3-21-2019

Control Strategies for Multi-Evaporator Vapor Compression Cycles

Sunderlin D. Jackson

CONTROL STRATEGIES FOR MULTI-EVAPORATOR VAPOR COMPRESSION CYCLES

THESIS

Sunderlin D. Jackson, Second Lieutenant, USAF

AFIT-ENY-MS-19-M-221

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense or the United States Government. This is an academic work and should not be used to imply or infer actual mission capability or limitations.

AFIT-ENY-MS-19-M-221

CONTROL STRATEGIES FOR MULTI-EVAPORATOR VAPOR COMPRESSION CYCLES

THESIS

Presented to the Faculty

Department of Aeronautics and Astronautics

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

in Partial Fulfillment of the Requirements for the

Degree of Master of Science in Astronautical Engineering

Sunderlin D. Jackson, BS

Second Lieutenant, USAF

March 2019

AFIT-ENY-MS-19-M-221

CONTROL STRATEGIES FOR MULTI-EVAPORATOR VAPOR COMPRESSION CYCLES

Sunderlin D. Jackson, BS
Second Lieutenant, USAF

Approved:

Anthony Palazotto, PhD (Chairman)	Date
Meir Pachter, PhD (Member)	Date
Nicholas Niedbalski, PhD (Member)	Date

Abstract

Next-generation military aircraft must be able to handle highly transient thermal loads that exceed the ability of current aircraft thermal subsystems. Vapor compression cycle systems are a particular refrigeration technology that is an attractive solution for dealing with this challenge, due primarily to their high efficiency. However, there are several barriers to realizing the benefits of vapor cycles systems for controlling thermal loads in military aircraft. This thesis focuses on addressing the challenge of controlling vapor cycles in the presence of highly transient evaporator heat loads. Specifically, a linear quadratic regulator (LQR) is designed for a simple vapor cycle system, and closed-loop performance is compared with a set of proportional-integral (PI) controllers. Simulation results show significant advantages of using the LQR method, and the same approach is repeated for a larger dual-evaporator vapor cycle system. The LQR method retains some of its benefits, but several issues associated with relying on a single linear model for the full nonlinear system are identified, and recommendations for future work are made at the end.

Acknowledgements

First, I would like to thank Dr. Anthony Palazotto for serving as my advisor over the past year. Your mentorship has fostered a passion within me for academic research. Thank you Dr. Meir Pachter, for sharing your many years of experience in controls engineering, and for providing many valuable insights as a member of my advisory committee.

I'd also like to extend thanks to Dr. Brandon Hencey at AFRL for answering many of my questions about troubleshooting control challenges in ATTMO, as well as Dr. John Wen and his team at Rensselaer Polytechnic Institute for sharing some of their experimental and simulation data with me, which helped me get my work off the ground. Last, but certainly not least, thank you Dr. Nicholas Niedbalski for serving as a member of my advisory committee. You took the time to help me find a research area that I was truly passionate about, and for that I am truly grateful.

Sunderlin D. Jackson

Table of Contents

Page

List of Figures

List of Tables

List of Acronyms

AFRL Air Force Research Laboratory

VCSRF Vapor Cycle System Research Facility

ATTMO AFRL Transient Thermal Modeling and Optimization

VCS Vapor Compression System

SST Saturated Suction Temperature

SDT Saturated Discharge Temperature

HVAC Heating, Ventilation and Air Conditioning

SISO Single-Input Single-Output

MIMO Multi-Input Multi-Output

PI Proportional-Integral

LQR Linear Quadratic Regulator

LQE Linear Quadratic Estimator

MPC Model Predictive Control

PAO Polyalphaolefin

CONTROL STRATEGIES FOR MULTI-EVAPORATOR VAPOR COMPRESSION CYCLES

I. Introduction

1.1 Background

The various mechanical and electrical systems present in modern aircraft have a tendency to generate heat. Left alone, the buildup of this heat would degrade performance, overheat components, and eventually lead to system failure. It is for this reason that all aircraft have some form of thermal subsystem, to regulate temperatures throughout the vehicle. Currently, most aircraft rely upon passive air cycle cooling to maintain thermal requirements. Unfortunately for the aircraft designer, it is anticipated that the design requirements for next-generation military aircraft will require the use of more advanced thermal control systems. Specifically, next-generation aircraft will need to handle highly transient thermal loads while meeting tighter mass, size and efficiency constraints [2]. In order to realize these goals, research is being conducted to evaluate the effectiveness of going to active thermal control, with different cooling architectures.

The Vapor Compression System (VCS) is an attractive solution to this challenge because they take advantage of evaporative cooling, and they incur less weight, cost and volume, compared to appropriately scaled versions of traditional thermal management systems [3]. The standard vapor cycle consists of four components: the valve, evaporator, compressor and condenser. The cycle is

visualized in terms of its components in Fig. 1, but it can also be visualized in terms of its pressure-enthalpy diagram, shown in Fig. 2.

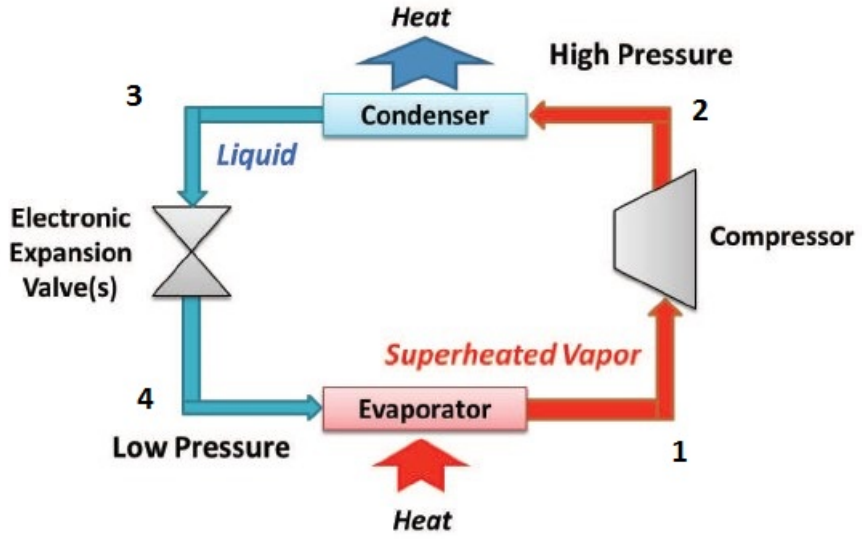

Figure 1. Simple Vapor Cycle Architecture

A refrigerant, such as R134a, flows throughout the cycle, carrying heat from a load to a sink. The fluid begins the cycle as a low pressure gas at the exit of the evaporator (1). It is then compressed to a high pressure (1-2) and passed through the condenser where it is cooled to a liquid (2-3). From here, it flows through the valve (3-4), expanding the fluid and bringing it to a two-phase state, at low pressure. Heat is then reintroduced to the fluid while passing through the evaporator (4-1), and the cycle repeats. Idealizing assumptions include isobaric phase change across the heat exchangers, isenthalpic expansion in the valve, and isentropic compression in the compressor.

VCS has a long and successful history of application in the Heating, Ventilation and Air Conditioning (HVAC) and automotive industries, among others. Phase change between liquid and gas requires a large amount of heat for a relatively small amount of refrigerant, thereby allowing the vapor cycle to achieve higher efficiencies

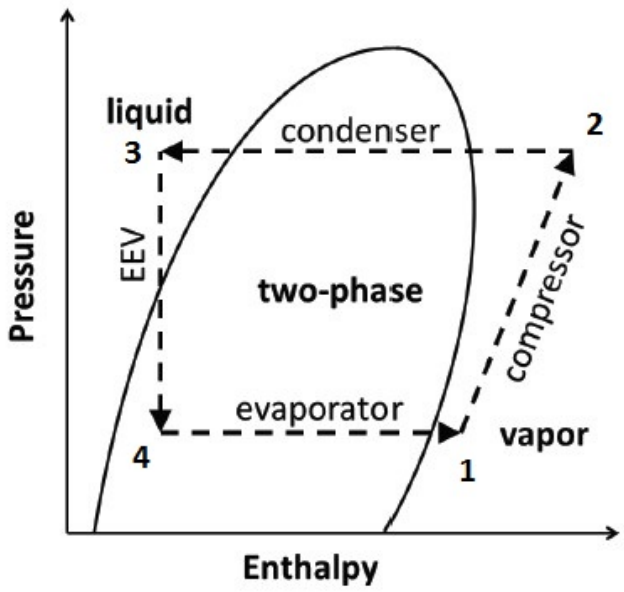

Figure 2. Pressure Enthalpy Diagram

than alternatives like the air cycle. Additionally, there is great potential in controlling the system to handle highly variable loads on the evaporator, so that cooling capacity is delivered as demand arises. However, this technology has some hurdles to overcome before it can be used effectively to address the aforementioned requirements for next-generation military aircraft.

Traditional applications for VCS have been characterized by large thermal inertias, relatively stagnant ambient conditions, and steady heat loads. Thus, there are well-defined operation and control methods for VCS operating under these conditions. However, military aircraft do not satisfy any of those conditions. High power electronics can double the evaporator heat load in a matter of seconds. Ambient temperatures can fall 20° C as the aircraft rapidly climbs in altitude. These issues might be overcome by over-sizing the VCS components, but then size constraints are violated. Additionally, there may be constraints imposed by stealth on the allowable heat that can be sunk to ambient. All of these characteristics make it harder to control the vapor cycle and maintain required operating temperatures

at the evaporator load. There are even cases where these issues can cause component failure, with the most well-known case being superheat loss [4]. Designing VCS capable of dealing with these instances of high transience, while still meeting design constraints is an ongoing area of research. Table 1, adapted from Doty et al. [1], provides a concise summary of the fundamental research challenges in aircraft thermal controls.

Control Limitation	**Aircraft thermal system example**
System and signal uncertainty	Aging aircraft systems changing dynamic behavior over time
Time delays	Transport delays of energy or mass for a liquid loop to affect a cold plate
Sensing/actuation limits	Maximum pump speed or valve opening in a liquid loop
Collocation of sensing and actuation	Temperature sensors in an electronics bay with cooling governed by an ACM system actuated remotely from the sensor
Underactuation	Spatially distributed fuel temperature profiles with point sensing and actuation
Nonlinearity	Variation in ACM effectiveness as a function of air-vehicle altitude

Table 1. Control limitations and their prevalence in aircraft thermal systems[1]

1.2 Problem Statement

This thesis will focus on the problem of controlling the actuators in a vapor cycle. By varying the valve opening and the compressor speed, the controller seeks to drive the system to a particular state. In most applications, this is handled by separate Single-Input Single-Output (SISO) control loops, each driven by separate PI controllers. This simple approach succeeds when the plant has a large thermal mass and the loads are slowly varying, but performance suffers greatly in the presence of highly transient loads, even becoming unstable in some cases [5]. The feedback gains in each PI loop are typically tuned heuristically, based on prior

experience, but tuning around performance objectives is often a process of trial-and-error, with no metric on whether or not success is even possible. Additionally, this process is exacerbated as additional evaporators are added to the system to cool distributed loads with different operating temperatures.

1.3 Research Objectives

This thesis aims to evaluate the benefits of applying model-based control to a VCS, as compared to a baseline PI controller. Academic research into VCS control has shown the benefits of applying more advanced control strategies, particularly in the field of automotive and building HVAC systems ([6],[7],[8]). Extending these techniques to applications with highly transient loads has also shown significant improvements in controller performance [9]. However, the prospect of iterating rapidly between controller and plant design is challenging, due to the difficulties of accurately modeling the complex, nonlinear dynamics present in vapor cycles. The use of the AFRL Transient Thermal Modeling and Optimization (ATTMO) toolbox can alleviate this modeling challenge, and so developing a systematic process for using the toolbox to design controllers is of great interest.

1.4 Methodology

Ngo et al. [10] established such a pathway for designing model-based controllers in ATTMO. This thesis utilizes the same pathway in order to design a Linear Quadratic Regulator (LQR) controller with the goal of rejecting a pulsed heat disturbance on a VCS plant, modeled in ATTMO. LQR controllers are Multi-Input Multi-Output (MIMO) controllers and therefore are not limited to controlling "pairs" of signals in the way that SISO loops are. Cross-couplings in the dynamics are entirely ignored when using SISO controllers, and manifest as "unmodeled"

dynamics in the closed-loop system, causing a deterioration in performance. MIMO controllers directly consider cross-couplings, and utilize them in the control design.

Our control design process began by defining the actuators, available feedback signals, and desired performance objectives. The process started with a Simple VCS, consisting of only the four necessary components, in order to gain familiarization with the method. An ATTMO model of the system was linearized in MATLAB, resulting in a linear state space model, with as many states as the full nonlinear model. This system was too large to be tractable for control system design, and was reduced in size via model-reduction techniques. With the reduced-order linear model, an LQR controller with a state estimator was designed and tuned on a test case where a large pulse disturbance in the heat load was sent to the system. To reach the desired performance objectives, the LQR controller was modified to include a model of exogenous disturbances in the form of the heat load, so that large increases in the heat can be compensated for without requiring explicit measurement of the heat load.

Once that process was complete and the improved performance was verified, the design was repeated with a Dual-Evaporator VCS, based on a model of the Vapor Cycle System Research Facility (VCSRF). The VCSRF was built by the Air Force Research Laboratory (AFRL) Aerospace Systems Directorate, at Wright Patterson Air Force Base, Ohio, and was designed as a test-bed for evaluating different control and operation methods for VCS. This test-bed has been modified and used for several studies in the past, many of which were aimed at evaluating new methods for controlling vapor cycles in operating conditions representative of those seen by military aircraft. Indeed, the baseline PI controllers against which we compare our results are taken from this past work [2]. By introducing the added complexity of

some of the VCSRF components, the intent was to show how the controller's performance scales with the complexity of the plant.

1.5 Assumptions and Limitations

Throughout this work, the controllers are designed and evaluated purely in simulation. The expectation that these controllers will work on the physical system is closely tied to an assumption that the models have been sufficiently validated on physical systems. This work does not contain any model validation, but there is already a significant amount of ATTMO model validation existing in the literature. For this reason, it is assumed that the ATTMO models are sufficiently accurate, such that this approach would work on a real VCS. However, in order to be sure of this assumption, further work is required to verify controller performance on the actual hardware.

Additionally, certain assumptions are made about the nature of the feedback signals provided by the sensors for state estimation. The feedback signals are all temperatures, but in physical systems there is often a discrepancy between measured temperatures and desired temperatures, either in terms of a time delay or an offset. No considerations for sensor inaccuracy, or actual sensor noise, are taken in this work, but these issues can be dealt with by appropriately modifying the model to include these phenomena.

Lastly, since we are designing linear controllers, we assume that the operating conditions remain sufficiently close to the setpoint such that the linearized dynamics are an appropriate description of the actual dynamics. There are regions of the state space that are more nonlinear than others (e.g. low superheat), and so judicious choosing of the setpoint is an important part of the design process. In

practice, this can be addressed by using multiple linear controllers that are gain-scheduled to handle the changes in dynamics.

1.6 Thesis Overview

The remainder of this thesis is organized as follows. Chapter II provides a review of existing literature on VCS control, in order to provide the appropriate context for our work. This begins with a review of the research efforts undertaken to develop accurate models of vapor cycles. The study of leveraging different control methods in VCS is then shared, and our particular contribution to the literature is given.

Chapter III discusses the methodology used in our models and in the design of our controllers. In particular, an outline of the linear optimal control theory associated with LQR controllers is given, along with a discussion of how integral control and disturbance estimation is introduced to the design process. The process for extracting a tractable linear model from the full nonlinear system is also outlined in this section, along with a table containing the controller specifications. Towards the end of the chapter, a brief outline of the PI control architecture is shared.

Chapter IV contains the closed-loop simulation results for the Simple VCS with the LQR and PI controllers. Performance comparisons are made, with emphasis on the benefits of using the LQR method. The same process is then repeated with the Dual-Evaporator VCS, and the limitations of the LQR method are highlighted in the simulation results.

Finally, Chapter V summarizes the research effort, and concludes with closing points and recommendations for future work.

II. Literature Review

2.1 Vapor Cycle Modeling

The vast majority of classical control techniques require a model of the plant dynamics in order to design an effective controller. Therefore, a natural starting point for choosing feedback gains to control the compressor and valves in a vapor cycle system would be in modeling the various components. As discussed by Kania et al. in their overview of the ATTMO toolbox [11], the challenge in vapor cycle modeling lies in the two-phase heat exchanger models. There are a number of instabilities that can develop in two-phase flows, both over fast and slow timescales as shown by Kakac in [12], and these issues are compounded by interactions between different subsystems. For example, it is critical in many vapor cycles that superheat be maintained, otherwise liquid can leave the evaporator, causing a mechanical failure in the compressor. Identifying what characteristics to include in the mathematical model of a VCS has always been a challenge, and often control engineers have resorted to model-free control techniques, such as PI, which can be tuned without a model of the dynamics. These approaches work under conservative operating conditions, but there are many performance sacrifices that are made when designing controllers like this.

The first tractable model of heat exchangers was the moving boundary model, developed by Wedekind et. al. [13]. This approach simplifies the two-phase flow by dividing it into discrete zones, with zones for liquid, gas and two-phase. Further research went into refining this modeling approach to improve transient behavior, and indeed continues today [14]. In many academic works today the moving boundary model forms the basis for the plant model, but it is not the only approach. Jain et al. [15] used linear models, analogous to electrical circuits, to

allow for scalability in the size of the system, sacrificing accuracy in exchange for system size. Linear models also see a large amount of use when designing controllers via gain-scheduling, as in Rasmussen et al. [16]. Here, the parameters in the linear models are often identified via system identification. This thesis instead uses a finite-volume approach to modeling the heat exchangers which is more computationally expensive, but has improved accuracy over other approaches [17].

2.2 Vapor Cycle Control

Towards the end of the 20th century, the academic controls community began applying different techniques to the problem of vapor cycle control. The earliest such work was He's dissertation at the Massachusetts Institute of Technology [6]. This work developed an LQR controller based on a linearized model of the moving boundary model. The results showed great performance over traditional control methods in experimental lab work. Later, Dr. He and his advisor Dr. Asada investigated the use of other control methods, including the use of nonlinear control methods like feedback linearization [18]. All of this work was targeted for application in the HVAC industry, although the principals are likely transferable to other applications of VCS. He and Asada published many papers illustrating the benefits of improving VCS control systems, leading other groups to expand the research.

Alleyne and his team at the University of Illinois Urbana-Champaign has a long record of academic contribution to both the modeling and the control side of this problem. The research group has published a number of influential papers on modeling vapor cycles for the control purposes. Often, the team uses linearized dynamics to design controllers, and so they have studied the benefits of choosing optimal setpoints about which to linearize [19]. Additionally, they have studied edge cases encountered during actual operation of vapor cycles, such as start-up and

shutdown transience [20]. On the theoretical side, the team has improved upon existing moving-boundary models to overcome some of the numerical issues associated with them [14], and they have found ways to improve sensor-actuator feedback pairings via the use of relative gain arrays [21].

More recently, Alleyne's team and others have investigated the benefits of applying more modern control techniques to this problem. Several papers applying model predictive control have been published by the team which frame the problem as a constrained optimization problem, where some cost function is formulated to capture the control objectives, and is minimized over the choice of control actions. This approach is often called Model Predictive Control (MPC), and seems to have the most promise when a linear model is used and a priori knowledge about future disturbances is readily available [22]. Of course, the real system is nonlinear, and nonlinear optimization tools are generally less robust, although some research efforts have shown promise in implementing nonlinear MPC in VCS control, with the work of Gräber et al. being a commonly cited example [23].

Wen and his team at the Rensselaer Polytechnic Institute also have a significant record of contributions to the field of vapor cycle control. Like many other teams, they have studied the benefits of applying modern techniques like H_2 [24], and H_∞ [25], usually designed with numerical linearizations of moving-boundary models. However, the team has applied these techniques to modified vapor cycles, with added components like accumulators and recuperators [26]. The accumulator can be used to control the amount of the refrigerant in the system, allowing for large changes in the maximum possible cooling, and the recuperator can be paired with this component to control the quality of the fluid entering the accumulator. These added components introduce additional degrees of control to the system and improve system robustness and efficiency, at the cost of added complexity. Indeed,

the control system design space should include consideration of additional actuators in order to reach desired performance objectives. In recent years, the team has extended their research to evaluate the effectiveness of linear MPC, when knowledge of upcoming transience in the load is available [27], and they have also studied multi-layered control loops to drive the system to more efficient setpoints [26].

2.3 AFRL Vapor Cycle System Research Facility

The AFRL Aerospace Systems Directorate at Wright Patterson Air Force Base has conducted a significant amount of research in aircraft thermal control systems in recent years. The Vapor Cycle System Research (VCSRF) facility was built by AFRL in order to run experimental tests, evaluating different control methods for VCS. Michalak et al. evaluated a cycle-based control strategy against the performance of a traditional superheat and capacity control strategy on the VCSRF system, and showed that significant performance improvements were achievable by using the cycle-based control approach [2]. However, this approach still consists of using a few SISO PI loops, as opposed to a MIMO approach. Work by AFRL has also investigated tip-to-tail aircraft thermal modeling in Simulink, beginning with the Integrated Vehicle and Energy Technology (INVENT) Program in 2008 [28]. This program studied aircraft thermal system performance in conjunction with other aircraft subsystems, and many of the simulation tools formed the early work that later led to the development of what is now the ATTMO simulation environment [29].

The ATTMO toolbox was developed for AFRL, and has been used in a number of simulation studies on VCS control. This toolbox provides a convenient method for modeling vapor cycles, allowing for rapid development of complex models, and modification in the Simulink environment. Additionally, the accuracy of the toolbox

uses a high fidelity control-volume modeling approach in the heat exchangers. With this toolbox, Ngo et. al. established a pathway for designing model-based controllers [10], where they designed a relatively simple LQR controller for a model of the Dual-Evaporator VCSRF facility. Later, three more papers came out that utilized the same pathway for designing various MPC controllers for different VCS architectures [30], solidifying the advantage of using the toolbox for designing controllers. A logical next step for the research effort is in the application of these controllers to actual hardware, to show that even the more advanced techniques, which rely on accurate models, can still be applied to hardware.

2.4 Contribution to Literature

The contribution of this thesis to the literature is centered around the idea of using linear MIMO controllers to improve disturbance rejection in vapor cycles. Experimental work at AFRL has studied the problem of evaporator heat disturbance rejection with PI control loops [2], but the body of academic research discussed in this chapter suggests that performance improvements are achievable with a MIMO control method, like LQR. The LQR method is "linear" in that it is designed to control a linear system, and it is "multi-input multi-output" in that the controller uses multiple actuators to control multiple outputs. This thesis uses the same Simulink pathway suggested by Ngo et. al. [10] to design a controller for a model of a Dual-Evaporator VCS. To the author's knowledge, only four academic papers have utilized this approach, as mentioned above. The ATTMO toolbox has significant value in evaluating different control architectures in simulation before implementation on hardware, and this thesis aims to further explore the benefits of using ATTMO to design controllers for aircraft thermal systems.

III. Methodology

3.1 Chapter Overview

This chapter covers our methodology, which we break into two parts. Section 3.2 covers our approach to modeling vapor cycles. This includes a brief overview of the two vapor cycle architectures, but the majority of this section is devoted to the algorithm for extracting a small linear model from the large nonlinear simulation.

Section 3.3 covers controller design. This focuses mainly on the LQR design steps, along with the extra steps required to augment the controller with integral control (which is not always covered in a discussion of LQR). There are also subsections on PI control and controller specifications at the end of this section.

3.2 Modeling

3.2.1 Nonlinear Models

In this thesis, two different architectures (i.e. "plants") are considered. The first is dubbed the "Simple VCS", which consists of the four basic components required to operate the cycle. This plant was studied first because it is the simplest vapor cycle possible.

The second plant is the Dual-Evaporator Vapor Cycle, and it was used to study the performance of the LQR method on a larger plant, more representative of what might be encountered in a real application. Figure 3 shows a comparison of the two architectures. For the sake of simplicity, the remainder of Section 3.2 focuses on the Simple VCS, but the process for working with the two plants is essentially the same.

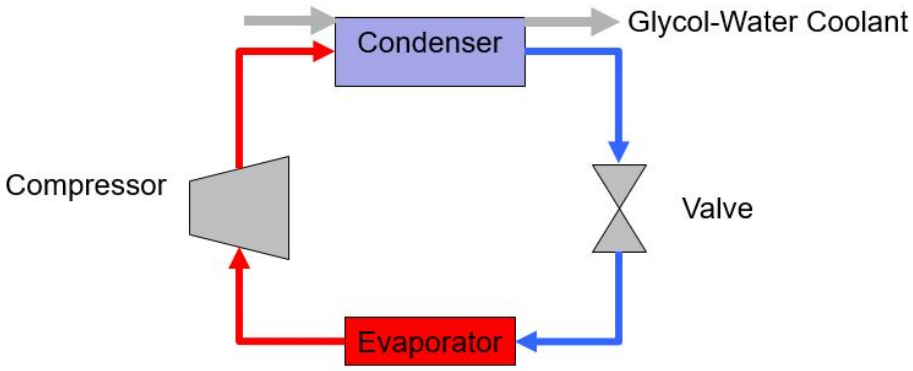

(a) Simple Vapor Cycle

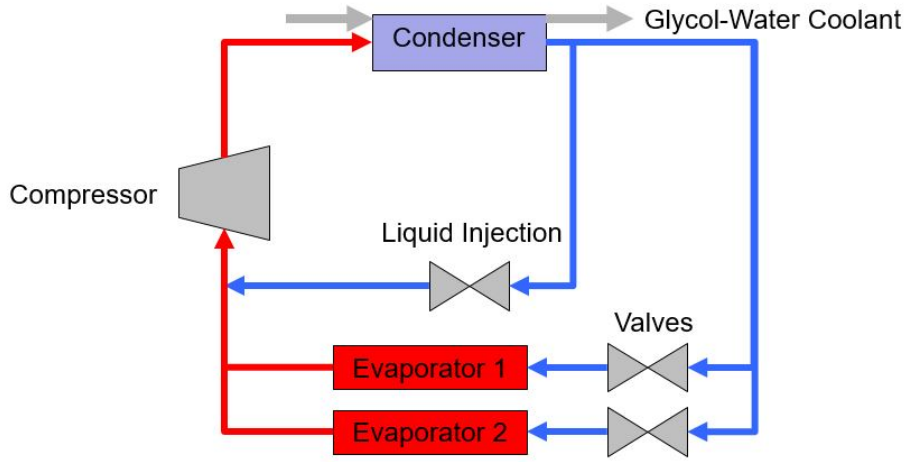

(b) Dual-Evaporator Vapor Cycle

Figure 3. Vapor Cycle Architectures

As discussed in Chapter 1, the Simple VCS is comprised of four components: evaporator, compressor, condenser and valve. The role that each component plays in the cycle is shown in the pressure-enthalpy diagram in Fig. 4. The ATTMO toolbox in Simulink is used to model the system. Static components are used to model the valve and compressor, whereas the heat exchangers are modeled dynamically, via a finite-volume method. The Simple VCS model is shown in Fig. 5 in the ATTMO user interface. A glycol-water mixture is used to cool the condenser, and Polyalphaolefin (PAO) is used on the load side of the evaporator.

A key distinction of this approach is the use of finite-volume heat exchanger models. Typical control-oriented vapor cycle modeling is done with moving-boundary heat exchanger models, which simplify the governing partial differential equations by reducing the flow into discrete zones occupied by each fluid state (i.e. gas, two-phase, and liquid). This results in a system of ordinary differential equations of relatively low-order (often < 10 states for each heat exchanger), which sacrifices accuracy in exchange for tractability. Finite-volume methods instead discretize the heat exchanger into many pieces, and conservation equations are then applied to each piece. This results in higher fidelity models which are slower and more complex, in exchange for improved descriptions of transient behavior. See Rasmussen [17] for a detailed discussion of the different approaches to heat exchanger modeling.

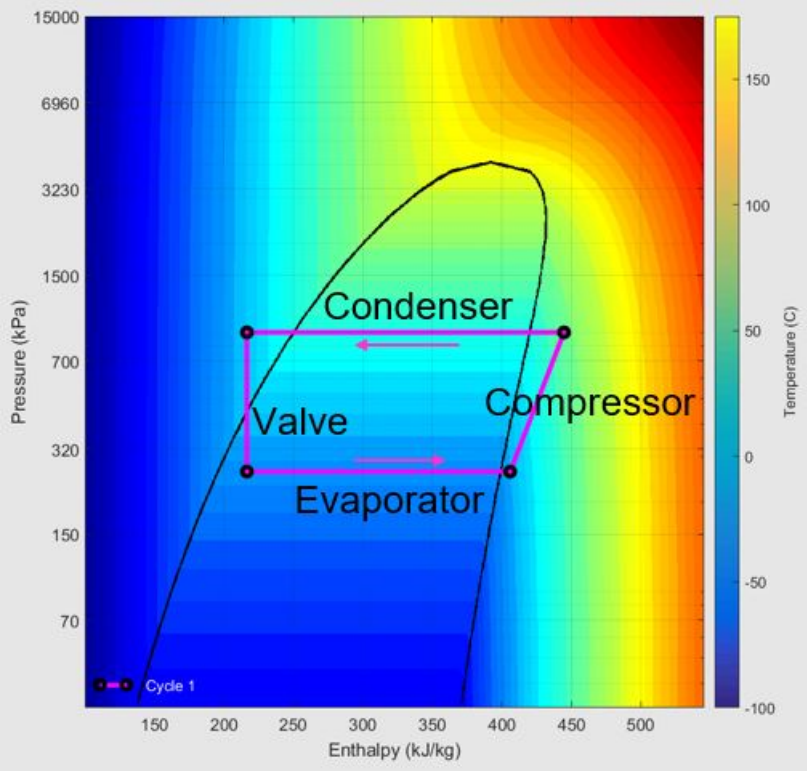

Figure 4. Pressure Enthalpy Diagram

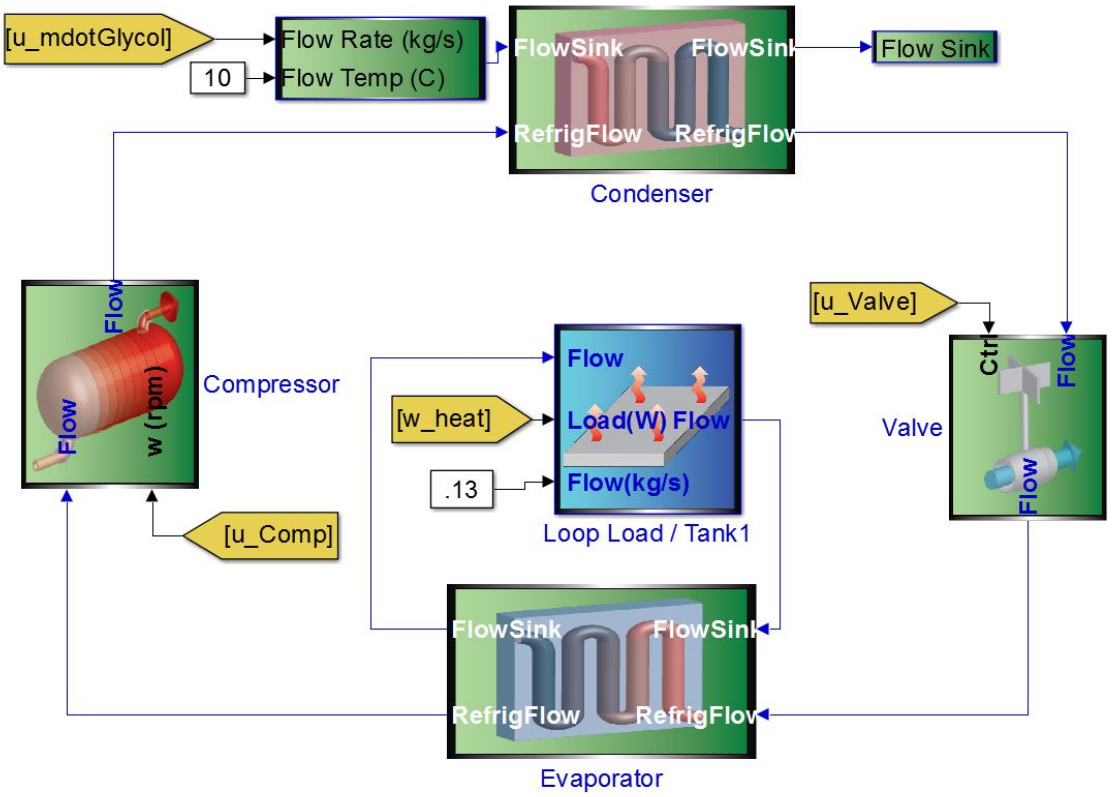

Figure 5. ATTMO VCS Model

3.2.2 Linearizing Models

In state space form, the model has a state vector $\tilde{x}$ which evolves according to some set of nonlinear dynamics $f_x(\tilde{x})$. The dynamics is also a function of $\tilde{u}$, which represents the system inputs (both controls and disturbances). Note that these terms are all vector-valued.

$$\dot{\tilde{x}} = f_x(\tilde{x}) + f_u(\tilde{u}) \tag{3.1}$$

Additionally, we use $\tilde{y}$ to denote the system outputs which we are interested in controlling (e.g. the evaporator load temperature). These outputs have some nonlinear relationship to the states $h_x(\tilde{x})$ and inputs $h_u(\tilde{u})$. Note that this is an algebraic equation, instead of a differential equation.

$$\tilde{y} = h_x(\tilde{x}) + h_u(\tilde{u}) \tag{3.2}$$

The field of control theory is essentially focused on the task of finding some control law for $\tilde{u}$, such that the state dynamics are stable (and perhaps meeting some level of desired performance). However, tools from control theory are limited when it comes to large nonlinear systems like ours. Additionally, the functions f and h are almost never explicitly known, since they are the result of a collection of nonlinear blocks spread out through a large Simulink diagram. So instead of attacking the nonlinear problem directly, we instead linearize it and deploy tools developed for linear systems. The resulting controller will work as long as the nonlinear system is kept "near" the operating point about which it was linearized.

To begin the linearization process, we choose some operating point to be a particular $\tilde{x}^*$ and $\tilde{u}^*$ such that the system is stable (i.e. $\dot{\tilde{x}} = f_x(\tilde{x}^*) + f_u(\tilde{u}^*) = 0$). Taking a Taylor Series Expansion of Eq. 3.1 and 3.2 about this operating point, and keeping only the linear terms, we get a linear system of the following form:

$$\dot{\tilde{x}} \approx 0 + \frac{\partial f_x}{\partial \tilde{x}}(\tilde{x} - \tilde{x}^*) + \frac{\partial f_u}{\partial \tilde{u}}(\tilde{u} - \tilde{u}^*) \tag{3.3}$$

$$\tilde{y} \approx y^* + \frac{\partial h_x}{\partial \tilde{x}}(\tilde{x} - \tilde{x}^*) + \frac{\partial h_u}{\partial \tilde{u}}(\tilde{u} - \tilde{u}^*) \tag{3.4}$$

To simplify the above equation, we make the following definitions:

$$x := \tilde{x} - \tilde{x}^* \tag{3.5}$$

$$u := \tilde{u} - \tilde{u}^* \tag{3.6}$$

$$y := \tilde{y} - \tilde{y}^* \tag{3.7}$$

$$A := \frac{\partial f_x}{\partial \tilde{x}} \tag{3.8}$$

$$B := \frac{\partial f_u}{\partial \tilde{u}} \tag{3.9}$$

$$C := \frac{\partial h_x}{\partial \tilde{x}} \tag{3.10}$$

$$D := \frac{\partial h_u}{\partial \tilde{u}} \tag{3.11}$$

This leads to the following compact representation of the linear system:

$$\dot{x} = Ax + Bu \tag{3.12}$$

$$y = Cx + Du \tag{3.13}$$

Note that the full nonlinear state $\tilde{x}$ has been broken into two parts, where the state "bias" $\tilde{x}^*$ is added to an incremental state x, with similar distinctions made for u and y. The linear system then describes the dynamics of these incremental states x in response to incremental changes u, and the effect on the incremental output y. To summarize, the system of equations in Eq. 3.12 and 3.13 are what we are talking about when we discuss the "linear" model, and the full ATTMO simulation is what we mean when we refer to the "nonlinear" model. The difference between the two models appears when x and u become "large" such that the higher order terms in the Taylor Series Expansion become non-negligible.

Relating all of this back to our Simple VCS, the elements of $\tilde{u}$ and $\tilde{y}$ are comprised of:

$$\tilde{u} = \begin{bmatrix} Compressor\ speed \\ Valve\ opening \\ Glycolwater\ flowrate \\ Evaporator\ heat\ load \end{bmatrix} \tag{3.14}$$

$$\tilde{y} = \begin{bmatrix} Condenser\ Pressure \\ Evaporator\ Pressure \\ EvaporatorLoad\,Temperature \end{bmatrix} \tag{3.15}$$

Note that the pressures in Eq. 3.15 are often referred to as the Saturated Discharge Temperature (SDT) and Saturated Suction Temperature (SST), respectively, since temperature and pressure are constant during the phase change process. The names are in reference to the compressor, which sucks from the evaporator and discharges into the condenser. Rather than measuring these temperatures directly in a real system, the intent is that these signals would be found by measuring pressures at the heat exchanger outlets, and converting them to the associated saturation temperature, based on the isobaric heat exchanger assumption. In other words, the SST is found by converting the pressure at the evaporator outlet to an associated saturated gas temperature via a table-lookup, and the SDT is found similarly from the condenser outlet temperature. Lastly, it's worth noting that the evaporator load temperature in Eq. 3.15 is often called the PAO temperature, in reference to the name of the fluid delivering heat to the evaporator.

In practice, it is possible to find the matrices $\{A, B, C, D\}$ in Eq. 3.12 and 3.13 by hand, but the *linearize()* function in MATLAB instead uses a numerical perturbation algorithm to estimate these terms much more quickly. The estimate is based on a finite-difference approximation, shown in Eq. 3.16 for B, with similar approximations made for A, C, and D.

$$B := \frac{\partial f_u}{\partial \tilde{u}} \approx \frac{f_u(u^* + \delta u) - f_u(u^* - \delta u)}{2\delta u} \tag{3.16}$$

Note in Eq. 3.16 that the δu term is vector-valued. In order to ensure that the perturbations are of the appropriate magnitude, we must include scaling gains on the perturbed inputs. Without these scaling terms, the algorithm may skew the dynamics in favor of the small signals. For example, the compressor speed will experience changes in the 1000's of rpm's, whereas the glycol-water mass flow rate will rarely deviate by more than 0.5. Without scaling gains, a perturbation of $\delta u = \pm[0.1, 0.1, 0.1, 0.1]$ would result in a linear model with very "flat" responses to the compressor, and very "steep" responses to the glycol-water mass flow rate. Therefore, the scaling gains ensure that all of the inputs and outputs be of order 1.

With the inputs u and outputs y identified for the ATTMO model and the appropriate scaling gains chosen, the *linearize()* function in MATLAB can be used to obtain a linear system of the form shown in Eq. 3.3 and 3.4. The next task is then to verify that the linear model is actually a good representation of the local nonlinear dynamics about the operating point. To answer this question, we send varying amounts of u into both systems, and compare the outputs in y. Figure 6 shows the results of performing this study.

In Fig. 6, the nonlinear system starts at its operating point, and the linear system starts at $x = 0$. Note that both of these correspond to the same starting point, and the linear outputs are scaled and added to the output bias y^* to allow for comparison. From here, a series of pulses across all four elements in u is sent into the system. We'd like to see that each system responds exactly the same to the inputs, indicating that the linear model is a very good approximation of the nonlinear model. That is not what is shown in the figure, but the two systems respond similarly enough to expect reasonable performance from a controller

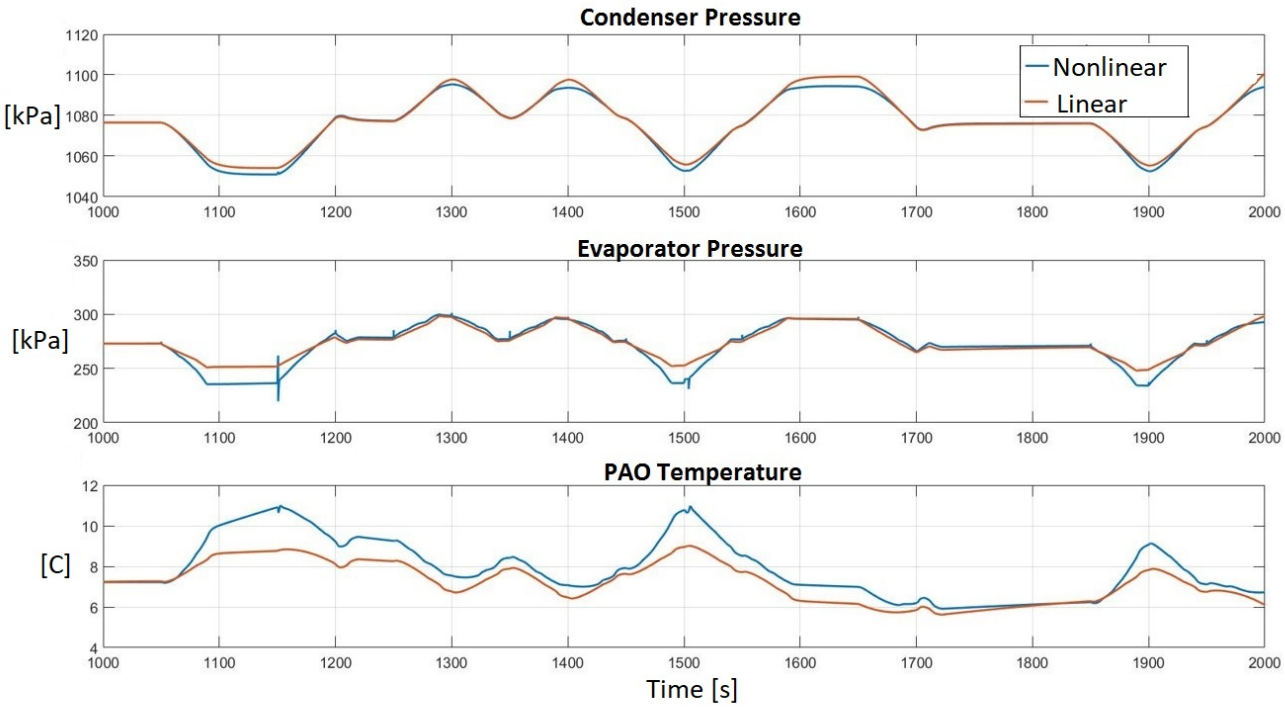

Figure 6. Open-Loop Comparison: Linear vs. Nonlinear

designed with the linear model in mind. It is important to note however, that the height of the pulses used in this test is important. The height of the pulses varies for each element in u, and as these heights get larger and larger, the difference between linear and nonlinear responses in Fig. 6 will necessarily grow. This is due to the inherent nonlinearities in the system, and we should expect some level of performance degradation whenever we experience these large deviations from the operating point. The beginnings of these issues are shown in the bottom plot of Fig. 6, where the linear model isn't as sensitive as the nonlinear model.

At this point, we have a good linear approximation of the full nonlinear system, but the linear system has around one hundred states. The majority of these states are not relevant to the control system designer, who may only be interested in the states associated with the inputs and outputs. Here, we use model reduction to

extract those parts of the model which are most important to us. Roughly speaking, the importance of each state can be measured along two dimensions: controllability and observability. Those states which are strongly affected by varying u are more controllable than others. Similarly if varying some states causes strong variations in the output y, then those states are said to be observable. These concepts can be rigorously captured in the observability and controllability Gramians W_o and W_c, which result from performing singular value decomposition on the controllability and observability matrices, defined in any linear systems theory textbook.

A naive first approach to model reduction might focus on calculating the controllability Gramian W_c, and then removing those dimensions of the state space which have a eigenvalues below a certain threshold. This approach would toss out parts of the model which are not very controllable, but it runs the risk of throwing out dimensions of the state space which are strongly observable, and are important to keep. The process for finding a balanced realization solves this problem by finding a state space realization which aligns the two concepts of controllability and observability. A similarity transform matrix T can be found which transforms the state space system to an alternate realization:

$$\bar{x} = Tx \tag{3.17}$$

$$\dot{\bar{x}} = TAT^{-1}\bar{x} + TBu \tag{3.18}$$

$$y = CT^{-1}\bar{x} + Du \tag{3.19}$$

This new system in $\bar{x}$ has the same input-output relationship between u and y as the original system, but the states $\bar{x}$ are "balanced" in that their controllability

and observability properties are the same. In other words, T has the property that:

$$\bar{W}_c = TW_cT^T = \bar{W}_o = T^{-T}W_cT^{-1} = diag(g) \tag{3.20}$$

The diagonal entries in the Gramian matrix (sometimes called "Hankel singular values") can then be ranked and plotted, showing the decay in importance for each state in the system. Figure 7 shows this decay on a logarithmic scale.

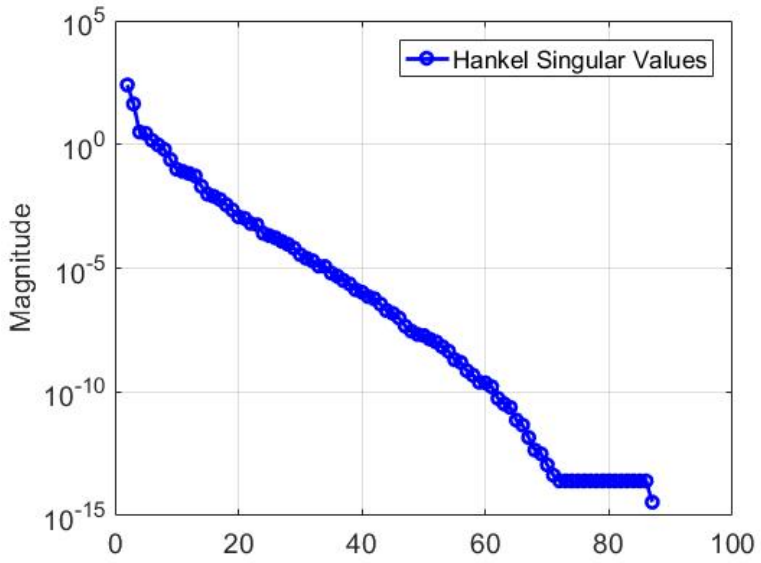

Figure 7. Hankel Singular Values

The *balreal()* command in MATLAB is used to find the similarity transform matrix T, but the control system designer is responsible for intelligently choosing the cutoff value for how many states to keep. Typically, one chooses to keep somewhere between 10 to 20 states, and compares the two system's Bode plots to verify that the dynamics are similar.

Lastly, there are two common methods for actually removing the states below the chosen threshold. Denote $\bar{x_1}$ as the subset of states which are to be kept, and $\bar{x_2}$ as the remaining states to be removed. The dynamics can then be written:

$$\begin{bmatrix} \dot{\bar{x}}_1 \\ \dot{\bar{x}}_2 \end{bmatrix} = \begin{bmatrix} \bar{A_{11}} & \bar{A_{12}} \\ \bar{A_{21}} & \bar{A_{22}} \end{bmatrix} \begin{bmatrix} \bar{x_1} \\ \bar{x_2} \end{bmatrix} + \begin{bmatrix} \bar{B_1} \\ \bar{B_2} \end{bmatrix} u \tag{3.21}$$

The truncation method simply removes $\dot{\bar{x}}_2$ entirely from the equation, whereas the residualization method sets $\dot{\bar{x}}_2 = 0$, and solving for the resulting system of equations. The residualization method is what we use in our study, because it has the benefit of matching the steady-state values of the full system. However, this method has the drawback that the model will have a nonzero feed-forward matrix D, which we'll see causes problems with estimation later on. The interested reader is recommended to consult textbooks on the theory behind model reduction for a more rigorous treatment of the subject.

Chapter 3.3 deals with control design, and often refers to a linear system $\{A, B, C, D\}$. It is important to keep in mind that this linear system is actually the reduced order linear system that we've just finished extracting, and *not* the complete linear system.

3.3 Control Design

3.3.1 Control Design Overview

Chapter 3.3 covers the process for designing controllers, and it is broken up into three parts. Section 3.3.2 makes up the majority of the discussion, and covers all of the theory behind the LQR controllers used in this thesis. Section 3.3.3 covers the PI control architecture. Finally, Section 3.3.4 discusses the controller specifications, which are used to evaluate the performance of the two types of controllers.

3.3.2 LQR Control Design

3.3.2.1 LQR Overview

The LQR controller is a linear optimal controller, which is designed to be optimal with respect to a particular linear system and cost functional. While we have been calling this controller a regulator, it's really the combination of a regulator and a state estimator. Often, this combination is called a Linear Quadratic Gaussian [31], but we have refrained from using that term because it is often associated with systems with white noise disturbances. The regulator is simply a feedback matrix K_x, mapping between the current state estimate, to an associated control action. However, the state vector is an abstract mathematical object which cannot be measured directly, and so it is instead estimated from the sensor measurements with a Kalman filter.

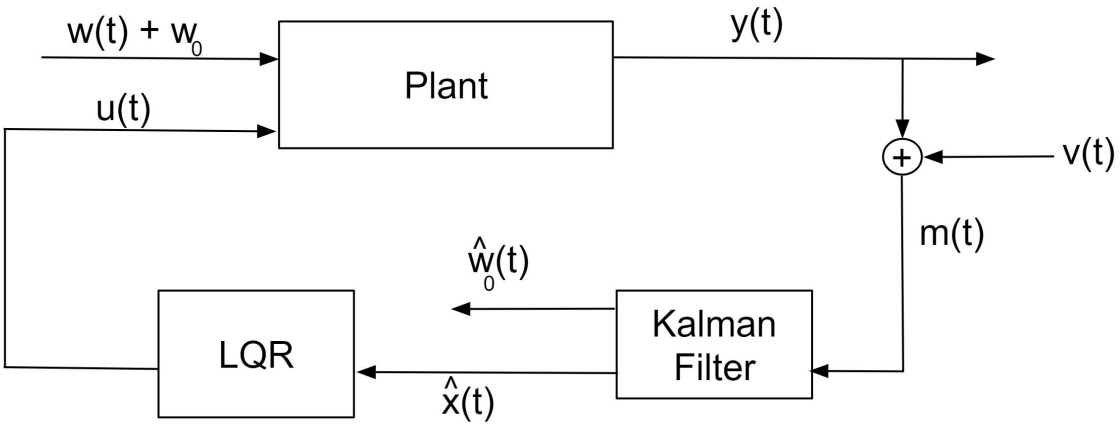

Figure 8. LQR Control Architecture

For context, the complete LQR control architecture is shown in Fig. 8. The "plant" is the system which requires controlling (i.e. the VCS). In practice, this block contains the full nonlinear system, but throughout this control design process, we treat the plant block as if it only contained the linearized system extracted throughout the course of Chapter 3.2. In Chapter 3.2, we lumped both the control

and the disturbance inputs in the vector u to keep things simple, but we now separate the terms input a control vector $u(t)$, and a vector of disturbances w_0. There are two fictional white noise disturbances $w(t)$ and $v(t)$ that are shown in the diagram, but these terms are not representative of anything physical. They are simply used to tune the Kalman filter, and they will be discussed more later in this section. The plant outputs $y(t)$, which represents the controlled outputs (e.g. SDT). The measurement vector $m(t)$ also comes from the plant, and represents the same signals in $y(t)$, with additive white noise from $v(t)$. In our simulations $m(t) = y(t)$, but in real systems, there is often differences between the two signals, due to noise or time delays. Continuing along the feedback path, the Kalman filter provides an estimate of the disturbance bias $\hat{w_0}(t)$, which goes unused, and an estimate of the linear state vector $\hat{x}(t)$. This state estimate is fed through the LQR gain matrix, which provides the control action $u(t)$.

The remainder of this section will walk through the design process for the LQR controller. Section 3.3.2.2 discusses the theory behind the regulator design, and Section 3.3.2.3 discusses the similar theory behind the estimator design. These two topics are covered in most introductions to the theory behind LQR, but Section 3.3.2.4 provides a discussion of the extra steps required to introduce integral control and disturbance bias estimation, which are not always shared in introductory treatments of the subject. It is important to revisit Fig. 8 throughout this section, to anchor each concept in the overall controller architecture.

3.3.2.2 Regulator Design

The field of optimal control focuses on finding control laws that minimize some cost functional for a given set of dynamics. With LQR, the cost functional takes the

form shown in Eq. 3.22, where the plant has the linear dynamics shown in Eq. 3.12 and Eq. 3.13.

$$J = \frac{1}{2}\int_0^\infty x^T(t)Qx(t) + u^T(t)Ru(t)dt \tag{3.22}$$

The cost functional J shown in Eq. 3.22 is quadratic in x and u, with a positive definite matrix R and positive semi-definite matrix Q. J is therefore strictly positive, and it can be shown that the feedback gain matrix K_x in Eq. 3.23 can be found to minimize J, for a linear system linear system with dynamics A and B_u.

$$u(t) = -K_x x(t) \tag{3.23}$$

Here, notice that we break B into two parts, B_w and B_u, to capture the fact that the plant in Fig. 8 has the two inputs w_0 and $u(t)$. The regulator design process ignores the disturbance action from B_w, which is instead considered in the estimator design process.

For a given set of weights $\{Q, R\}$ and dynamics $\{A, B_u\}$, the LQR design process yields a unique feedback gain matrix K_x which minimizes the value of J for any bounded input. Generally, software is used to solve the algebraic Riccati equation required to find K_x, as with the *lqr()* command in MATLAB. Tuning of the controller takes place in the diagonal matrices Q and R. When the elements of Q are large relative to the elements of R, the cost function J is weighted more heavily towards the states, and so the feedback gain matrix K_x will use large amounts of control to drive the states down to zero. Conversely, when R is large relative to Q, the controller will be much more passive with its use of control. Again, it's worth pointing out that this approach is only optimal with respect to the

linear system, but closed-loop performance should be adequate, so long as the nonlinear system is kept near the linearized operating point.

At this point, we have the feedback gain matrix K_x, which corresponds to the LQR block in Fig. 8. Next, we describe the method for obtaining an estimate of the state vector $\hat{x}(t)$.

3.3.2.3 Estimator Design

The state vector feedback $x(t)$ shown in Eq. 3.23 represents an abstract mathematical term, which does not correspond directly to anything physical. An estimator can be designed to provide estimates of these states (denoted with a hat $\hat{x}(t)$), based on available measurements $m(t)$. There are numerous ways to provide estimates of the state vector, one of which is called a Linear Quadratic Estimator (LQE), also known as a Kalman filter. The design process begins with a linear model of the dynamics, with additive white noise on the inputs and outputs.

$$\dot{x}(t) = Ax(t) + B_u u(t) + B_w(w(t) + w_0) \tag{3.24}$$

$$m(t) = Cx(t) + v(t) \tag{3.25}$$

The two above equations represent the estimator's model of the plant dynamics in Fig. 8. Here, $w(t)$ and $v(t)$ are white noise disturbances on the plant and sensors, respectively. Additionally, we ignore the w_0 term for now. The white noise signals have spectral densities:

$$E[w(t)w^T(t+\tau)] = S_w\delta(\tau) \tag{3.26}$$

$$E[v(t)v^T(t+\tau)] = S_v\delta(\tau) \tag{3.27}$$

The design element of this estimator takes place in the spectral density matrices S_w and S_v, which play a similar role to the Q and R matrices from the regulator design process. Using the *lqr()* command to design the estimator is possible, since the LQE and LQR approaches have some duality associated with them, and the result is an optimal Kalman gain L [31]. That is to say that for a given set of weights $\{B_w S_w B_w^T, S_v\}$ and dynamics $\{A^T, C^T\}$, the LQE design process yields a unique Kalman gain matrix L which minimizes the mean square estimation error. Traditionally, the spectral density matrices are measures of real sources of noise, which drive the design of the estimator. However, in our simulation case study, these matrices are used only as tuning parameters for finding the Kalman gain matrix. Once a set of values are chosen for the spectral density matrices, the algebraic Riccati equation is solved, L is found, and the estimator's dynamics are propagated to provide a state estimate $\hat{x}(t)$.

$$\dot{\hat{x}}(t) = A\hat{x}(t) + B_u u(t) + L(m(t) - C_m \hat{x}(t)) \tag{3.28}$$

At this point, we have a working regulator and estimator, completing the feedback loop in Fig. 8. However, we have entirely neglected the effect of the disturbance bias w_0. The next section will provide a method for dealing with this term.

3.3.2.4 Augmented LQR

Referring back to the block diagram in Fig. 8, we see that we have completed the design of the Kalman filter and LQR blocks. However, we have ignored the disturbance bias w_0, which can result in degraded performance. This is due to the fact that the estimator has no knowledge of the disturbance bias, and the result will be a steady-state error in the state estimate. For the simulation case studies in

Chapter IV, there is a large pulse in the heat load, which can be modelled as a bias in the disturbance. In many practical applications, these disturbances cannot be measured accurately in real time, and so some form of estimation is required. One way to address this problem is via a combination of integral control and disturbance estimation [31]. In this section, we'll discuss how the estimator is augmented to provide estimates of w_0 and how the regulator is augmented to provide integral control action.

First, we will discuss the process for introducing an estimate of the disturbance bias. Ignoring the fictional noise term $w(t)$ in Eq. 3.24, the linear dynamics can be written as:

$$\dot{x}(t) = Ax(t) + B_u u(t) + B_w w_0 \tag{3.29}$$

We seek to design an estimator that not only provides state estimates, but also estimates the disturbance bias w_0. To do this, we rearrange the dynamics in the previous equation such that a new state $\hat{w}_0(t)$ is appended to the state vector estimate. After doing this, the augmented estimator dynamics are written as:

$$\begin{bmatrix} \dot{\hat{x}}(t) \\ \dot{\hat{w}}_0(t) \end{bmatrix} = \begin{bmatrix} A & B_w \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{x}(t) \\ \hat{w}_0(t) \end{bmatrix} + \begin{bmatrix} B_u \\ 0 \end{bmatrix} u(t) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \tilde{w}(t) \tag{3.30}$$

Note that in the above equation, the zeros are padding matrices of the appropriate dimension (i.e. each zero represents a matrix of zeros such that matrix-vector multiplication makes sense). Upon inspecting the above equation, we see that dynamics of the state estimate now account for the effect of the bias, and that the estimate of the bias is driven by a fictional noise term $\tilde{w}(t)$. This noise term is not representative of anything physical, and it is only used as a design

parameter to tune how the bias estimate converges (otherwise $\hat{w}_0(t)$ would have no dynamics). The bias estimate is not fed to the regulator, but it is important to capture the $B_w w_0$ term in order to provide accurate state estimates. Lastly, we point out that the bias is written as a constant w_0 to be estimated, even though the evaporator heat loads in Chapter IV are time-varying. So long as the loads allow enough time for the estimator to converge on a useful estimate, the method should work as if the w_0 were constant.

The second modification to the controller is the introduction of integral control in the regulator, which helps to mitigate any steady state errors caused by the disturbance bias. First, a reference signal $r(t)$ is introduced, representing the desired plant output. The error $e(t)$ and the integral error $e_I(t)$ are then defined as:

$$e_I(t) := \int_0^\infty e(t)dt = \int_0^\infty r(t) - y(t)dt \tag{3.31}$$

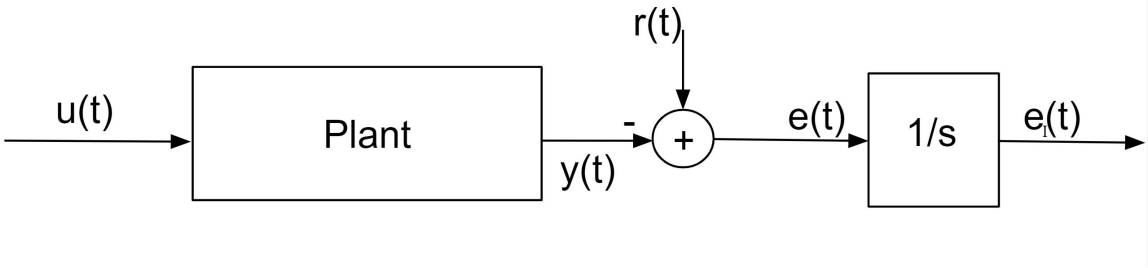

Figure 9. LQR Integral Control Augmentation

A block diagram of this modification is shown in Fig. 9, showing that we've added an integrator block $\frac{1}{s}$ that acts on the newly defined error term. For our purposes, $r(t)$ is simply zero, but this approach can be modified to track different output setpoints by choosing nonzero values for $r(t)$. With that said, it is worth mentioning that such setpoint changes were investigated in this research, and there were issues with keeping these setpoint changes from interfering with the bias estimation. In any case, by setting $r(t) = 0$, the error dynamics are:

$$\dot{e}_I(t) = -Cx(t) \tag{3.32}$$

This integral error term is then appended to the state vector, and the augmented plant is written as:

$$\begin{bmatrix} \dot{x}(t) \\ \dot{e_I}(t) \end{bmatrix} = \begin{bmatrix} A & 0 \\ -C_y & 0 \end{bmatrix} \begin{bmatrix} x(t) \\ e_I(t) \end{bmatrix} + \begin{bmatrix} B_u \\ 0 \end{bmatrix} u(t) = A_I x_I(t) + B_I u(t) \tag{3.33}$$

The process for designing the LQR controller then follows the same path as before, with the augmented plant shown in Eq. 3.33. When building the cost functional J, the weighting matrix Q is enlarged to accommodate the augmented state vector, and all of the elements in Q are set to zero, except for those weighting the integral error states. R is kept the same, and the A and B_u matrices used in the *lqr()* command are instead the augmented $\{A_I, B_I\}$ matrices. The resulting feedback gain matrix K can be broken down to illustrate the fact that we now have integral control action.

$$u(t) = -Kx_I(t) = -K_x x(t) - K_I \int_0^\infty r - y(t)dt \tag{3.34}$$

This concludes the discussion of the theory behind the model-based controller's design. To summarize, the LQR controller is broken into two parts: the regulator and the estimator. The regulator is essentially a non-square matrix K mapping a state estimate to the appropriate control action u. The plant is augmented with an integral error term, in order to introduce integral control action to the regulator.

The estimator uses a linear model of the state dynamics to provide a dynamic estimate of the state vector $\hat{x}(t)$, based on sensor feedback measurements $m(t)$. When designing the estimator, we use a linear model of the disturbance bias's effect

on the system (B_w), in order to append an estimate of the disturbance bias to the state estimate. This mitigates the steady-state estimation error caused by the disturbance bias.

When actually designing the controller, one typically starts by finding good values for the tuning matrices $\{Q, R\}$ to design the regulator with full-state feedback, on the reduced-order linear system. Once these values are found, the estimator can be inserted before the regulator, driven by feedback measurements $m(t)$. In this way, effective values for $\{S_w, S_v\}$ can be found. Once all of these tuning parameters are found to be satisfactory, the reduced-order linear system can be replaced in simulation with the full nonlinear system (with appropriate scaling and biasing of $u(t)$ and $y(t)$ about the operating point), and a final round of tuning occurs. So long as the linear model is a good approximation of the nonlinear system, this final round of tuning shouldn't require much effort.

3.3.3 PI Control Design

To provide context for the LQR controller's performance, a set of PI controllers were designed [2] to provide a performance baseline. These PI controllers were implemented in separate SISO loops, shown in Fig. 10, with the controller gains being chosen based on heuristics and manual tuning. In the Simple VCS, the pairings are as follows: the compressor controls the evaporator pressure, the glycol-water flow rate controls the condenser pressure, and the valve controls the PAO temperature. Each of these pairings correspond one of the PI blocks in Fig. 10.

Each PI block has a single element of the output vector $y(t)$ being fed into it, and a single element of the control vector $u(t)$ being output by it. The key difference between the PI architecture in Fig. 10 and the LQR architecture in Fig. 8 is the fact that the PI blocks are decoupled, whereas the LQR feedback is coupled.

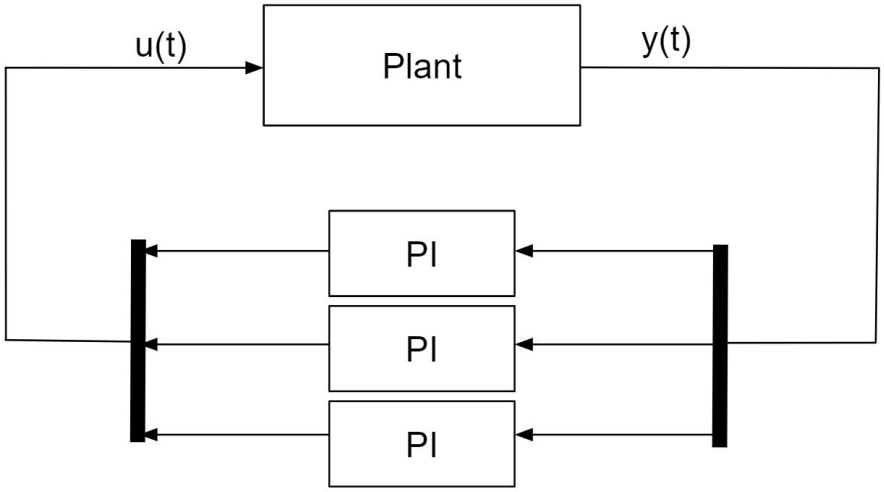

Figure 10. PI Controller Architecture

3.3.4 Controller Specifications

Table 2 provides a set of specifications for the controllers in the Simple VCS case study. This table consists of target envelopes, within which the particular output signals must be maintained. Brief departures from these envelopes is tolerable, but undesirable. Hard constraints are specified for some of the signals, with these constraints being associated with hardware failure modes for the system. None of these constraints are explicitly included in the controller design, but these specifications provide a way to compare performance between two controllers, and to decide whether or not a particular simulation run was successful. The process of controller tuning is the process of finding controller gains such that these performance specifications are met. The values below are for the Simple VCS system, but similar constraints are used for the Dual-Evaporator VCS used in the second set of simulation case studies.

Term	Target Envelope	Hard Constraint
Low side temperature	26.5 F ± 2 F	
High side temperature	97 F ± 3 F	[∼,115] F
Evaporator sink temperature	45 F ± 2 F	
Condenser Subcooling		[10,∼] F
Evaporator superheat		[0,∼] F
Compressor speed		[4000,18000] rpm
Valve opening		[0,100] %
Glycol-water flowrate		[0,1] kg/s

Table 2. Simple VCS Controller Specifications

IV. Results and Analysis

4.1 Overview

This chapter contains the results of the simulation work undertaken in this thesis. The results are broken up into two parts. First, the simple vapor cycle system is simulated under a pulse disturbance in the evaporator heat load. Comparisons are made between the closed-loop performance of the LQR and PI controllers. In the second section, the same process is performed with the Dual-Evaporator vapor cycle system, with a more varied disturbance profile across the two evaporators.

4.2 Simple Vapor Cycle Results

The steady-state evaporator heat load for the Simple VCS was 5 kW, based on the steady-state heat load for which the heat exchangers were sized in the ATTMO simulation. For the simulation case studies, pulses above this 5 kW load were investigated, and it was found that a doubling of the heat load was roughly where the first signs of performance degradation were observed in one of the control loops. The results for this case study are broken up into two parts. First, we show the control action for each controller in Fig. 11, and then we show the output performance in Fig. 12 and 13.

In the top-left of Fig. 11, we show the heat pulse as a solid-red line, which pulses from 5 kW up to 10 kW for 50 seconds. Also displayed in the top-left plot is the LQR controller's estimate of the heat load (shown as a dashed magenta line), which corresponds to the $\hat{w}_0$ term discussed in Chapter 3.3. Close inspection of this plot shows that the heat estimate takes about 1 second to react to the loading event, with a 15% overshoot. When the pulse is turned off, the estimate takes 2

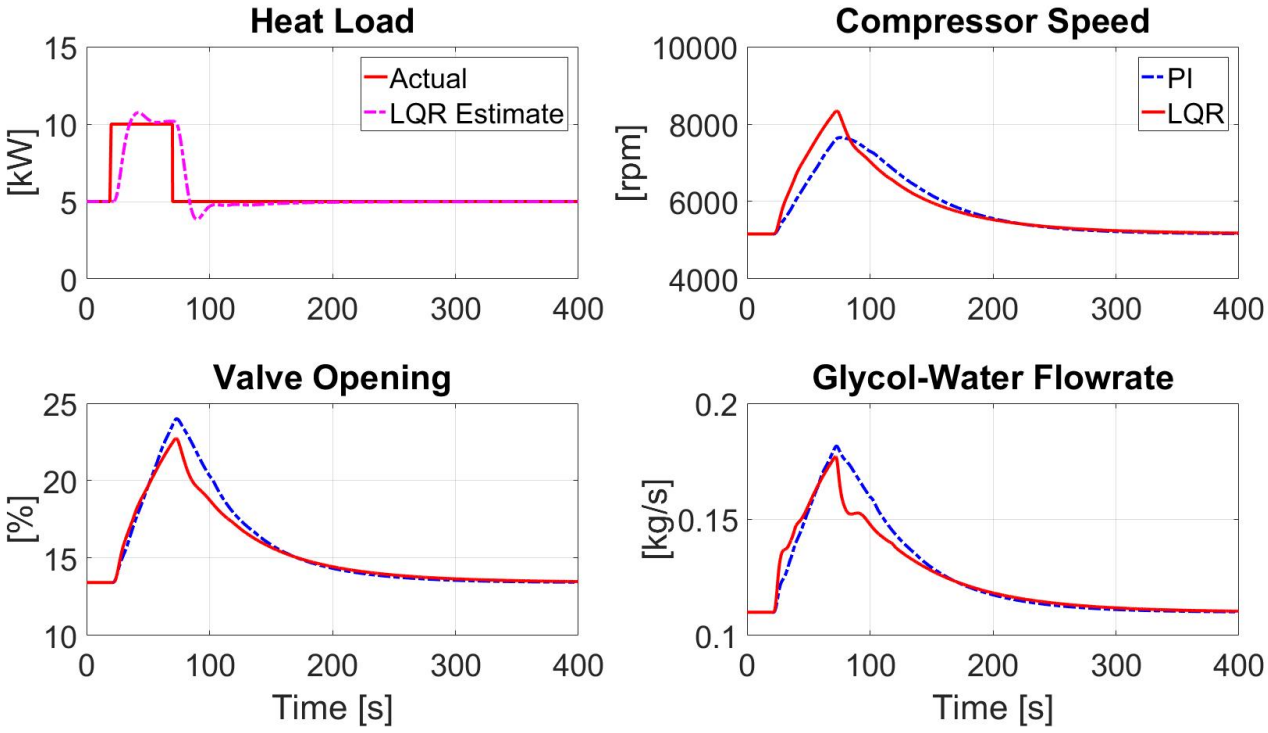

Figure 11. Simple VCS: Disturbance, and Controller Responses

seconds to respond to the event and has a 23% overshoot. This asymmetry in time delay and overshoot performance is likely due to the fact that the estimator is based off of the linearized dynamics at the operating point (i.e. at 5 kW), and so the estimator's performance is degraded when the system is under 10 kW of load.

The other subplots in this figure show the control actions for the PI and LQR controller, across the three actuators. PI control trajectories are shown as dashed-blue, with LQR trajectories shown in solid-red. These plots all show roughly the same pattern: shortly after the loading event begins at 20 seconds, both the PI and LQR controllers speed up the compressor, open up the valve, and increase the glycol-water flow rate. All of these plots peak just after the pulse turns off at 70 seconds, and eventually return to the steady-state values. Note that these steady state values correspond to the bias term u^* discussed in Chapter 3.2.

Comparing the two sets of control trajectories, we see that the LQR controller is more aggressive with its use of the compressor (in that it reaches a higher peak value near 70 seconds), whereas the PI controller uses more of the valve. This has more to do with differences in tuning than it has to do with the controller designs, as the PI gains can be adjusted to achieve a closer match between these two trajectories. However, the LQR control trajectories do show an interesting feature that differentiates them from the PI control trajectories. When the loading event begins, the LQR controller responds quickly and aggressively, indicated by the steep slope of the initial responses. This rapid response is required in order to keep the evaporator sink temperature from rising, as the evaporator heat load has the strongest affect on this output. Following the rapid spike in control action is a change in the slope, where the actuators work together to control the other two outputs. This contrast is most apparent in the bottom-right plot of the glycol-water flow rate, but the trend is present across all three actuators. Lastly, the same feature is seen when the pulse ends, with a rapid response followed by a slower response.

The first set of outputs are shown in Fig. 12, again using dashed-blue to represent the PI trajectories and solid-red for LQR. These are called the "controlled" outputs, since they represent the terms in $y(t)$, which are used for controller feedback. The other set of outputs are shown in Fig. 13, which are called the "uncontrolled" outputs since they are not actively regulated by the controllers. Ideally, we'd like to see that all of the output trajectories are flat lines and remain within the target envelopes (represented as dashed black lines at constant values corresponding to the values shared in Table 2). However, we see that all of these signals begin to deviate from their setpoint when the loading event begins at 20 seconds. In the top plot, we see that the PI controller performs perfectly in regulating the condenser pressure (i.e. the SDT, shown as the top horizontal line in

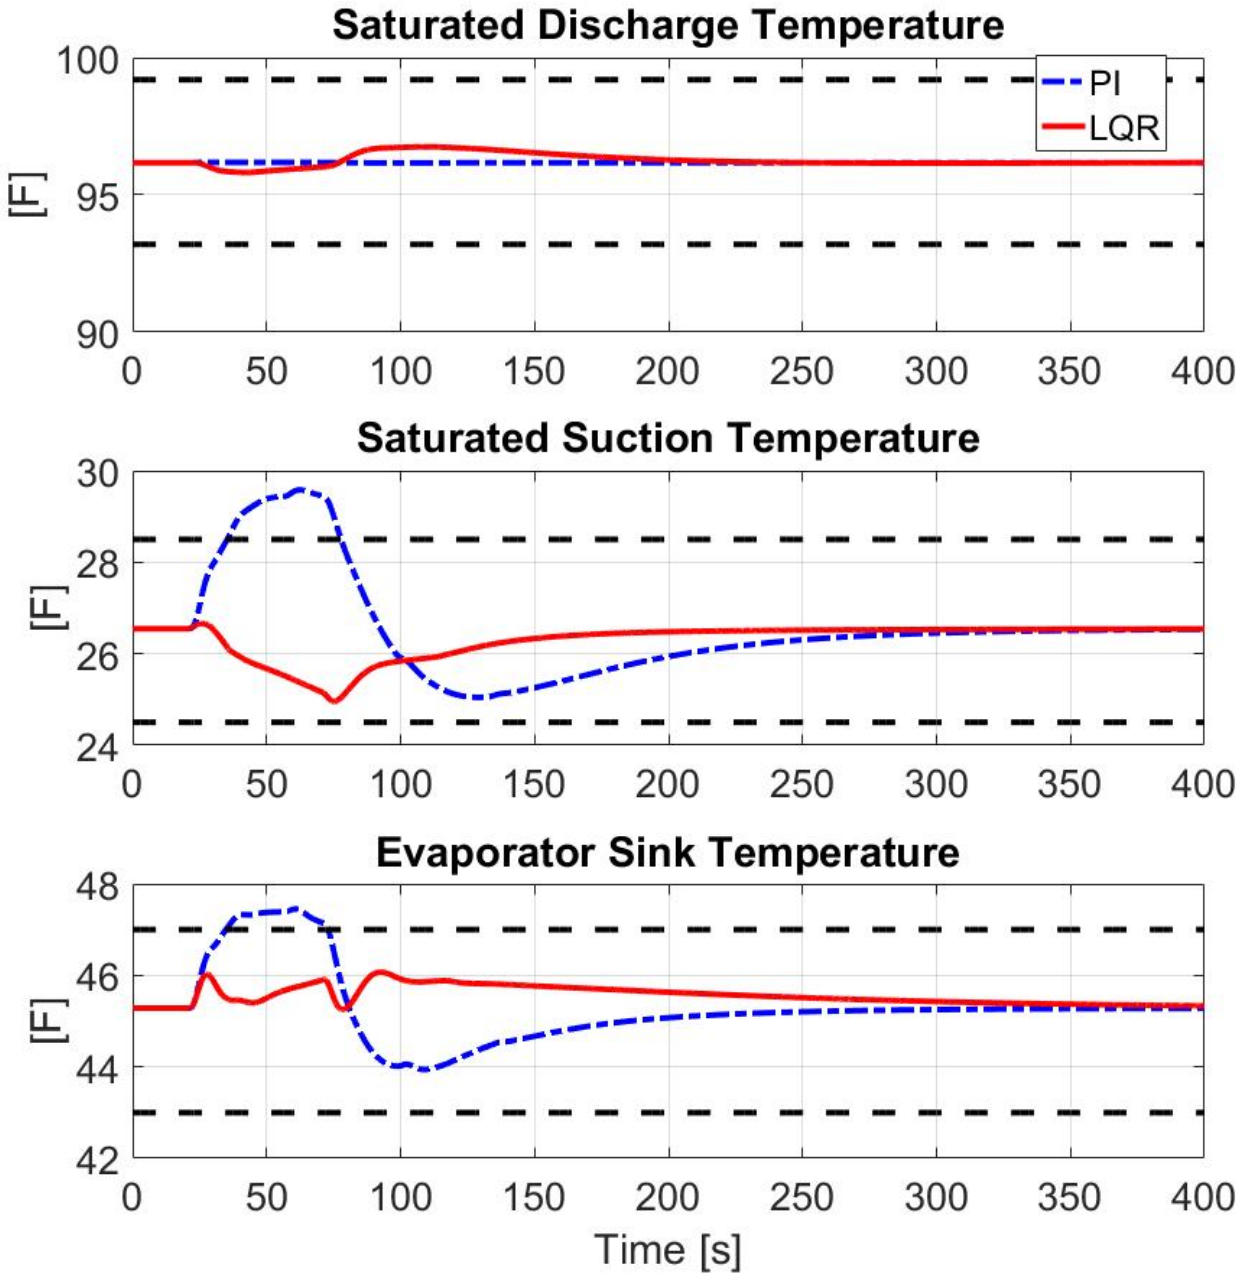

Figure 12. Simple VCS: Controlled Outputs

the vapor cycle pressure-enthalpy diagram in Fig. 2). However, both controllers keep the signal well within the desired constraints. The second plot shows the evaporator pressure (i.e. the SST, shown as the bottom horizontal line in Fig. 2),

and we see that the PI controller is unable to keep the signal within the constraints. This does not correspond to a system failure, but it does show that the LQR controller outperforms the PI controller in regulating the SST. In the bottom plot of Fig. 12, we see that the LQR controller also outperforms the PI controller in regulating the evaporator sink temperature (i.e. the PAO temperature, on the sink side of the evaporator). All of these factors indicate that LQR is outperforming PI in this case study, but we have to verify that all of the constraints in Table 2 are respected, which includes the two signals shown in Fig. 13.

Fig. 13 shows the uncontrolled signals of the evaporator superheat and the condenser subcooling. These two signals are not used for feedback, but it is important that they are kept above the hard constraints, shown in the plots as solid black lines. The bottom plot shows that both controllers fall below the target envelope for a period of time, but the PI loop additionally breaks the hard constraint, and loses superheat during a portion of the loading event. This can cause compressor failure, if liquid starts entering the device, and represents a critical failure in the PI controller's performance. It may be possible to tune the PI gains to improve that controllers performance, but the process of tuning the gains is typically based on heuristics and it can be very challenging to search for the correct set of gains, since the control architecture does not account for any cross-couplings across the different actuators.

From this case study, it is clear that the LQR controller outperforms the PI controller in regulating the SST, the evaporator sink temperature, and the evaporator superheat. However, as a final point of discussion for this section, it is worth pointing out one of the limitations of the LQR method. In Fig. 12, the evaporator sink temperature with LQR control shows a slow increase in the temperature, starting just before 50 seconds and ending once the pulse was shut off

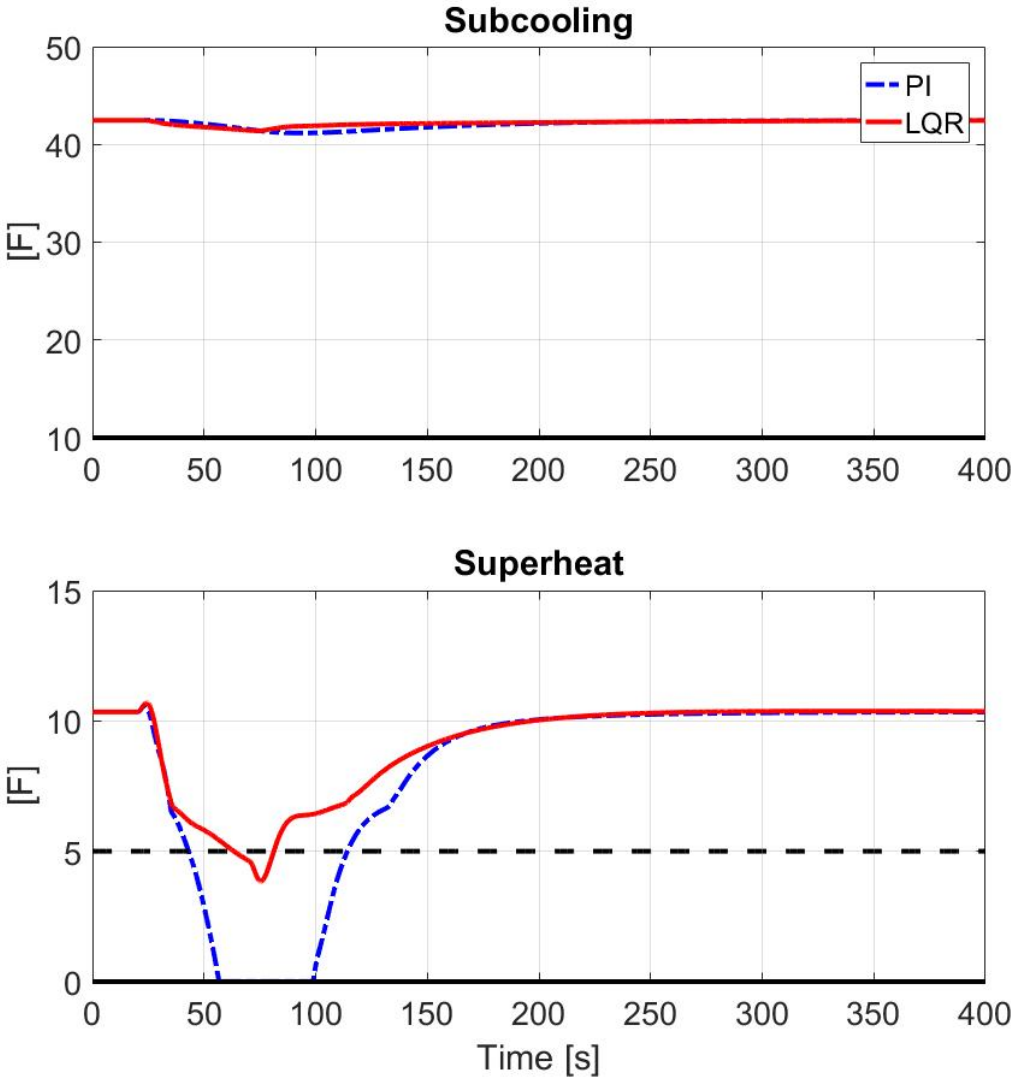

Figure 13. Simple VCS: Uncontrolled Outputs

at 70 seconds. The cause of this increase is due to a feed-forward term in the linear model of the dynamics, which was introduced when performing the residualization step discussed at the end of Chapter 3.2. This feed-forward term should ideally be accounted for in the estimator, but accounting for feed-forward dynamics in a controller is notoriously susceptible to instability caused by modelling errors [31]. Since the system is really nonlinear, modeling errors are absolutely present, and the expected instabilities were encountered in simulation case studies which tried to

account for the feed-forward term. For this reason, this LQR controller cannot reject heat pulses that last for a very long period of time in this system, typically above 200 seconds, whereas the PI controller was able to handle these loading events.

4.3 Dual-Evaporator Vapor Cycle Results

The LQR controller performs well on the simple vapor cycle system, with its primary drawback being its inability to handle long duration loading events, due to the feed-forward terms introduced in the residualization step performed during model reduction. The relationship between how long and how large of a loading event a controller can handle is unclear, and perhaps this limit is best found via simulation case studies, but it is clear that the LQR controller outperforms the PI controller over some envelope near the operating point. However, it is of great importance that the system work over a wide range of operating conditions, including variations in other terms such as the glycol-water temperature. As more complexity is added to the controller design, to capture more of the system dynamics, the controller inherently becomes more vulnerable to modeling errors. For this reason, the process of controller design was repeated for a larger system. The results of this study are shared in this section, and are meant to give an indication of how the LQR approach scales with the number of components in the vapor cycle, when working in the ATTMO simulation environment.

The larger vapor cycle is based on AFRL's VCSRF system, with two evaporators experiencing differing loads and a liquid injection valve feeding some of the subcooled liquid from the condenser outlet to the compressor inlet. This valve introduces an additional degree of freedom to the system, allowing for reduction in the system's superheat, to keep the compressor from operating over regions where it is inefficient.

The heat load profile for the Dual-Evaporator system is much longer and more varied than the load used in the previous section. Figure 14 shows the load for both evaporators, with the profile being taken from previous work by AFRL on studying VCS control strategies. The load is centered on 9 kW, and the system begins at the operating point x^*, with control u^*. At 900 seconds, the first pulse begins, with one load increasing by 1.8 kW and the other decreasing by the same amount. This asymmetry in loading requires the two evaporator valves to take different actions, and is meant to evaluate each controller's ability to handle uneven loads in a multi-evaporator system. Each step in the profile lasts 900 seconds, which gave sufficient time for the system to reach steady-state after each transition. At 1800 seconds, the loads reach their extreme at 12 kW and 6 kW, before beginning their return to 9 kW. The same pattern is then repeated at 3600 seconds, with the load profiles being flipped. The response to a $\pm$ 3 kW deviation was evaluated from 8,100-10,8000 seconds. The final stage of the profile has both evaporator loads being turned up to the maximum of 12 kW, to evaluate whether or not this maximum disturbance introduces any steady-state errors in the closed-loop systems.

The Dual-Evaporator VCS is then simulated under this loading event in closed-loop with each of the controllers, and output performance is compared in the following figures. The first set of controlled outputs is shown in Fig. 15, containing plots of the SST and SDT.

From the two plots in Fig. 15, we see that both controllers keep the signals well within the target envelopes. The PI controller's SST trajectory has two notable spikes at 11,700 and 12,600 seconds, corresponding to the times when the evaporator loads are pushed up to 12 kW, but the amplitude of these spikes is not large enough to cause any problems. Examining the LQR trajectories, we see the controller never allows for any large spikes, but there is a significant amount of drift

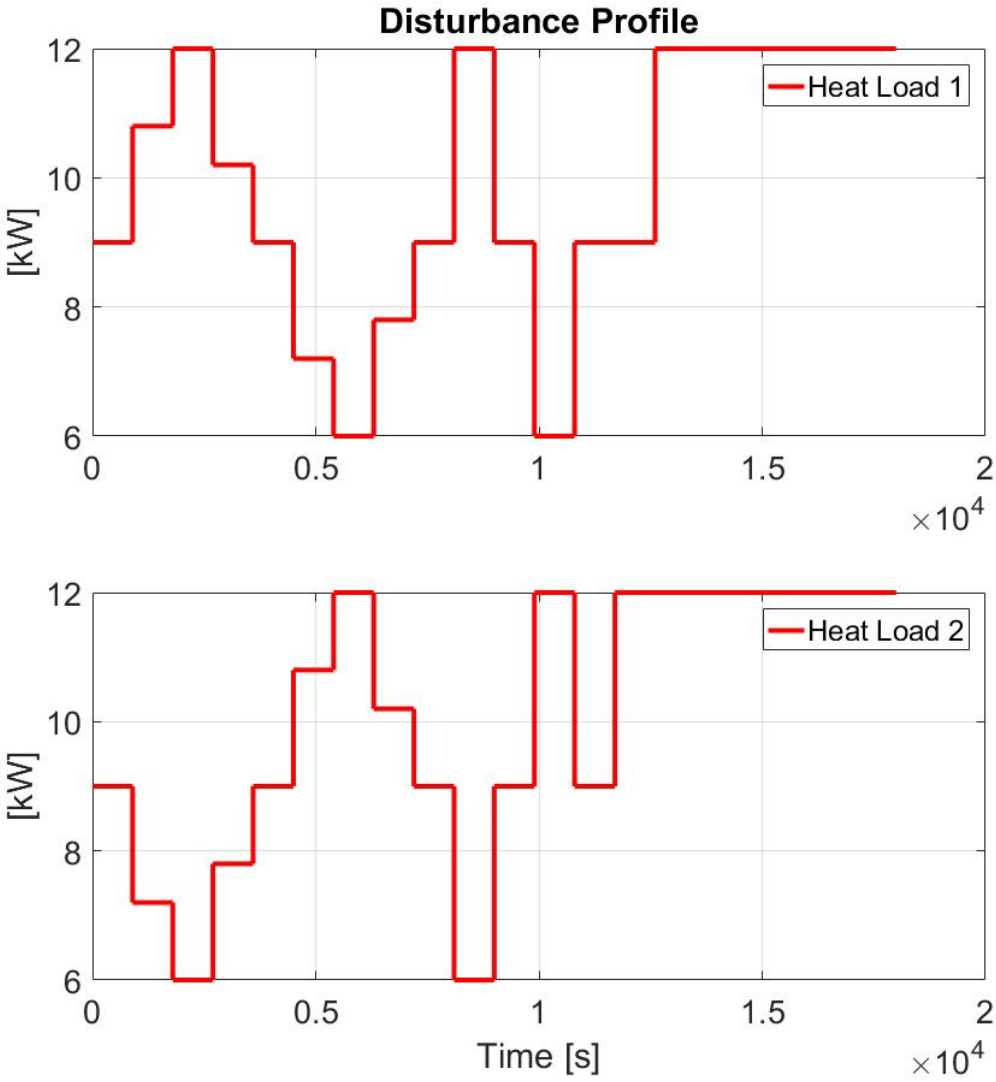

Figure 14. Dual-Evaporator: Evaporator Heat Load Profile

and steady state error present in the trajectory. This is most apparent in the SST trajectory, where the controller has a steady state error of almost 0.5^oF when both heat loads are at 12 kW. This is chiefly due to model differences between the linear and nonlinear system, and represents a limitation of the LQR approach. So long as there are nonzero disturbances acting on the system, there will be some nonzero steady state error in the outputs. For this particular system, these steady state

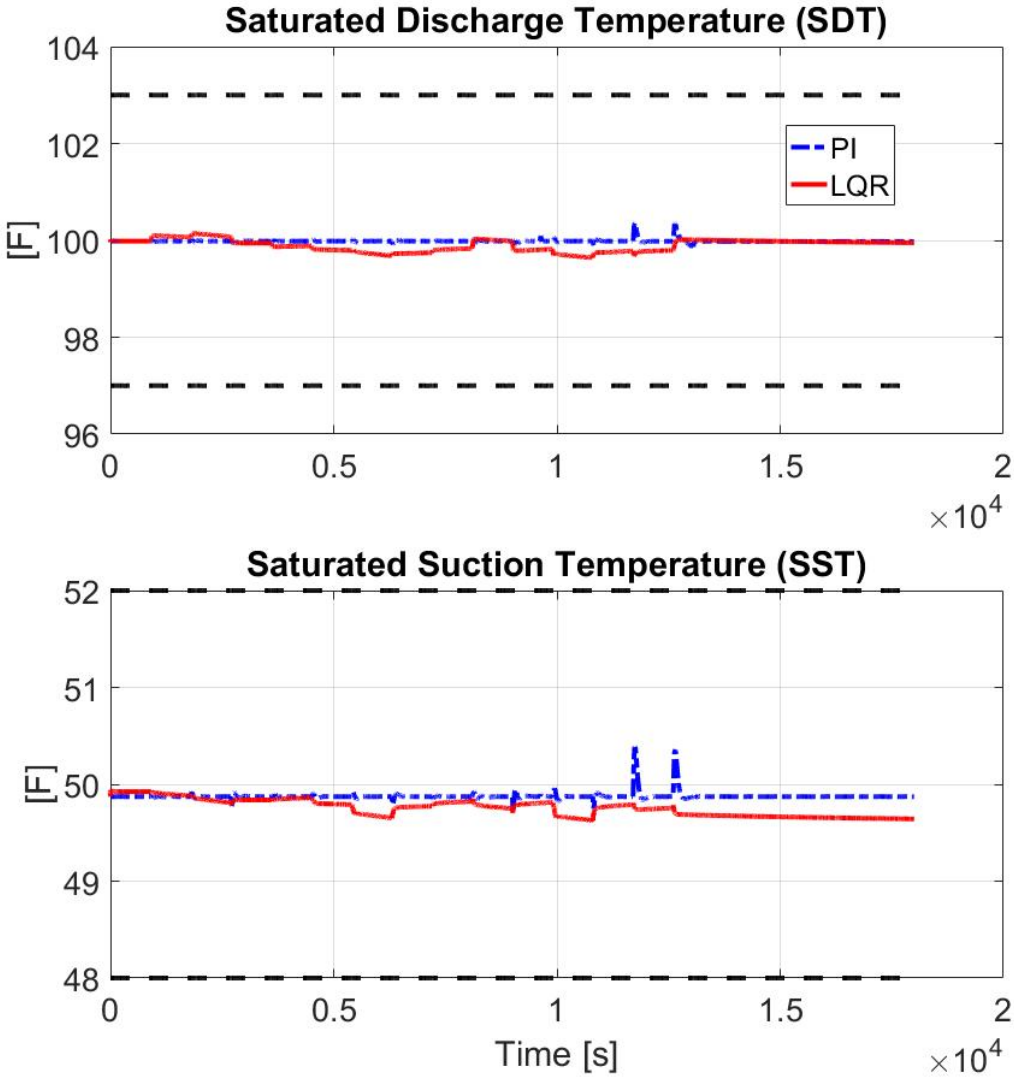

Figure 15. Dual-Evaporator: Evaporator and Condenser Temperatures

errors are within the target envelope, but larger disturbances may drive the LQR controller outside of the envelope in cases where the PI system would not fail.

Figure 16 shows two more controlled outputs used for feedback. These are the temperatures of the PAO loops where the heat actually enters the system, driving each evaporator. At the first heat pulse at 900 seconds, the dashed-blue line experiences a larger deviation from the setpoint, compared to the solid-red line. However, the PI controller quickly brings this signal back to the desired setpoint,

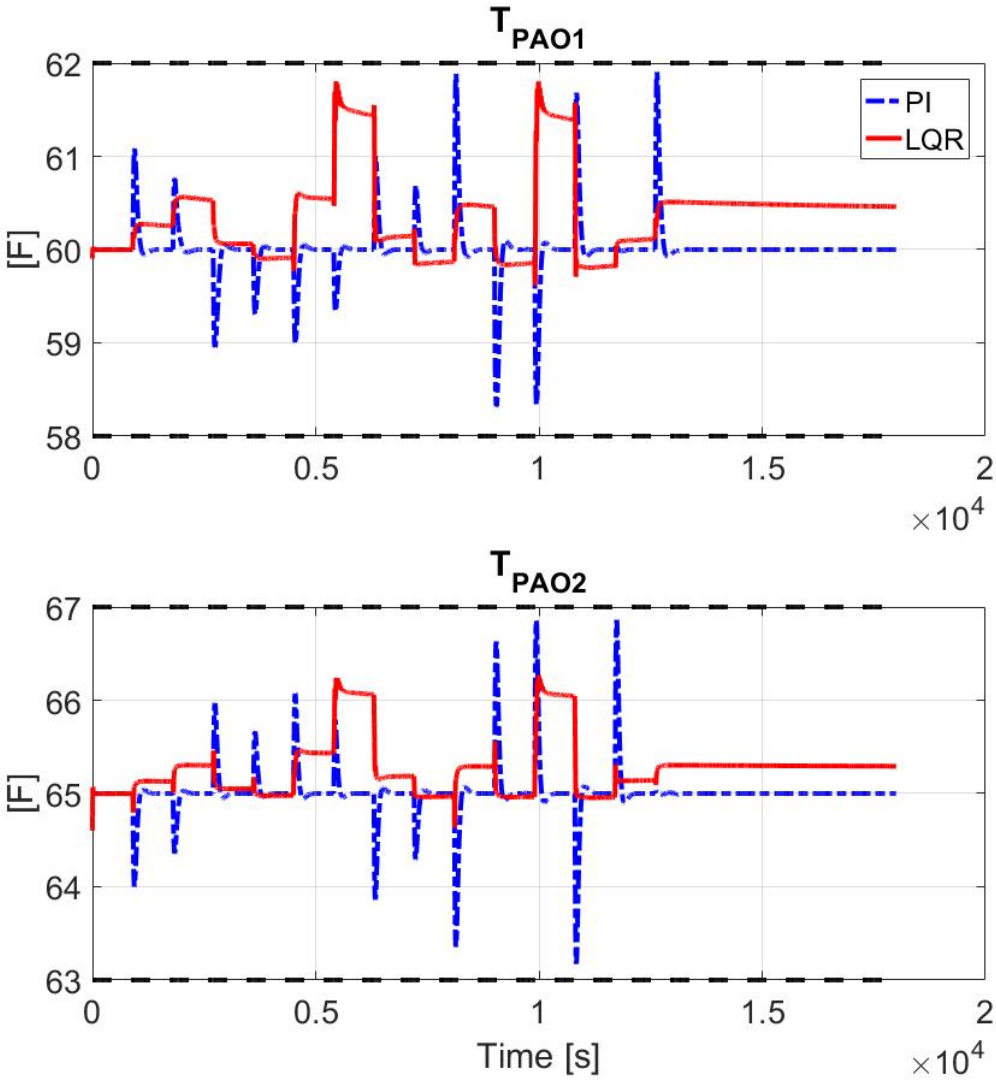

Figure 16. Dual-Evaporator: Evaporator Load Temperatures

where as the LQR controller is unable to drive the steady state error in the red line to zero. Every loading event throughout the profile has this feature, where the LQR controller has non-zero steady state error, but the PI controller allows a sharper initial deviation from the setpoint. At the end of the profile, the LQR controller has about 0.4^oF and 0.25^oF of steady-state error in the two signals. The steady-state error is likely caused by modeling errors, but there are two likely candidates for the source of the modeling error. The first is a modeling error caused by differences

between the nonlinear system and the linearization at the operating point. Unfortunately, this source of error is difficult to mitigate in the framework of linear control theory. However, it is also possible that the steady-state error is caused by the feed-forward terms introduced in the residualization step of model reduction. The source of the steady-state error was not identified in our research, but if it is due to the residualization step, there may be improvements possible by utilizing different techniques from model reduction or estimation, to address the feed-forward term.

Another noteworthy feature of these plots are the peaks in the solid-red curves, occurring around 5,400 and 9,900 seconds. Note that aside from these two peaks, the solid-red curves stay quite close to the operating point. The source behind this issue is uncertain, especially because there isn't anything particularly special about the disturbance driving the system when these peaks occur. If anything, we might expect these peaks to occur at the end of the simulation, when the heat loads are highest, as opposed to the middle of the simulation. However, if we examine the spikes closely, we see some high frequency oscillations occurring just after the spike. Examining the LQR controller's estimate of the heat load in Fig. 17, we see a similar high-frequency feature occurring at 5,400 and 9,900 seconds.

The solid lines in Fig. 17 show the actual heat load, with the dashed lines showing the Kalman filter's estimate of the heat load, based on the feedback measurements. We see that the early estimates are quite accurate, along with the estimates during the final loading event. However, there are two periods of very poor estimates, roughly matching the spikes in Fig. 16 at 5,400 and 9,900 seconds. It is reassuring that the LQR controller is still able to maintain the signal within the target envelope during these periods of poor heat estimation, but this correlation does suggest that the current limiting factor in our architecture is the

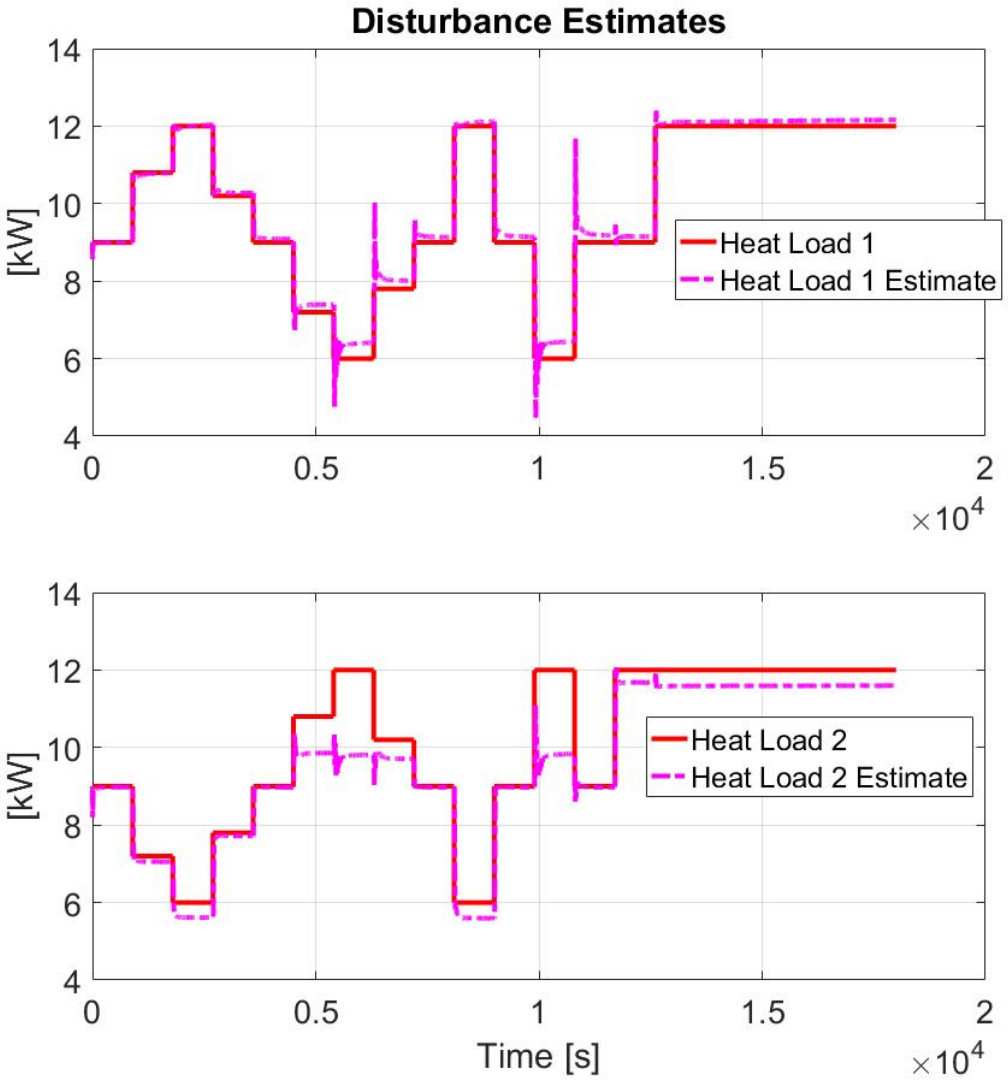

Figure 17. LQR Disturbance Estimates

estimator. Again, this issue is likely attributable to modeling errors from either the nonlinear system, or the feed-forward term.

Figure 18 shows plots for the superheat and subcooling during the case study. Note that the evaporator superheat is being used for feedback, so it can be thought of as "controlled", whereas the condenser subcooling is left uncontrolled. Both of these signals have hard constraints associated with them, with the 5^oF minimum superheat and 10^oF minimum subcooling shown as solid black lines. From the

bottom plot, we see that the amount of subcooling present in the system is well above the constraint, which is why the signal was not included in the feedback path. Additionally, control of subcooling in a vapor cycle is challenging with the current architecture, and might require the addition of additional components to the cycle.

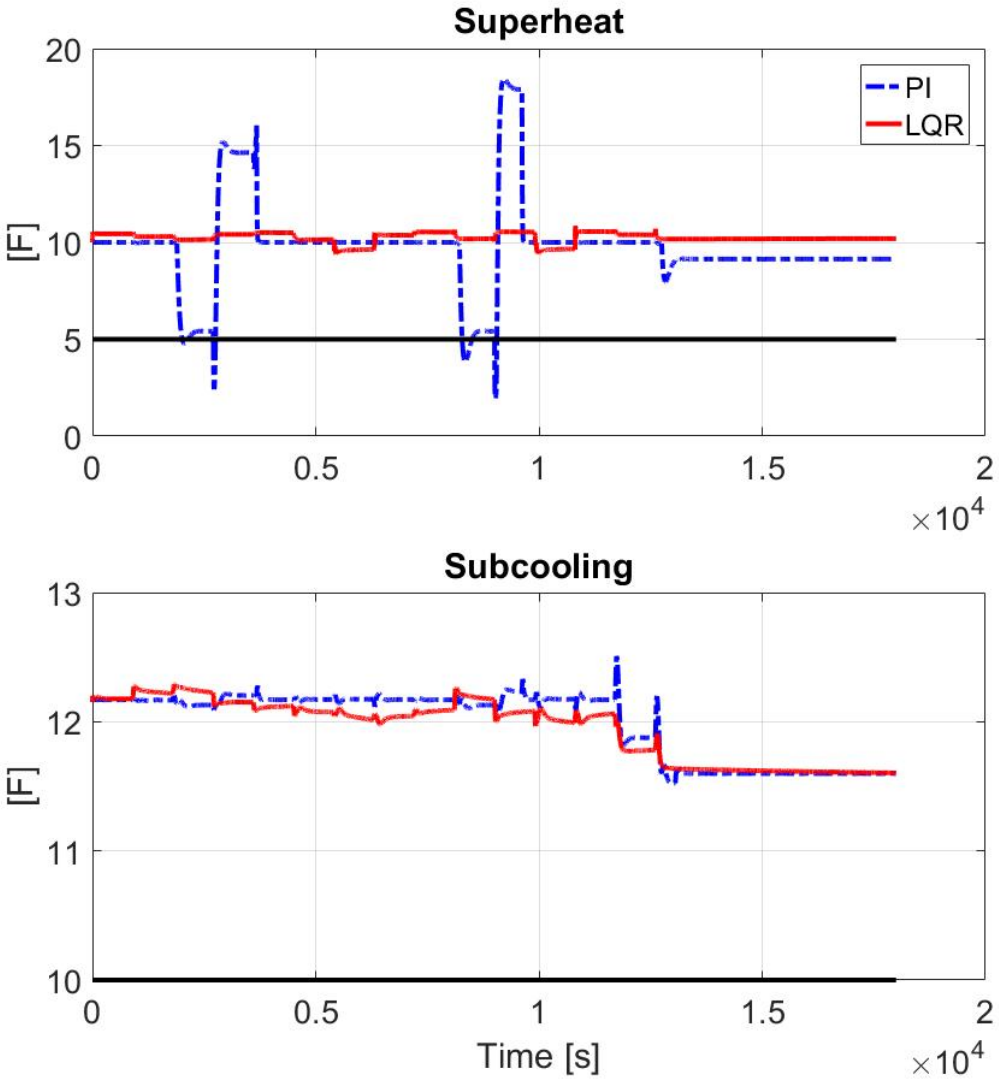

Figure 18. Dual-Evaporator: Superheat and Subcooling

On the other hand, the upper plot of superheat shows an important difference between the PI and LQR systems. In many real-world VCS control architectures, the evaporator valve directly controls evaporator superheat, due to the strong

coupling between these two signals. However, our PI system only uses the liquid injection valve to control superheat, and is severely limited in its control authority. This results in the relatively poor performance shown in the figure, compared to the LQR performance. This issue can be addressed in the context of a PI architecture by simply reassigning the sensor-actuator pairings to match a valve with the superheat. However, this comes with the trade-off of sacrificing control of one signal in exchange for another. This is a fundamental drawback of the PI architecture which directly pairs actuators with outputs, compared to the LQR architecture which uses all of the actuators in an attempt to control all of the outputs.

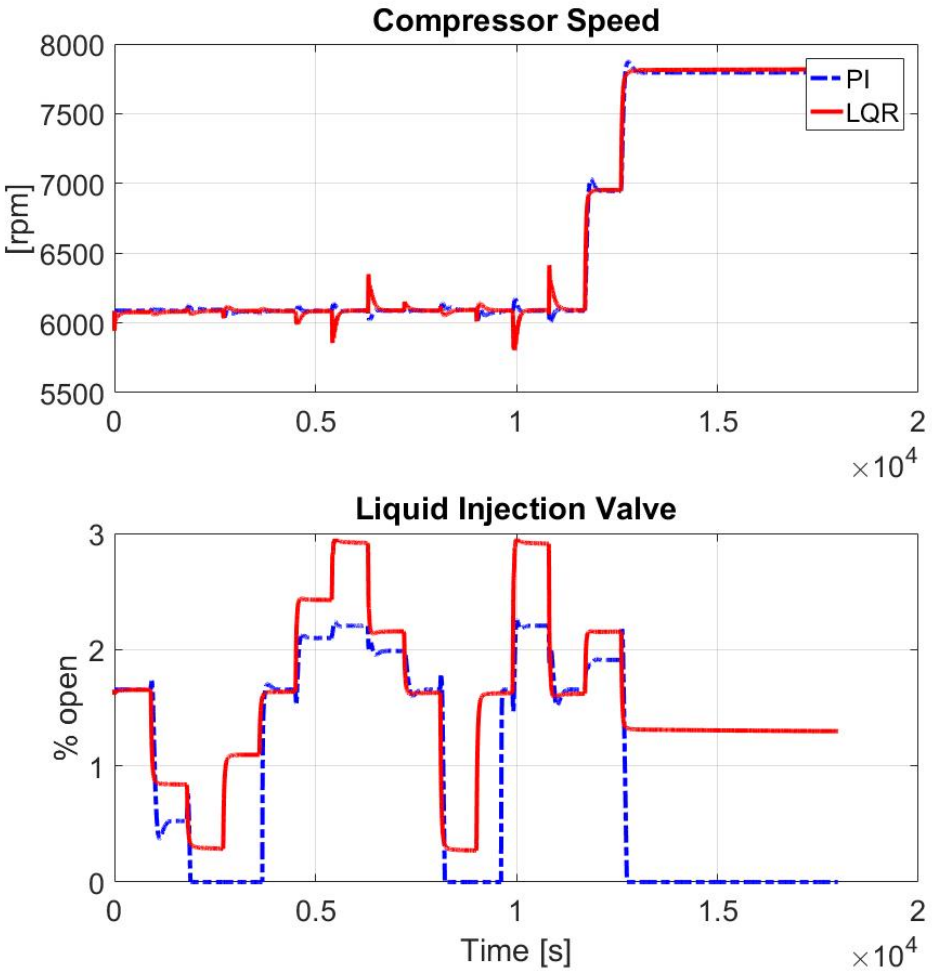

Figure 19. Dual-Evaporator: Compressor and Liquid Injection Valve Control

The last set of plots in this Chapter show the control actions taking place in the four actuators. Fig. 19 shows the compressor speed and the liquid injection valve. In the top plot, we see an interesting pattern where the LQR controller often does the opposite of what the PI controller does. This is due to the fact that the PI controller is only using the compressor speed to control the SST, whereas the LQR controller is using the compressor to control all of the feedback signals. Looking at the bottom plot in Fig. 19, we see that the LQR controller does not saturate the use of the liquid injection valve, whereas the PI controller turns off the valve three times, during periods of very high evaporator heat loading.

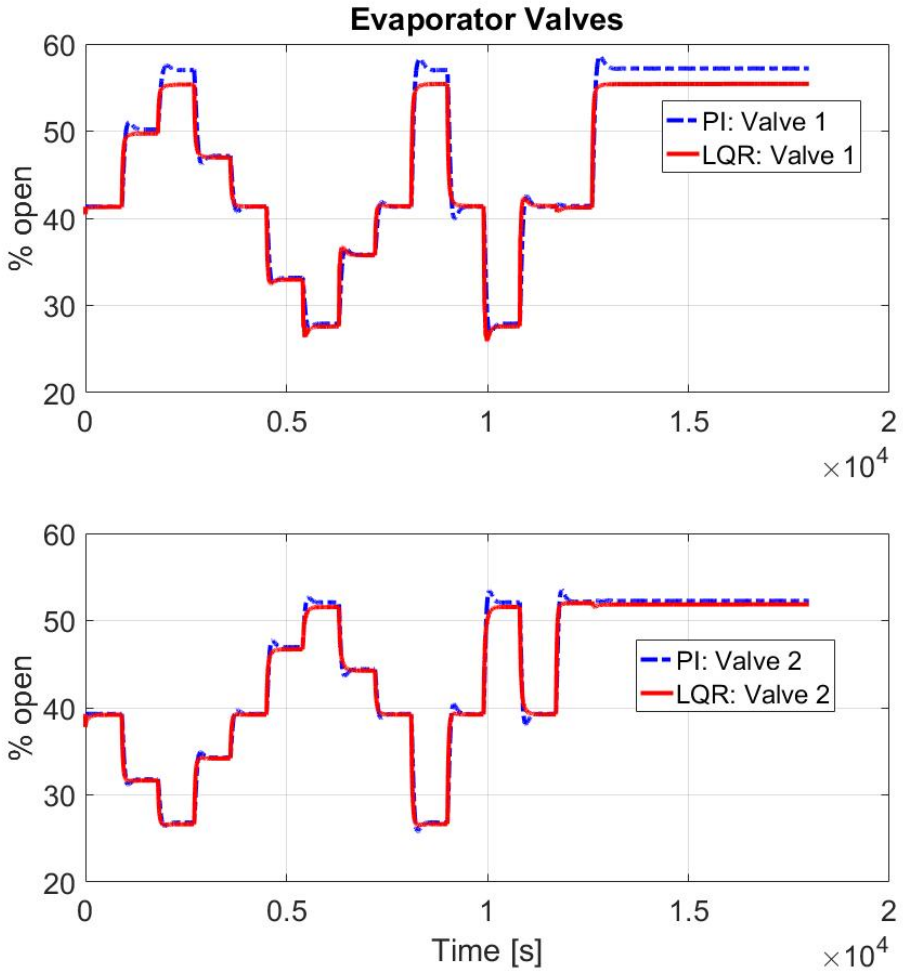

Figure 20. Dual-Evaporator: Evaporator Valves Control

Fig. 20 shows the evaporator valve opening for each of the evaporators in the system. Comparing the curves in these two plots, we see that the use of the valves controlling flow into each evaporator follow essentially the same paths. The only notable difference is that the PI controller is tuned to be slightly more aggressive with the maximum amount of valve opening, as can be seen by portions of the curves where the dashed-blue line is higher than the solid-red line.

Lastly, we have the glycol-water flow rate shown in Fig. 21. This figure exhibits similar trends to the compressor action, with the LQR controller doing the opposite of what the PI controller, again displaying a key difference between the two control approaches.

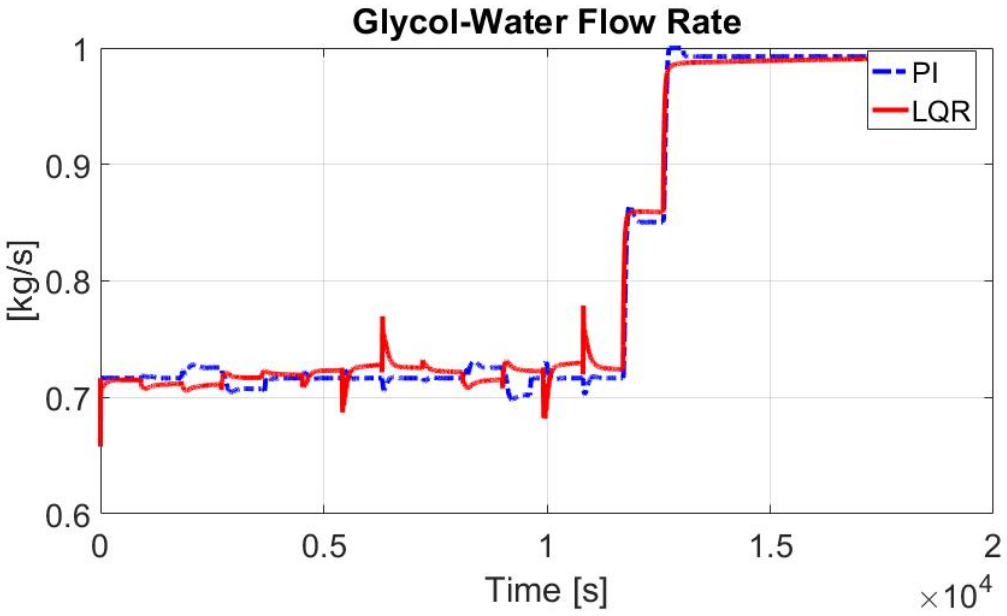

Figure 21. Dual-Evaporator: Glycol-Water Control

From these case studies, we've seen that the advantages of using LQR over PI are not so clear cut. The LQR controller manages to respect all of the constraints set by the controller specifications, whereas the PI controller fails to maintain the minimum level of superheat. However, the PI controller does perform much better at rejecting the steady-state errors in most of the outputs. The LQR method struggles

with steady-state errors, which may be caused by modeling errors introduced by the nonlinear system, or by the feed-forward term introduced during residualization. With that said, the issues seen with the LQR system may be overcome by simply adopting a more appropriate model reduction or estimation scheme.

V. Conclusions and Recommendations

5.1 Conclusions

This thesis presented the results of designing a model-based controller for a vapor compression cycle, in simulation. Linearization of the nonlinear model was discussed, including details on model reduction and validation of the similarity between the models. The process for designing LQR and PI controllers was shared, with a focus on the elements associated with designing the former. These controllers were first designed for a simple vapor cycle, with the minimum four components, and performance was evaluated in the presence of a 50 second pulse in the heat load. In this case study, the LQR controller outperformed the PI controller in regulating the three feedback signals. Additionally, the PI controller was unable to maintain superheat, which would have led to hardware failure in a real system. From these results, it is clear that the LQR method has the potential to outperform traditional PI methods in rejecting evaporator heat loads.

To study how the LQR approach works for larger system, the same process was applied to a larger vapor cycle system, with an added evaporator and a liquid injection valve. The linearization process and the model reduction steps followed the same path as in the previous study with the simple vapor cycle. However, it was more challenging to find a minimum number of states when performing model reduction, due to the fact that the system is larger and more states are required to capture the salient dynamics. For even larger vapor cycles, this step may be the limiting factor and it's likely that other approaches may be required if one aims to design a MIMO controller.

Comparing the results in the Dual-Evaporator case study, the PI controller performed well across all of the outputs, except in maintaining minimum levels of

superheat. The LQR controller performed much better at maintaining superheat, due to the fact that it was able to use the evaporator valves to do so, and it maintained all of the other signals within the target envelopes. However, the LQR controller did exhibit significant amounts of steady state error at the end of the simulation case study, where both evaporator heat loads were set to their maximum amounts. This is due to differences between the linear and nonlinear model, and presents a shortcoming of the LQR approach.

Based on our analysis of these two case studies, we can make the following conclusions:

- For the Simple VCS, the LQR controller significantly outperforms the PI controller. Tuning the PI controller to overcome the performance issues may be possible, but the process for finding suitable PI gains is often arduous and without guarantees of success. Using model-based control design methods like LQR is therefore a viable method for improving controller performance in VCS.

- In the Dual-Evaporator VCS case study, the LQR method fell short of the PI controller's performance in all categories except for superheat regulation. These performance issues are likely rooted in modeling errors, and may be overcome with better estimation methods for estimating unmeasured disturbances.

- Our results are contingent upon accurate VCS models. For this controller design pathway to reach its full potential, there must be significant confidence in the ATTMO toolbox's ability to capture all of the cross-couplings between the different inputs and outputs in the system. Without accurate modeling of these couplings, MIMO controllers will suffer performance degradation.

Indeed, the benefits of using most classes of model-based controllers, which extend far beyond simple LQR approaches, are contingent upon reliable models.

- One challenge in designing an LQR controller for the Dual-Evaporator VCS was performing model reduction. There is no reason to expect that this task will get easier as more components are added to the system. Therefore, we must develop improved methods for performing the model reduction step, while balancing the competing traits of model size and accuracy.

5.2 Recommendations for Future Work

In this research the LQR approach showed promise in the simple vapor cycle case, but adding a second evaporator already has begun to introduce challenges to the method. In order to realize the full potential of vapor cycles in military aircraft, the control system must be robust over a wide range of operating conditions, and the results presented here already show signs of performance degradation off of the operating point. This final section of the thesis will present some recommendations for future research, to address what we identify as the key challenges associated with realizing the full potential of vapor cycle systems for military aircraft.

In order for LQR controllers to see use in real vapor cycle systems, with many evaporators operating across a wide range of operating conditions, the process for extracting a tractable linear model must be improved. Sources of nonlinearity in the system should be removed as much as possible, both at the level of hardware design, and during the process of setting up the model linearization. Much of this research was characterized by a mindset of trying to "linearize everything", and there are likely great improvements possible if only a subset of the vapor cycle dynamics were to be used in the controller design. Indeed, for very large systems, some of the

literature suggests that centralized MIMO control is less effective than distributed forms of control, which do not lean so heavily on accurate linear models of large nonlinear systems.

The task of finding "good" values for the weighting matrices in the cost functional is relatively simple, with satisfactory results found after only a few educated guesses. There is perhaps some room here for performance improvement by automating the tuning process, in closed loop feedback with the linearized model. Some work on this has already been done and shown promising results in the field of model predictive control [10], and it is not unreasonable to expect similar performance improvements in the LQR approach.

Lastly, and perhaps most importantly, is the question of how to achieve control over a wide region of the nonlinear system's operating conditions. Without resorting to nonlinear control methods or model predictive control, it may be possible to extend the results of LQR by adopting a gain-scheduling approach, where multiple linear models are used to capture the variations in the system dynamics. This approach has been investigated by a number of other researchers, but the method is almost always characterized by large amounts of effort required to design the many sets of linear controllers. Additionally, finding efficient scheduling variables to ensure smooth control hand-offs between the different controllers is still a challenge associated with gain-scheduling.

Bibliography

[1] J. Doty, K. Yerkes, L. Byrd, J. Murthy, A. Alleyne, M. Wolff, S. Heister, and T. S. Fisher, "Dynamic Thermal Management for Aerospace Technology: Review and Outlook," *AIAA Journal of Thermophysics and Heat Transfer*, vol. 31, no. 1, pp. 86–98, 2017.

[2] T. Michalak, S. Emo, and J. Ervin, "Control strategy for aircraft vapor compression system operation," *International Journal of Refrigeration*, vol. 48, pp. 10–18, 2014.

[3] S. Engelking and H. Kruse, "Development of Air Cycle Technology for Transport Refrigeration," *International Rfrigeration and Air Conditioning Conference*, pp. 348–356, 1996.

[4] B. P. Rasmussen and A. G. Alleyne, "Dynamic modeling and advanced control of air conditioning and refrigeration systems," *Air Conditioning and Refrigeration Center*, 2006.

[5] D. T. Pollock, Z. Yang, J. T. Wen, Y. Peles, and M. K. Jensen, "Modeling and Control of Single and Multiple Evaporator Vapor Compression Cycles for Electronics Cooling," *American Control Conference*, 2013.

[6] X.-D. He, S. Liu, H. H. Asada, and H. Itoh, "Multivariable Control of Vapor Compression Systems," *HVAC&R Research*, 1998.

[7] M. Anderson, M. Buehner, P. Young, D. Hittle, C. Anderson, J. Tu, and D. Hodgson, "MIMO robust control for HVAC systems," *IEEE Transactions on Control Systems Technology*, vol. 16, no. 3, pp. 475–483, 2008.

[8] R. Shah, B. P. Rasmussen, and A. G. Alleyne, "Application of a multivariable adaptive control strategy to automotive air conditioning systems," *International Journal of Adaptive Control and Signal Processing*, vol. 18, pp. 199–221, 2004.

[9] D. T. Pollock, M. A. Williams, and B. M. Hencey, "Model predictive control of temperature-sensitive and transient loads in aircraft vapor compression systems," *American Control Conference*, pp. 575–580, 2016.

[10] A. Ngo, S. S. Patnaik, B. M. Hencey, and J. R. Cory, "A Simulink Pathway for Model-Based Control of Vapor Compression Cycles," *ASME 2015 Dynamic Systems and Control Conference*, 2015.

[11] M. Kania, J. Koeln, A. Alleyne, K. McCarthy, N. Wu, and S. Patnaik, "A Dynamic Modeling Toolbox for Air Vehicle Vapor Cycle Systems," *SAE International*, 2012.

[12] S. Kakac and B. Bon, "A Review of two-phase flow dynamic instabilities in tube boiling systems," *International Journal of Heat and Mass Transfer*, vol. 51, 2008.

[13] G. L. Wedekind, B. L. Bhatt, and B. T. Beck, "A system mean void fraction model for predicting various transient phenomena associated with two-phase evaporating and condensing flows," *International Journal of Multiphase Flow*, vol. 4, pp. 97–114, 1978.

[14] T. L. McKinley and A. G. Alleyne, "An advanced nonlinear switched heat exchanger model for vapor compression cycles using the moving-boundary method," *International Journal of Refrigeration*, vol. 31, pp. 1253–1264, 2008.

[15] N. Jain, S. Sundaram, and A. G. Alleyne, "Stability analysis for decentralized control of multi- evaporator vapor-compression cycle systems," *Decision and Control (CDC) IEEE 51st Annual Conference*, 2012.

[16] B. P. Rasmussen and A. G. Alleyne, "Gain scheduled control of an air conditioning system using the Youla parameterization," *IEEE Transactions on Control Systems Technology*, vol. 18, no. 5, pp. 1216–1225, 2010.

[17] B. P. Rasmussen, "Dynamic modeling for vapor compression systems-Part I: Literature review," *HVAC&R Research*, vol. 18, no. 5, pp. 934–955, 2011.

[18] X.-D. He and H. H. Asada, "Heating, Ventilation and Air Conditioning (HVAC) System and Method using Feedback Linearization," *US Patent No. 7,076,962 B2*, 2006.

[19] N. Jain and A. G. Alleyne, "Thermodynamics-based optimization and control of vapor-compression cycle operation: Optimization criteria," *2011 American Control Conference*, pp. 1352–1357, 2011.

[20] B. Li, "Dynamic Modeling and Control of Vapor Compression Cycle Systems with Shut-Down and Start-Up Operations," *Thesis, University of Illinois at Urbana-Champaign*, 2009.

[21] M. C. Keir and A. G. Alleyne, "Feedback structures for vapor compression cycle systems," *American Control Conference*, pp. 5052–5058, 2007.

[22] J. P. Koeln and A. G. Alleyne, "Scalable Model Predictive Control for Multi-Evaporator Vapor Compression Systems," *American Control Conference*, pp. 392–397, 2014.

[23] M. Gräber, C. Kirches, J. P. Schloder, and W. Tegethoff, "Nonlinear model predictive control of a vapor compression cycle based on first principle models," *7th Vienna International Conference on Mathematical Modelling*, 2013.

[24] D. T. Pollock, Z. Yang, J. T. Wen, Y. Peles, and M. K. Jensen, "Model-based control of vapor compression cycles for transient heat-flux removal," *International Journal of Heat and Mass Transfer*, vol. 77, pp. 662–683, 2014.

[25] Z. Yang, D. T. Pollock, and J. T. Wen, "Multivariable Control of Vapor Compression Cycle with Transient Heat Flux," *ASME International Mechanical Engineering Congress and Exposition*, 2013.

[26] D. T. Pollock, "Inner / Outer-Loop Control Methodolgy for Mulitple Evaporator Dryout Avoidance Under Transient Heat Loads," *Dissertation, Rensselaer Polytechnic Institute*, 2016.

[27] J. T. Wen, D. T. Pollock, and Z. Yang, "Hierarchical Systems Level Thermal Management for Multiple High Transient Heat Loads," 2017.

[28] E. Walters, M. Amrhein, T. O'Connell, S. Iden, P. Lamm, K. Yerkes, M. Wolff, K. McCarthy, B. Raczkowski, J. Wells, W. Borger, and B. Wampler, "INVENT Modeling, Simulation, Analysis and Optimization," *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, no. January, pp. 1–11, 2010.

[29] M. Bodie, G. Russell, K. McCarthy, E. Lucas, J. Zumberge, and M. Wolff, "Thermal Analysis of an Integrated Aircraft Model," *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, no. January, pp. 1–12, 2010.

[30] J. R. Cory and A. D. Ngo, "A Model Predictive Control Law for a Vapor Compression Cycle System," *ASME Dynamic Systems and Control Conference*, 2016.

[31] J. Burl, *Linear Optimal Control*. Addison Wesley Longman, Inc, 1999.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
21-03-2019	Master's Thesis	Sept 2017 - March 2019

4. TITLE AND SUBTITLE

Control Strategies for Multi-Evaporator Vapor Compression Cycles

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Jackson, Sunderlin D, 2LT
Niedbalski, Nicholas P, Dr
Pachter, Meir N, Dr
Palazotto, Anthony N, Dr

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Air Force Institute of Technology
Graduate School of Engineering and Management (AFIT/EN)
2950 Hobson Way
Wright-Patterson AFB OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER

AFIT-ENY-MS-19-M-221

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

AFRL Aerospace Systems Directorate
POC: Dr. Nicholas Niedbalski
1950 5th St. B20018A RA213
WPAFB, OH 45433-7251
nicholas.niedbalski@us.af.mil 312-986-5548

10. SPONSOR/MONITOR'S ACRONYM(S)

AFRL/RQ

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Distribution Statement A. Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Next-generation military aircraft must be able to handle highly transient thermal loads that exceed the ability of current aircraft thermal subsystems. Vapor compression cycle systems are a particular refrigeration technology that is an attractive solution for dealing with this challenge, due primarily to their high efficiency. However, there are several barriers to realizing the benefits of vapor cycles systems for controlling thermal loads in military aircraft. This thesis focuses on addressing the challenge of controlling vapor cycles in the presence of highly transient evaporator heat loads. Specifically, a linear quadratic regulator (LQR) is designed for a simple vapor cycle system, and closed-loop performance is compared with a set of proportional-integral (PI) controllers. Simulation results show significant advantages of using the LQR method, and the same approach is repeated for a larger dual-evaporator vapor cycle system. The LQR method retains some of its benefits, but several issues associated with relying on a single linear model for the full nonlinear system are identified, and recommendations for future work are made at the end.

15. SUBJECT TERMS

vapor cycle, linear quadratic gaussian, thermal control

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dr. Anthony Palazotto, AFIT/ENY
U	U	U	UU	73	19b. TELEPHONE NUMBER *(Include area code)* 937-255-3636x4599 Anthony.Palazotto@afit.edu

Reset

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18