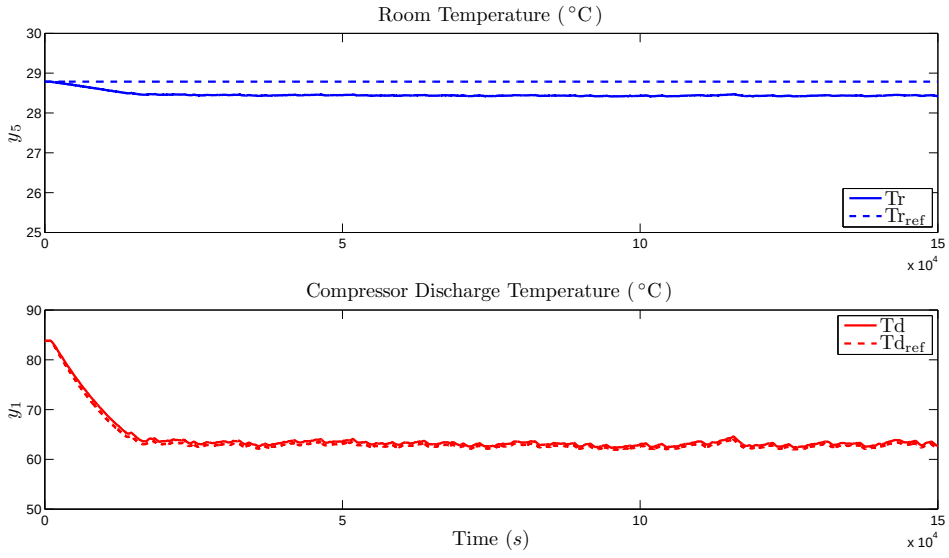

FIGURE A.1: The MPC ensures that the performance variables follow the given setpoints.

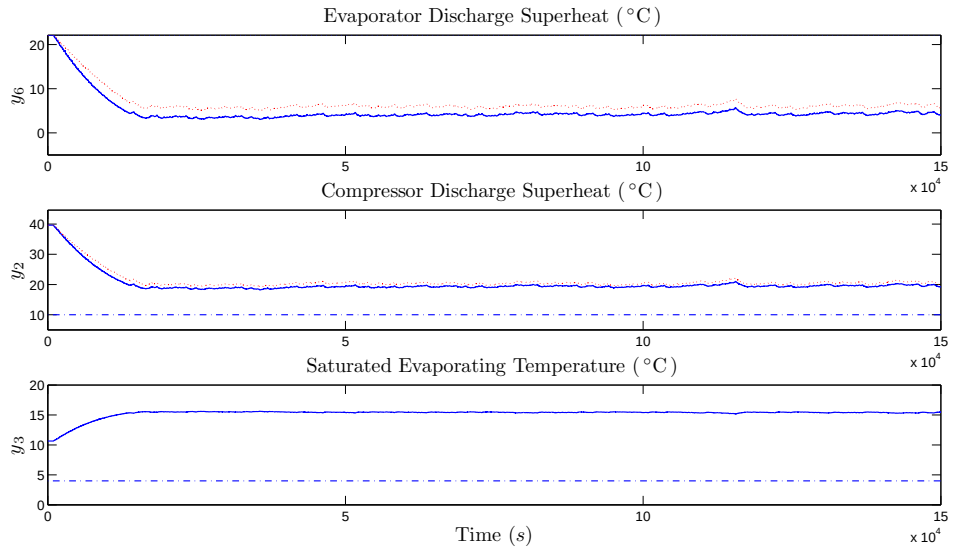

FIGURE A.2: Trajectories shown in solid blue and dotted red correspond to the outputs as computed by the observer and Thermosys respectively. Note that the outputs are kept within their constraints shown by the dashed blue line.

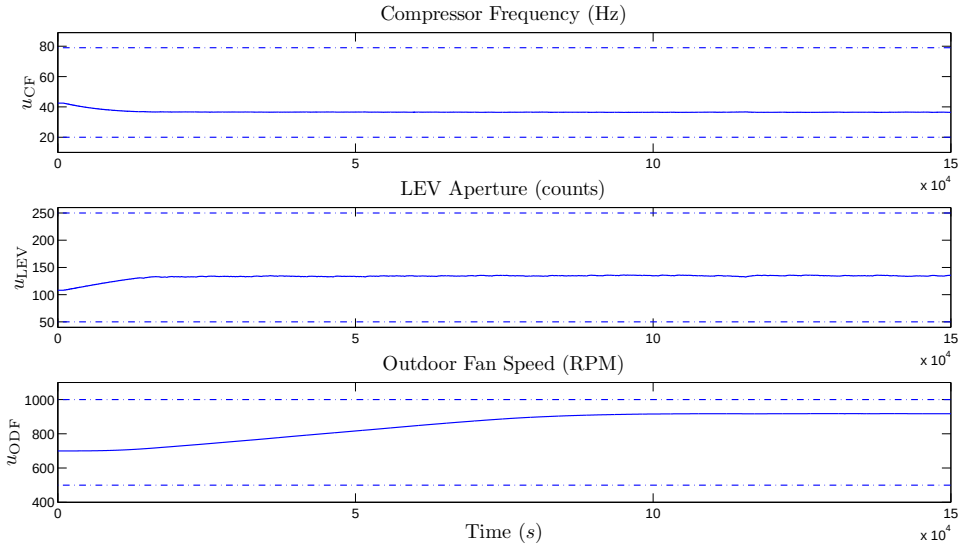

FIGURE A.3: As the LEV opens to increase cooling capacity, the compressor can reduce its frequency, which lowers power consumption. The inputs are kept within their constraints due to the MPC.

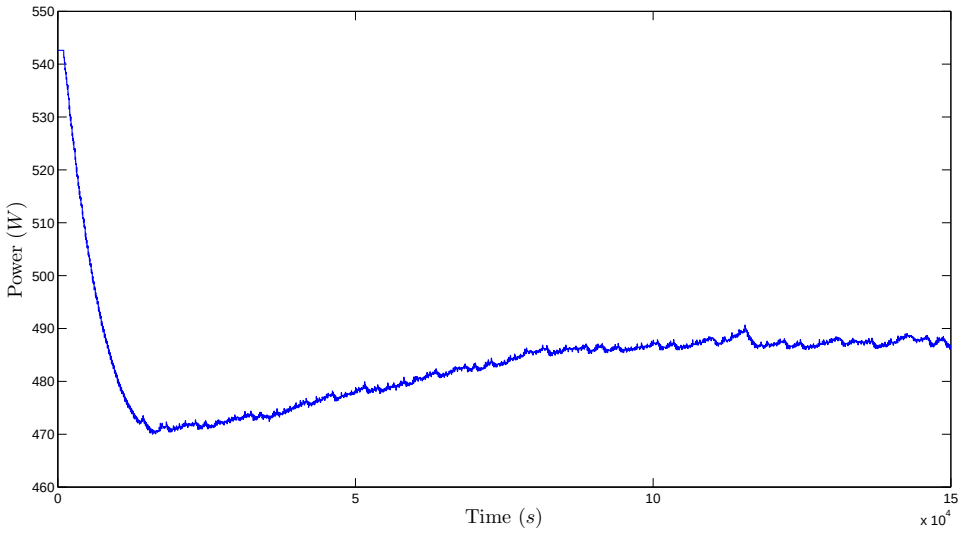

FIGURE A.4: Power consumption using the hierarchical ESC-MPC routine for a heat load of 2300 W.

Bibliography

E F Camacho and C Bordons. *Model Predictive Control*. Advanced Textbooks in Control and Signal Processing. Springer London, 2004. ISBN 9781852336943.

A.N. Venkat, I.A. Hiskens, J.B. Rawlings, and S.J. Wright. Distributed MPC Strategies With Application to Power System Automatic Generation Control. *IEEE Transactions on Control Systems Technology*, 16(6):1192–1206, November 2008. ISSN 1063-6536. doi: 10.1109/TCST.2008.919414.

Jing Zeng, Ding-Yu Xue, and De-cheng Yuan. Research and Development Trend of Distributed MPC. In *2008 Fourth International Conference on Natural Computation*, volume 4, pages 417–421. IEEE, 2008. ISBN 978-0-7695-3304-9. doi: 10.1109/ICNC.2008.846.

Alberto Bemporad and Manfred Morari. Robust Model Predictive Control: A Survey. In A Garulli and A Tesi, editors, *Robustness in identification and control SE - 16*, volume 245 of *Lecture Notes in Control and Information Sciences*, pages 207–226. Springer London, 1999. ISBN 978-1-85233-179-5. doi: 10.1007/BFb0109870.

G C Goodwin, S F Graebe, and M E Salgado. *Control system design*. Prentice Hall, 2001. ISBN 9780139586538.

D Q Mayne, J B Rawlings, C V Rao, and P O M Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, June 2000. ISSN 0005-1098. doi: http://dx.doi.org/10.1016/S0005-1098(99)00214-9.

A Ferramosca, D Limon, I Alvarado, T Alamo, and E F Camacho. MPC for tracking with optimal closed-loop performance. *Automatica*, 45(8):1975–1978, August 2009. ISSN 0005-1098. doi: http://dx.doi.org/10.1016/j.automatica.2009.04.007.

M Leblanc. Sur l'électrification des chemins de fer au moyen de courants alternatifs de fréquence élevée. *Revue générale de l'électricité*, 12(8):275–277, 1922.

K B Ariyur and M Krstic. *Real-Time Optimization by Extremum-Seeking Control*. Wiley-interscience publication. Wiley, 2003. ISBN 9780471468592.

Miroslav Krstić and Hsin-Hsiung Wang. Stability of extremum seeking feedback for general nonlinear dynamic systems. *Automatica*, 36(4):595–601, 2000.

Ying Tan, Dragan Nešić, and Iven Mareels. On non-local stability properties of extremum seeking control. *Automatica*, 42(6):889–903, 2006.

Ying Tan, Dragan Nešić, and Iven Mareels. On the choice of dither in extremum seeking systems: A case study. *Automatica*, 44(5):1446–1450, 2008.

M Guay, S Dhaliwal, and D Dochain. A time-varying extremum-seeking control approach. In *American Control Conference (ACC), 2013*, pages 2643–2648, 2013.

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

M Guay and T Zhang. Adaptive extremum seeking control of nonlinear dynamic systems with parametric uncertainties. *Automatica*, 39(7):1283–1293, July 2003. ISSN 0005-1098. doi: http://dx.doi.org/10.1016/S0005-1098(03)00105-5.

Darryl DeHaan and Martin Guay. Extremum-seeking control of state-constrained nonlinear systems. *Automatica*, 41(9):1567–1574, September 2005. ISSN 0005-1098. doi: http://dx.doi.org/10.1016/j.automatica.2005.03.030.

V Adetola and M Guay. Parameter convergence in adaptive extremum-seeking control. *Automatica*, 43(1):105–110, January 2007. ISSN 0005-1098. doi: http://dx.doi.org/10.1016/j.automatica.2006.07.021.

Y. Tan, Dragan Nešić, I.M.Y. Mareels, and A. Astolfi. On global extremum seeking in the presence of local extrema. *Automatica*, 45(1):245–251, 2009.

Glauce De Souza, Darci Odloak, and Antônio C. Zanin. Real time optimization (RTO) with model predictive control (MPC). *Computers & Chemical Engineering*, 34(12):1999–2006, December 2010. ISSN 00981354. doi: 10.1016/j.compchemeng.2010.07.001.

V Adetola and M Guay. Adaptive output feedback extremum seeking receding horizon control of linear systems. *Journal of Process Control*, 16(5):521–533, June 2006. ISSN 0959-1524. doi: http://dx.doi.org/10.1016/j.jprocont.2005.07.004.

V Adetola and M Guay. Integration of real-time optimization and model predictive control. *Journal of Process Control*, 20(2):125–133, February 2010. ISSN 0959-1524. doi: http://dx.doi.org/10.1016/j.jprocont.2009.09.001.

Paul Frihauf, Miroslav Krstic, and Tamer Başar. Finite-horizon LQ control for unknown discrete-time linear systems via extremum seeking. *European Journal of Control*, 19(5):399–407, September 2013. ISSN 0947-3580. doi: http://dx.doi.org/10.1016/j.ejcon.2013.05.015.

D Burns and C Laughman. Extremum Seeking Control for Energy Optimization of Vapor Compression Systems. In *International Refrigeration and Air Conditioning Conference*, 2012.

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

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

Yudong Ma, F Borrelli, B Hencey, B Coffey, S Bengea, and P Haves. Model Predictive Control for the Operation of Building Cooling Systems. *Control Systems Technology, IEEE Transactions on*, 20(3):796–803, May 2012. ISSN 1063-6536. doi: 10.1109/TCST.2011.2124461.

H S Sane, C Haugstetter, and S A Bortoff. Building HVAC control systems - role of controls and optimization. In *American Control Conference, 2006*, pages 6 pp.–, June 2006. doi: 10.1109/ACC.2006.1656367.

Pengfei Li, Yaoyu Li, and John E Seem. Efficient Operation of Air-Side Economizer Using Extremum Seeking Control. *Journal of Dynamic Systems, Measurement, and Control*, 132(3):31009, April 2010. ISSN 0022-0434. doi: 10.1115/1.4001216.

V. Tyagi, H. Sane, and S. Darbha. An extremum seeking algorithm for determining the set point temperature for condensed water in a cooling tower. In *2006 American Control Conference*, page 5 pp. IEEE, 2006. ISBN 1-4244-0209-3. doi: 10.1109/ACC.2006.1656368.

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

Martin Guay, Ruud Beerens, and Henk Nijmeijer. A time-varying extremum-seeking control approach for discrete-time systems with application to model predictive control. In *IFAC*, 2014.

Martin Guay. A time-varying extremum-seeking control approach for discrete-time systems. *Journal of Process Control*, 24(3):98–112, March 2014. ISSN 09591524. doi: 10.1016/j.jprocont.2013.11.014.

G C Goodwin and K S Sin. *Adaptive Filtering Prediction and Control*. Dover Books on Electrical Engineering Series. Dover Publications, Incorporated, 2009. ISBN 9780486137728.

Walter Weiss, Daniel Burns, and Martin Guay. Realtime Optimization of MPC Setpoints using Time-Varying Extremum Seeking Control for Vapor Compression Machines. In *15th International Refrigeration and Air Conditioning Conference at Purdue*, 2014.

James B. Rawlings, David Angeli, and Cuyler N. Bates. Fundamentals of economic model predictive control. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 3851–3861. IEEE, December 2012. ISBN 978-1-4673-2066-5. doi: 10.1109/CDC.2012.6425822.

Alleyne Research Group. THERMOSYS 4 Toolbox, 2012. URL http://arg.mechse.illinois.edu/thermosys.