ADVANCEMENTS IN THE INDUSTRIAL INTERNET OF THINGS FOR ENERGY EFFICIENCY

A Dissertation

by

TREVOR JON TERRILL

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	Bryan Rasmussen
Committee Members,	Jerald Caton
	Juan Carlos Baltazar
	Prabhakar Pagilla
Head of Department,	Andreas Polycarpou

August 2018

Major Subject: Mechanical Engineering

ABSTRACT

The Internet of Things is an emerging field that leverages the connections of everyday objects for the betterment of society. A subfield of this trend, the Industrial Internet of Things (IIoT), has been referred to as an industrial revolution that enhances both productivity and safety in the industrial environment. While still in its early stages, identified improvements have the potential to markedly improve manufacturing productivity.

Energy efficiency within manufacturing plants has traditionally received little focus. The Industrial Assessment Center Program demonstrates the potential energy improvements that can be realized in manufacturing plants, but these assessments also highlight some of the traditional barriers to energy efficiency. Some of these barriers include the lack of data to justify actionable improvements, unclear correlations between improvement costs and potential cost savings, and lack of knowledge on how energy improvements provide ancillary benefits to the plant. The IIoT has the potential to increase energy efficiency implementation in manufacturing plants by addressing these challenges.

This dissertation discusses the framework in which energy efficiency enhancements within the IIoT environment can be realized. The dissertation initially details the potential benefits of IIoT for energy efficiency and presents a general framework for these improvements. While proposed IIoT frameworks vary, they all include the core elements of improved sensing capabilities, enhanced data analysis, and

intelligent actuation. In addition to presenting the framework generally, the dissertation provides detailed case studies on how each of these framework elements lead to improved energy efficiency in manufacturing.

The first case study demonstrates improved sensing capabilities in the IIoT framework. A non-intrusive flow meter for use in compressed air and other fluid systems is presented. The second case study discusses Autonomous Robotic Assessments of Energy, which use advanced data analysis to autonomously perform a lighting energy assessment in facilities. The third case study is then presented on intelligent actuation, which uses a novel k-means algorithm to autonomously determine appropriate times to actuate compressors for air systems in manufacturing plants. Each of the presented case studies includes experimental tests demonstrating their capabilities.

DEDICATION

I would like to dedicate my thesis to my marvelous wife and our four darling kids.

ACKNOWLEDGEMENTS

I would like to thank first and foremost my advisor and committee chair, Dr. Rasmussen for all his support. I would also like to acknowledge the other faculty and staff I have interacted with during my time at Texas A&M.

I would like to thank all those who have helped us have a great experience here. Some in particular I would like to note include Jim Eggebrecht at the Industrial Assessment Center, my fellow lab members at the TFCL, fellow IAC students, and friends from church.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supported by a dissertation committee consisting of Professors Bryan Rasmussen, Jerald Caton, and Prabhakar Pagilla of the Department of Mechanical Engineering and Professor Juan Carlos Baltazar of the Department of Architecture.

Undergraduate students in the AggiE Challenge program assisted with collection and analysis of spectrometer data in Chapter 4. Additionally, students in this program created the integrated prototype system depicted in Figure 61 and Figure 62.

All other work conducted for the dissertation was completed by the student independently.

Funding Sources

Graduate study was supported by the AggiE Challenge program and the Industrial Assessment Center program at Texas A&M University.

TABLE OF CONTENTS

Page

LIST OF FIGURES

Page

LIST OF TABLES

Page

CHAPTER 1

INTRODUCTION TO THE INDUSTRIAL INTERNET OF THINGS FOR ENERGY EFFICIENCY

The Internet of Things (IoT) is a recent emerging field that envisions the connection and utilization of everyday objects for the betterment of society. The global impacts from the IoT are expected to be widespread, with some terming it a technological revolution [1-3]. Recent advances in low cost communication and RFID technologies have enabled the networking of everyday objects, prompting a new field in which all things are connected together in a global network. At the core of the IoT is the embedding of electronics in everyday objects, connecting them to the global infrastructure and making them "smart." The envisioned fields of application are diverse, ranging from personal smart homes to transportation and commercial applications. Xu, He, & Li identify some of the many widespread industries that stand to benefit from IoT, including transportation systems, environmental monitoring, healthcare, traffic monitoring, and industrial systems [4].

The birth of IoT is rooted in the expansion of RFID technology, which enabled the fast, low-cost tracking of things in commercial and industrial environments. The field has grown rapidly, as indicated both in available market solutions and research in the literature [4]. The scope of IoT has widened recently to a broad range of potential applications and available sensors, but the field is still in its infancy. In a survey paper of many IoT applications that are available in market, Perera observes that most

applications are narrow in focus [5]. Other researchers have observed that the field is still widely fragmented, with one-off solutions being developed for single applications [6]. Effort is now being performed to establish uniform frameworks and communication standards to facilitate functionality and communication between devices produced from different companies. The end vision of IoT involves use of information from a global connection of smart devices to adapt the environment for improved human productivity, safety, and convenience. Specifically, these applications will autonomously identify changes in the operating environment and adjust to provide desired service [6].

1.1 Industrial Internet of Things

Many subfields of IoT have been identified, including applications pertaining to the industrial environment. The Industrial Internet of Things (IIoT) is focused on applications in manufacturing environments. The IIoT is also commonly referred to as the Industrial Internet or Industry 4.0. Industry 4.0 refers to the next perceived revolution in industry, with previous revolutions pertaining to mechanization, assembly lines, and system automation.

While the IoT at large is still in its early stages, this is especially true in the IIoT. Perera et. al. explore IIoT solutions that are commercially available in the marketplace, in contrast to solutions that have only been proposed in literature. [7]. They find that many products are available in residential and commercial markets, but very few are available in industrial facilities. Although some applications from the commercial sector could be used in an industrial environment, such as intelligent thermostats and air quality monitoring, these technologies lie on the peripheral of industry production and energy

use. The paper only identified a single supplier of IoT technologies for manufacturing equipment, with no identified suppliers for energy subsystems (such as compressed air systems, motors, or boilers).

More potential applications in the IIoT have been identified and are being researched in the literature. Within these application areas, some of the identified industrial solutions include improved climate control for energy efficiency in buildings, product tracking for enhanced manufacturing capability, and improved safety by monitoring equipment and processes [2]. Many solutions in the IoT rely primarily on RFID or similar tracking methods of inventory or processes [4, 8]. While RFID and other solutions are being explored specifically for the industrial sector, currently-available technologies in the industrial field are notably behind their commercial and residential counterparts [5].

1.1.1 Challenges in IIoT

While there is much discussion on potential productivity improvements in manufacturing through IIoT [9], some have expressed doubt at the ability of the IIoT to significantly improve manufacturing productivity [10]. There are several reasons for this scepticism on the Internet revolution for manufacturing. One illustrative difference comes from comparing the contrasting usage pattern between the industrial sector and the commercial/residential sectors. Residential and commercial buildings operate with similar equipment between buildings and predictable schedules based on building purpose. Manufacturing processes vary widely depending on the type of manufacturing process. These differences include the extent of process automation, product heating and

cooling needs, production operation hours, and support system requirements. This variance creates difficulty in predicting production schedules or energy use trends based solely on the type of industry. Specific solutions are at times unique to a facility and not readily applied to other facilities [11].

Another potential reason for the expressed doubt in IIoT impact is the trend for industrial plants to lag behind consumer and commercial sectors in the adoption of new technology. This is especially true in smaller manufacturers, where concerns over disruptions in the production process overwhelm potential benefits of IIoT solutions. Relatively few potential solutions to date have been identified with impactful potential in industrial facilities, and Stanovic et al. note how there are many potential unidentified solutions due to this hesitation in technology adoption. [12].

Other researchers have highlighted some additional unique challenges in IIoT compared to the IoT for traditional domestic applications [13, 14]. First, IIoT hardware often needs to integrate with other systems in the industrial plant. In contrast, domestic IoT solutions typically entail standalone, vertically integrated solutions. Second, domestic IoT safety considerations primarily focus on protecting user data, while IIoT needs to additionally prevent unauthorized access to equipment. Finally, as faulty operation of IIoT can potentially present safety hazards to human personnel, the technology must operate correctly without continuous access to the internet or connectivity with other smart devices. Faults in many residential or commercial IoT solutions create inconvenience or revenue loss, while faulty IIoT solutions can result in human endangerment. A high-tech analysis firm examined additional necessary

requirements for implementation of IIoT. They found that, when compared to typical IoT solutions, IIoT solutions must be more robust to extreme environments, have longer service lifetimes, be more resilient to failure, and be more secure against security threats [14].

1.1.2 Growth Areas of IIoT

With these challenges, there are a couple of avenues where IIoT is likely to see faster growth. Wang et. al. focuses on the IIoT as a paradigm for safety-critical applications in industrial facilities [13]. Specifically, IIoT applications have recently emerged in environments where equipment failure results in dangerous situations. In these environments, the IIoT has potential for immediate impact in identifying new faults more quickly or even prior to the fault. As opposed to IIoT solutions that enhance production with potential introduced safety risks, focusing on safety critical applications can enhance worker safety, which in turn enhances plant profitability. Wang et. al. specifically discusses safety critical application for the IIoT in the maritime industry [13], but there are many potential safety applications in a wide variety of industries such as chemical plants and mining [4].

A second avenue for early adoption of IIoT is in predictive maintenance and automatic fault detection [15, 16]. Predictive maintenance of equipment reduces unexpected shutdowns, decreases overall plant downtime, and increases productivity. There are additional safety benefits to production workers, as equipment failure and unexpected production line changes create safety risks. Finally, predictive maintenance

can reduce operational costs through reduced equipment replacement or refurbishment costs.

1.1.3 Energy Applications in the Industrial Internet of Things

Much of the literature that discusses energy in the framework of IoT and Big Data solutions reference Smart Grid concepts as the primary research area [17, 18]. While much research is necessary on the supply and distribution side of smart grids, there are plethora of opportunities on the demand side beyond facility load forecasting. To date, IIoT solutions that optimize energy use in manufacturing facilities have received limited attention in the literature.

Improved energy efficiency of facility energy systems will reduce facility demand and consumption, providing benefits to the demand side of the smart grid. The Industrial Assessment Center (IAC) program performs energy assessments of regional small- to medium-sized manufacturers. These assessments identify potential energy efficiency improvements for the manufacturer. The impact of the national program is well documented in both decreased utility costs and student training [19].

The IAC program demonstrates potential energy savings that may be achieved in manufacturing facilities, but it also highlights some of the reasons why energy efficiency in industry has received little attention in the literature. During energy assessments, there is little available time for data collection or in-depth study of the facility's energy use. A survey of production managers, including some who were recipients of an IAC visit, reveal that retrofit recommendations are considered strongest when there are other ancillary benefits beyond reduced energy costs [20]. Energy efficiency investments are

often dropped at the management level because the endeavour is seen as unprofitable. The surveys revealed that often energy efficiency is not the principal driver of capital projects, but it is typically only considered in the context of other plant improvements.

Production output and employee safety entail the primary foci of production managers, as these factors directly affects the bottom line of the company. Jiang el al. note that research performed in Big Data for energy applications all have a common purpose in cost minimization [17]. With a lack of knowledge prior to and following an energy efficiency upgrade, the only metric to gauge potential impacts of the upgrade is through utility bills. The large variations in energy usage on a daily and monthly usage make discerning energy savings from a specific project using utility bills impractical.

Historically, the potential cost savings from energy reductions are the primary point that is emphasized in the adoption of energy-efficient technologies. These cost savings reduce overall operating expenses for manufacturers and contribute to the realized profits for manufacturers. However, one identified challenge in deploying Energy Efficiency Measures (EEMs) for energy efficiency is the initial cost to implement the EEM. While many EEMs benefit from a very small payback period [19], the initial cost of the consultation by experts and the cost of equipment can be prohibitive for the implementation of the technology [21].

1.1.3.1 Marrying Energy Efficiency with Safety and Fault Detection in the IIoT

A proposed approach to improve implementation of energy efficiency of industrial facilities in the IIoT framework is to focus on solutions that either offer energy efficiency as an ancillary benefit on major energy support systems or else improve the

process of performing system diagnostics. These solutions, which increase production or decrease downtime, are more likely to be implemented.

Additionally, early IIoT solutions are likely found where solutions can be uniformly applied to a wide variety of industrial plants, provide increased worker safety, or incorporate machine fault detection. Energy support systems, such as boilers, compressed air and electric motors, are commonly found in most plants. Solutions that apply to these support systems can be readily implemented across many plants with little or no modification.

As additional sensing becomes inexpensive and readily available, the cost decreases for health monitoring of equipment, optimizing production performance, or ensuring correct system operation. Lee et. al. have identified health monitoring and management of equipment as a key application area in the industrial environment. Proper equipment health monitoring can detect issues and potential failures prior to human operators, resulting in the avoidance of potential costly downtimes on the production floor. The researchers specifically identify fault detection and diagnosis as a key application area in equipment health monitoring [15]. These same data obtained from health equipment monitoring may be utilized to simultaneously make energy efficiency improvements to the system. Although EEMs are economically justifiable by themselves with a relatively low payback period, coupling energy efficiency solutions with fault detection or safety improvements provides increased incentive for managers to invest in these upgrades. These implemented projects will still benefit from a small

payback period, which reduces overall plant operational costs while still providing the benefits of equipment monitoring.

Many of the safety-related faults in industrial systems initially start with soft faults prior to a hard fault. Soft faults reduce system efficiency but still meet operational requirements. Identifying and correcting soft system faults prevent hard faults of the system, generating additional energy and cost savings to manufacturers.

1.2 Internet of Things Framework

The IoT as a research field is still in its infancy in regards to standardization and frameworks. There are many proposed frameworks and solution approaches to IoT in the literature, many of which are summarized by Sundmaeker et al. [2]. This lack of a uniform framework creates challenges in designing IoT solutions. However, most all the frameworks contain several key components the constitute the core of an IoT solution [1-3, 6, 17, 18]. The referenced frameworks are not a comprehensive list of identified frameworks. Instead, they provide a sample of proposed approaches to illustrate that a large number of researchers have proposed similar approaches in this emerging field.

The key components identified by most frameworks are visualized in Figure 1. The foundational three pillars of IoT are 1) the sensing layer in which the data on the physical world are collected, 2) the network layer in which data are transmitted and processed, and 3) the application layer in which the data are used to accomplish some useful purpose [1, 8].

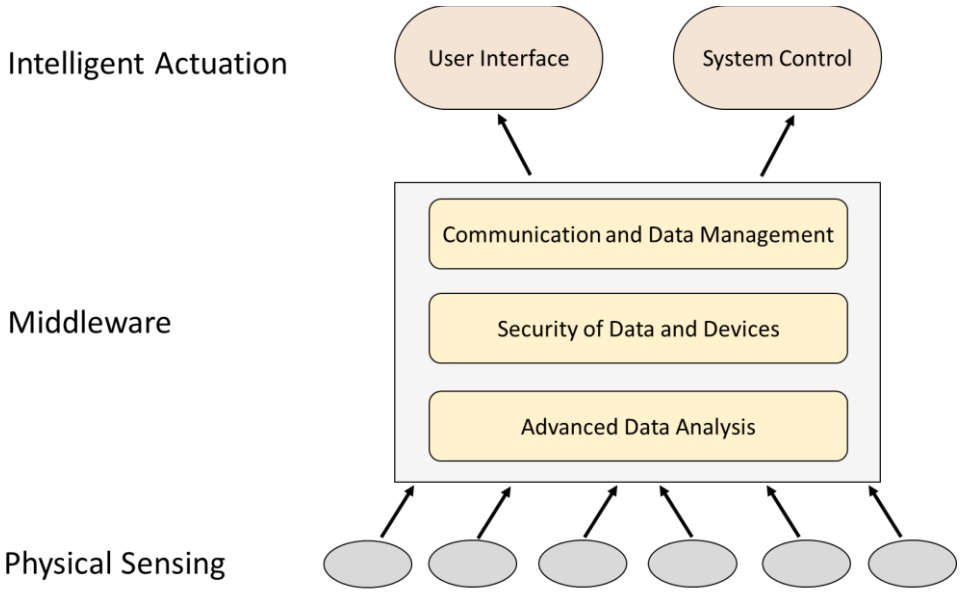

Figure 1. General framework for realizing the Internet of Things.

As mentioned earlier, many potential application areas have been identified for future IoT solutions. While the application areas range from consumer trends, transportation, intelligent homes, and industrial applications, most every proposed framework includes these core areas of sensing, middleware, and actuation. Within middleware, most frameworks include elements of communication, data analysis, and ensuring security of devices and user data. The physical sensing layer interfaces with the world and provides insight on the current state of physical processes. As the sensors have limited communication and processing power, they are connected to a global network through the middleware. All the sensed data are available in the middleware, where greater computing power is available. The middleware has many different terms, but it typically consists of data processing, storage, communication protocols, analysis, and cyber security [22]. Through data analysis, improvements are identified and either

implemented autonomously through actuators or conveyed to operators through a user interface. This actuation or user interface constitutes the actuation layer, where improved decision making is achieved.

The three layers that are identified above are key variables to realizing the full potential of the IoT both generally and specifically to industrial systems. Other variables that must be considered include security among devices, expanding the sensing and communication loops to outside the closed plant, ensuring robustness with aging sensors and processes, and ensuring effective communication protocols between interconnected devices [12]. This dissertation focuses on advances in the Industrial Internet of Things as it applies to energy systems. Specifically, the dissertation focuses on the three elements of data collection, autonomous data analysis, and intelligent control based on the analyzed data. The author recognizes the importance of advancement of communication and cybersecurity technologies in the IIoT, but these topics are not a focus of the presented work. Each of the five topics (Cybersecurity, Communication, Sensing, Data Analysis, and Intelligent Actuation) are discussed in the following sections.

1.2.1 Communication and Cyber Security

One large identified challenge to IoT is the reconciliation of various communications and networking standards [2, 6, 18, 23]. Communication between smart devices and equipment becomes increasingly important in industrial applications. Breakdown in communication between devices can have severe consequences, including loss of plant production or endangerment of human life. This problem is complicated by the large networks of devices and associated distributed communication. The challenge

in further exacerbated by IIoT solutions being developed with proprietary software and algorithms [11]. Multiple researchers and organizations are working to set uniform standards, as this is viewed as a key step to wide acceptance and implementation of IoT solutions, some of them specific to IIoT [4, 24]. Communication standards present a unique challenge in manufacturing systems, where processes and equipment vary so widely that uniform IIoT solutions are nearly impossible to identify. Communication standards are a necessary precursor to widespread IIoT solutions, as communication protocols for manufacturing equipment is likely to be equipment specific.

Issues related to cyber security is an additional challenge that is readily identified by those creating IoT frameworks [2, 6, 17, 18, 23]. While readily identified, limited attention has been given to the topic [5]. There are two regularly cited challenges in regards to cyber security. The first is in the protection of individual consumer data, as sensitive information can cause significant harm to individual users or manufacturers. The second challenge is in the security of the IIoT devices. If security of the IIoT devices is breached, unauthorized users can control specific equipment on the manufacturing floor. This presents monumental economic and safety risks. Even small, imperceptible changes in a manufacturing line can result in final products that do not meet final specifications. The source of the alteration may be difficult to identify, and the resulting loss in productivity and product waste can be economically devastating to manufacturers. While both communication and security challenges are paramount to establishing an effective IIoT framework, these topics are not the focus of the

dissertation. The following chapters focus on technical aspects of data sensing, autonomous data analysis, and intelligent actuation.

1.2.2 Advanced Data Sensing

The field of Big Data uses large volumes of collected data to mine useful correlations and make informed decisions. IIoT relies on Big Data from the physical world to enable production, safety, or energy enhancements. This is only possible if the physical world can be sensed inexpensively so that sufficient data can be captured. Measurements that are currently expensive limit IIoT solutions, as advanced networking and communication can't substitute for a lack of sensing. Insufficient measurement and monitoring has been found to be a key factor that limits implementation of energy-saving technologies and best practices [11, 20].

Researchers have remarked on the increasing availability of sensed data, and they use this assumption of low-cost data when building an IoT infrastructure [4]. Further development in sensing technologies is therefore necessary to realize the full potential of IIoT solutions. While there are many readily-available sensors for use in monitoring equipment and processes, many desirable measurements are difficult to obtain. The desirable measurements may require prohibitively expensive equipment, be unreliable, or possibly can't be obtained with the best available technology. Additionally, many straightforward methods for sensing often have ad hoc additions for specific circumstances to account for uncertainties in estimation of parameters. Advances in technology will drive down costs and make additional measurements possible. This is particularly true for IIoT for energy efficiency. As the primary benefit

of energy efficiency is cost savings to the manufacturer, any deployment of technology with energy efficiency as the primary driver must be low cost.

When sensors for desired measurements are prohibitively expensive or otherwise unavailable, at times the desired measurement can be derived from other measurements. Often termed virtual sensors, physical laws and relationships from measured parameters are used to estimate the desired parameter. Virtual sensors are capable of determining parameters that are not immediately obvious or can't be directly measured with available technology. Use of virtual sensors can at times bring down sensing cost by deriving the desired measurements implicitly from other readily-available, low-cost measurements. Developments in advanced sensors, in particular virtual sensors, is an area of active research in energy systems. The following section describes work done in the industrial field regarding advanced sensors.

1.2.2.1 Examples of Advanced Sensing for Energy Applications

While regular advances are made on lowering cost and improving sensitivity of direct measurement technologies, there is a large body of literature on use of virtual sensors for low-cost advancements in sensor technology. This section presents several virtual sensors being developed or in use in the industrial sector related to energy efficiency. The presented examples are not meant to be an exhaustive list of available and researched sensors in the field. Instead, the purpose is to provide a sampling of other work done in the literature to provide context and relevance for the example case study presented later in Chapter 2 and Chapter 3 of the dissertation.

A common application of virtual sensors is found in fault detection and diagnosis of heating, ventilation, and air conditioning (HVAC) systems. Research has been done for years in the context of both energy efficiency and compressor health monitoring [25, 26]. Inexpensive temperature sensors are the primary sensors used to evaluate the HVAC systems. Virtual sensors, using these temperature measurements, can estimate the HVAC system pressure, refrigerant mass flow rate, and compressor power [27, 28]. The sensors are then used to detect and diagnose soft system faults that degrade system performance [29, 30]. The identification of soft faults helps reduce potential catastrophic faults, reducing unexpected shutdown of the equipment and a potential stop in production.

A second example is illustrated in Chapter 5 on estimation of compressor load factor values using only pressure information. This work is presented in the context of intelligent actuation, but the load factor is a necessary sensed measurement to realize the intelligent control. For reciprocating compressors and compressors operating under a load/unload configuration, the rates of pressure rise and drop can be used to estimation compressor loading. This only requires pressure sensing somewhere in the system to estimate compressor loading conditions.

A final example is found in estimation of motor loading factors and efficiencies in industrial processes. Estimating the load factor and efficiency of motors is a key element of predicting total motor energy use, but directly measuring the motor torque is difficult. Various approaches have been proposed and researched for estimating load factor through derived quantities, each with its own challenges and advantages. Some of

these approaches include using motor slip, motor electrical current, the airgap torque method, and statistical approaches [31, 32].

1.2.3 Autonomous Data Analysis

One challenge to energy efficiency in industrial facilities is the cost to identify and analyze potential retrofit opportunities. For both the external energy auditor and internal energy manager, the labor cost to evaluate potential energy-saving opportunities reduces the economic viability of such opportunities. In general, the increasing amount of data that are available in many fields require additional and automated data analysis techniques. The expected amount of data in the context of IIoT makes effective human analysis infeasible, and sensor data without proper analysis have little value [15, 33]. Large amounts of data on the sensed world provide a potential wealth of information, but advanced analysis must be accompanied with the data to mine useful knowledge.

Data from diverse, low-cost sensors presents a standing challenge on how to marry incompatible data to generate production, safety, or energy improvements in the industrial facility. The middleware in the IIoT framework is where these large datasets are analyzed and useful relationships identified. The middleware consists of all data handling, communication, and data processing. There are several fields of research, including Big Data, that investigate data handling and analysis techniques to derive useful, easily-comprehensible trends from the data. In general, the phase “Big Data to Smart Data” is used to summarize the field of data analysis in the context of IIoT [11]. While strides have been made, researchers emphasize that large research challenges still exist for the IoT. Without additional analysis techniques, improvements to industrial

facilities are difficult to realize, and advanced analysis tools are a necessary element to gain useful insight [1, 15]. A survey paper on Big Data asserts that the processing capacity of Big Data in the IoT field has fallen behind data collection, and development of Big Data analysis techniques should be urgently pursued in order to advance IoT [8]. Include in those listed challenges is automated or sophisticated data analysis techniques to productively analyze large streams of data while still accounting for human variability [8, 33, 34].

Some researchers have applied IoT and Big Data frameworks to increase the energy efficiency in commercial buildings. While work has been done for various energy subsystems to optimize energy use of building subsystems in the context of occupant thermal comfort, Moreno et al. propose a whole building solution to electrical energy efficiency [21]. This work illustrates the difficult nature in defining all the factors that contribute to energy usage, as the complete model setup required manually defined occupancy schedules, thermal zoning, selection of model inputs, and selection of model structure in order to yield the best results. These manually tasks are time intensive and limit widespread applicability of IIoT solutions for energy efficiency.

In addition to unclear relationships in sensed data, the large amount of sensed data can present an additional challenge when data come from unique sources with specific formatting. For example, weather data, utility data, and internal sensor data must all be converted to compatible types for data analysis use [17, 18]. This adds an additional layer of complexity in the data analysis that isn't found in traditional IoT solutions that use a single sensor type.

1.2.3.1 Industry Examples of Autonomous Data Analysis for Energy Applications

A first example of advanced data analysis is in health diagnostics of motors used in production and support processes. Rewound or replaced motors are a considerable cost to manufacturing facilities. Many of the root causes leading to motor failure come from inadequate lubrication in motor bearings [35, 36]. Evaluating motor bearing health to ensure proper lubrication can both increase energy efficiency of motors and decrease expensive motor overhaul costs. One approach to diagnose the health in motors is through a vibrational analysis of the motor [37]. Using low-cost accelerometers, advanced signal analysis is used to correlate the incoming vibration signal with motor bearing health [38].

Other industrial examples can be found in process or system optimization. A common approach in industry for fault diagnosis is using data-driven approaches. While these models require large amount of data, no prior knowledge about the system is necessary. Common data driven approaches include principle component analysis, partial least squares, and independent component analysis [39].

1.2.4 Intelligent Actuation

One of the early challenges seen in Big Data is the difficulty of improving decision making based on the wealth of information available. The data produced and analyzed in the IIoT environment are all done with the goal of improved decision making. If no useful improvements are made to a process or system, then a typical IIoT solution fails to achieve its potential impact. To this end, improved decision making is a

necessary piece of any IIoT solution. All of the data collection, communication, and analysis serves the purpose of enhancing productivity, safety, or efficiency.

Currently, changing needs and operating requirements in plants are manually identified, and actuation changes are initiated by plant personnel. In the IIoT, changing requirements are autonomously identified through data analysis. Controllers are automatically updated based on updated models from advanced data analysis [11]. In this respect, big data analytics, system identification, pattern recognition, and machine learning constitute methods for controller design. There are a wide variety of context-aware approaches in the literature, including unsupervised learning, supervised learning, rule-based logic, fuzzy logic, and probabilistic logic [33]. Each of these approaches have the potential to draw useful trends and improved actuation from large data sets.

One challenge in physical installation of actuators in IIoT applications regards the multiple layers of technology and equipment in the industrial process. New technology must be implemented into existing hardware and software of process systems [11]. The interdependence of equipment can present unanticipated consequences when actuating a single element in an entire system. In light of this challenge, focusing on energy support systems provides a feasible approach to initial adoption of IIoT. As mentioned previously in the chapter, energy support systems are readily found in most manufacturing plants. As such, implementation frameworks can be consistently applied across different manufacturing facilities. Additionally, corporate management typically do not perform a large overhaul of equipment specifically to increase plant energy efficiency. Instead, IIoT solutions for energy support systems can be implemented

incrementally in a facility with marginal additional cost. Finally, energy support systems often operate upstream and independent of the production process. For example, the interface of a pneumatic system with a manufacturing process consists of the available pressure and air flow capacity. As long as plant requirements are met for both pressure and flow, changes in the operation of the compressed air system will have no detrimental effect on the production process.

1.2.4.1 Industry Examples of Intelligent Actuation for Energy Applications

There are multiple available intelligent actuators in industrial applications to improve energy efficiency. One example is oxygen trim units used in boilers. Optimal boiler efficiency is achieved with proper air to fuel ratios. Boilers deviate from the optimal ratio due to changes in environmental and boiler conditions. Oxygen trim units automatically measure and adjust the air fuel ratio to achieve the optimal ratio. Another readily available actuator used in energy applications is automatic drain valves on compressed air systems. These valves sense the condensate build up in the compressed air line and purge the line when necessary. These drain valves eliminate compressed air loss compared to drain valves that use timers or other technologies.

1.3 Dissertation Outline

The remaining chapters of the dissertation are organized to provide examples of the three key elements of the IIoT framework. Specifically, technical analyses are presented that illustrate advanced data sensors, autonomous data analysis, and intelligent actuation. Each of these examples was developed to either improve the energy assessment process or illustrate a potential retrofit to improve the energy efficiency of

industrial buildings. Chapter 2 and Chapter 3 describe a non-intrusive flowmeter used in the diagnostics of compressed air leakage rate. Chapter 4 presents a prototype system capable of autonomously identifying and classifying facility lighting in order to make recommendations on energy efficient retrofits. Chapter 5 discusses a k-means pattern recognition algorithm for compressed air systems to intelligently adjust the compressor operating scheme for reduced energy use.

1.4 Summary

The Industrial Internet of Things is an emerging field in which sensors, middleware, and actuators are all networked together to create a smart environment. Potential improvements in the industrial field include enhanced manufacturing production, increased worker safety, and improved energy efficiency. This chapter discussed how energy efficiency fits in an IIoT framework by leveraging potential solutions that increase worker safety or improve preventative maintenance. This framework fundamentally consists of physical world sensors, middleware for communication and analysis, and actuators or user interfaces for improved performance. The remaining dissertation will present illustrative examples of each of the three foundational pillars of the framework.

CHAPTER 2

ADVANCED SENSORS

One aspect of the IIoT is the rapid increase in data storage and transmission capabilities. The ability to collect and store data is constrained by available sensors. A key component of advancing IIoT is the ability to sense important parameters and utilize that information for improved operation. While there are many readily-available sensors for use in monitoring equipment and processes, many desirable measurements are difficult to obtain. The desirable measurements may require prohibitively expensive equipment, be unreliable, or possibly can't be obtained with available technology. Advances in technology will drive down costs and make additional measurements possible. In many applications of the IoT or the IIoT, development of advanced sensors is a necessary component for widespread deployment of the technology to be economically viable. This is particularly true for IIoT for energy efficiency. As the primary benefit of energy efficiency is cost savings to the manufacturer, any deployment of technology must be low cost and result in a short payback period.

When sensors for desired measurements are prohibitively expensive or otherwise unavailable, at times the desired measurement can be derived from other measurements. Often termed virtual sensors, physical laws and relationships from measured parameters are used to estimate the desired parameter. This chapter describes a case study on how advanced sensors can derive desired measurements implicitly from other readily-available measurements. Developments in advanced sensors, in particular virtual

sensors, is an area of active research in energy systems and constitutes a key component in the realization of the IIoT for energy efficiency. The chapter presents work on a non-intrusive virtual flow meter. The need for an improved non-intrusive flow meter is discussed first, followed by a description of how the proposed system works. Simulations results are then presented, followed by results from experimental tests.

2.1 Case Study: Non-intrusive Compressed Air Flow Measurement using Thermal Signatures

In the Industrial Assessment Center, one of the most common energy recommendations is the repair of air leaks in compressed air systems. Accurate estimation of potential cost and energy savings to manufacturers requires a measurement of the air flow due to leaks in the system. This example case study presents work on a novel method to measure compressed air leakage rates in compressed air lines.

The method uses heat pulses and the resulting thermal response to non-intrusively derive the flow rate inside the compressed air line. The proposed system has the potential to simplify the process in determining leakage airflow rates in compressed air systems while maintaining similar or greater accuracy than current methods in the Industrial Assessment Center. Compared to these current methods, this system can simplify the calculations and reduce the possibility for error when taking multiple measurements. The methodology of the proposed method is explained, followed by simulations that validate the method. Simulation are performed to evaluate both a step response and frequency approach to maximize potential measurement sensitivity. An experimental system and results are then presented to confirm simulation results.

2.1.1 Background and Introduction

Compressed air systems are one of the most common systems found in industrial facilities. In fact, the widespread use of compressed air in industry has led many people to term compressed air a fourth utility. Some estimates report that compressed air constitutes approximately 10% of all industrial electricity [40]. Due to inefficiencies in compression, the relative cost of compressed air is high [41]. Because of this high cost, any losses in the system result in substantial wasted electricity. There are several points of loss and inefficiencies in a compressed air system, including oversized compressors, high operating pressure, and inefficient compressed air drying [40, 42]. These inefficiencies in energy use result in waste, where losses in the system can range up to 60% [43].

One of the most common and worst contributors of losses is leaks in compressed air lines. Leaks in an unmaintained plant typically consume 10-40% of the total system usage [44]. The large amount of lost energy presents an opportunity for cost savings by repairing leaks. Additionally, plants often run the compressed air system at higher pressures to compensate for leaks in order to maintain downstream pressure at equipment. The higher pressure exacerbates compressed air leakage, inducing additional losses. Reducing these leaks has demonstrated significant cost savings for industrial plants [19].

Figure 2 displays data from the Texas A&M Industrial Assessment Center (IAC) on recommendations related to repairing compressed air leaks. The graph displays data taken from approximately 50 compressed air systems and shows the average leakage rate

for the facilities. The annual cost of the compressed air leaks depends on the cost of electricity, compressor efficiency, and operating hours of the plant. For a 2-shift plant operating 5 days a week with a cost of electricity and demand of $0.06/kwh and $4/kW·mo, respectively, a leakage rate of 100 SCFM would result in an annual cost of $6,800. If the compressor continues to run when the plant is shutdown, the resulting cost of compressed air leaks is over $12,500. Figure 2 demonstrates that most systems have an opportunity to reduce utility expenditures in compressed air systems by reducing the leakage rate of air. Repairing leaking systems requires minimal capital investment and expertise, resulting in short payback periods.

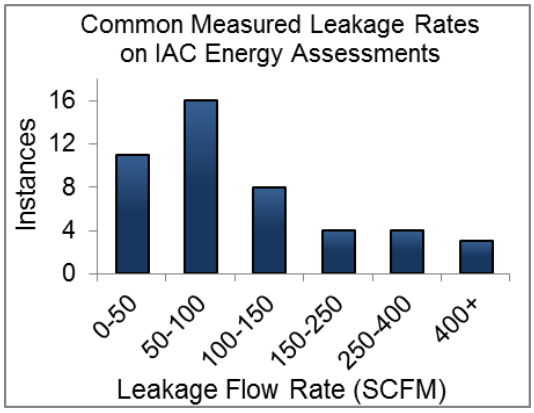

Figure 2. Measured leakage rates from approximately 50 compressed air systems in small- to medium-sized manufacturers. Data were taken from the Texas A&M University Industrial Assessment Center.

2.1.1.1 Current IAC Compressed Air Flow Estimation Methods

The primary measurement in estimating potential energy and cost savings is the quantification of the air flow rate due to leaks. Once the flow rate is determined, conversion factors are used to estimate the required electricity for compression. These

other factors and values are obtained from nameplates and utility information. The current methods for quantifying leakage rate either estimate total system rate or estimate flow from individual leaks. In estimating total system losses, the four primary methods are performing a leak-down test, measuring compressor power, using an ultrasonic flow meter, and installing an inline flow meter. All of these methods require production to be halted, so that all air flow in the compressed air lines comes from leaks in the system. Measuring individual leaks can be done while production is still running.

Each of the listed methods has disadvantages that create high uncertainty in the air flow measurement. The following paragraphs describe each method along with the relative advantages and disadvantages of the method. Table 1 summarizes these advantages and disadvantages.

2.1.1.1.1 Leak Dropdown Test

A leak-down test records the time for pressure to drop off the system when the compressor is shut off. By estimating the volume of the system, a correlation can be made to estimate total flow [42]. Determining the volume of all major compressed air lines and tanks can be difficult, resulting in large uncertainties. Careful verification needs to be performed with this method because the estimated leakage rate at times exceeds the capacity of the air compressor.

2.1.1.1.2 Ultrasonic Flow Meter

Commercially available clamp-on ultrasonic flow meters use ultrasonic waves to estimate the speed of the air flowing inside of the compressed air line. This method requires extensive setup and calibration along with knowledge of several parameters that

are difficult to obtain. Additionally, strict requirements are given on the location of the meter and correct attachment to the pipe. Small variations in this setup create large uncertainties in flow measurement.

2.1.1.1.3 Electrical Current Measurement

Power measurement of the compressor directly measures the electricity required for compression due to leaks. This method is common in professional energy assessments, but the unique time limitations on the one-day assessments and limited training of students in the IAC program restricts the use of this method. Additionally, this method can only estimate total system flow; leakage rate of individual lines cannot be obtained from the power meter.

2.1.1.1.4 Inline Airflow Measurement

Inline flow meters are the most accurate method to measure flow rate in compressed air lines, but they require drilling into the line for installation of the meter. Limited time for data collection on the one-day assessments and substantial resistance from the plant for permanent alteration to plant equipment limit the ability to use inline meters during IAC assessments.

2.1.1.1.5 Ultrasonic Detection of Individual Leaks

In identifying individual leaks, the most common method for finding leaks is with ultrasonic leak detection [40]. Compressed air leaks are loud and can be identified audibly when other equipment is stopped. When production is still running, ultrasonic leak detectors can aid in identifying leaks among other loud equipment. Although some work has been performed to estimate sizes of individual leaks based on ultrasonic

measurements, typically the size of each leak is manually estimated [45]. The uncertainty in leak size, together with the challenge in determining all leak points, makes using this method to determine total system losses difficult. Figure 3 shows an example of an ultrasonic leak detector.

Figure 3. Example of an ultrasonic leak detector used to find individual compressed air leaks.

Table 1: Summary of Current Leak Quantification Methods

Method	Advantages	Disadvantages
Ultrasonic Flow Meter	Commercially available	Highly sensitive to small perturbations Time intensive setup and calibration Unknown parameters must be guessed
Leak Dropdown Test	Simple physics correlation Limited required expertise	Large uncertainties in measurement
Electrical Current Measurement	Direct electricity measurement	Expertise required to install meter Only total system use is estimated
Inline Airflow Meter	Accurate measurement	System must be permanently altered
Individual Ultrasonic Detection	Can be done during production Individual leaks are identified	Leaks may be difficult to reach Time consuming

2.1.1.2 Operating Limitations in the IAC Program

The IAC program has two primary purposes in performing energy assessments. 1. Help manufacturers improve their competitiveness by identifying waste and cost-effective measures to reduce utility use and improve productivity, and 2. Train undergraduate and graduate students in energy management and efficiency. While the visits have significant documented savings for participating manufacturers, the purpose of training students places unique operating requirements on data collection and analysis for recommendations. First, the expertise of the auditors can be limited as some students join the IAC program during their undergraduate sophomore year. Second, a typical one-day assessment consists of a plant tour, interviews with facility personnel, identification of recommendation projects, and data collection. With these considerations, the operating requirements for quantifying air flow include minimal operator training, simplicity of setup & operation, nonintrusive data collection, and minimal time for data collection. These unique conditions for energy assessments dictate the restrictive operating requirements in obtaining compressed air flow of leaks.

2.1.1.3 Other Industries and Applications for a Non-intrusive Flow Meter

The unique requirements in the IAC program are atypical of traditional energy assessments, and the proposed flow meter may have limited market applicability in the energy assessment field. However, there are many industries in which the need for an improved non-intrusive flow meter has been identified. These industries generally fall into two categories. The first is one in which a one-time spot measurement is desired, and installing an inline flow meter is impractical or undesirable. The second main

category is in applications where the fluid inside the pipe is corrosive. Inline sensors quickly fail in these applications, and there is a strong maintenance and economic incentive for a non-intrusive meter. The identified applications for an improved non-intrusive flow meter are listed in Table 2.

Table 2. Other Applications for Non-intrusive Flow Meters

Identified Application	Applicable References
Geothermal industry for brine flows	[46]
Hot and chilled water on process and HVAC systems	[47-49]
Refrigerant flow on HVAC systems for fault detection and diagnosis	[47]
Diagnostic tool for pump performance in nuclear applications	[46]
Production water in oil and gas wells	[50]
Other oil and gas applications, including injected water, captured water and multiphase flows	[50-52]
Steam traps and steam systems	[53]

2.1.2 Proposed Design

The proposed non-intrusive flow meter uses the thermal response of a compressed air line to derived the flow inside of the line. The setup of the system is diagrammed in Figure 4. The device is wrapped around a compressed air line with the heaters and thermocouples. Insulation is also wrapped around the thermocouple and heating elements to limit free convection off the surface of the pipe. Similar to other current methods, production use is halted, so that only the airflow due to leaks is measured. Once production use is stopped, the compressor remains running to maintain pressure in the system. The exterior of the pipe is heated using thermal heaters, and the

resulting temperature profile is measured using thermocouples or other similar temperature measurements. When heat is applied to the pipe, the resulting temperature profile will change depending on air flow in the pipe. Increased airflow with increase the convection coefficient on the inside of the pipe, and the resulting temperature will decrease relative to the reference temperature profile. There are different methods on how to relate the differing temperature profiles to the flow rate. There are two methods approached in this research: a steady state method and a dynamic method. Once the estimated leakage rate is obtained, the compressor specific efficiency is used to estimate the energy and monetary loss due to air leaks.

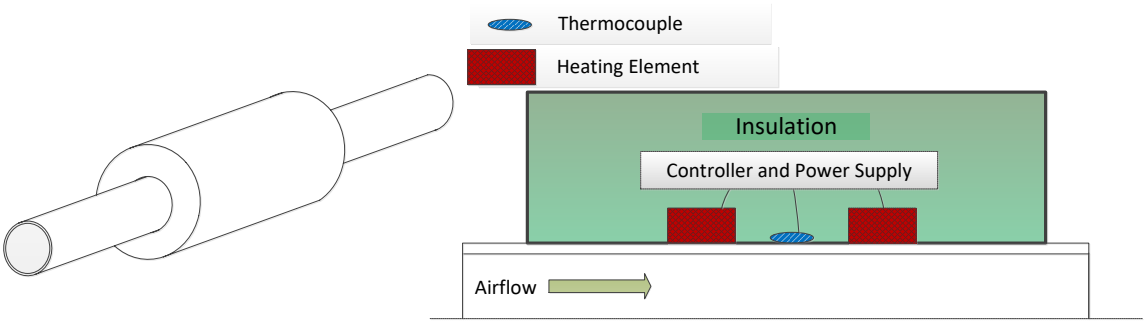

Figure 4. Proposed device to measured airflow based on the thermal response. A controller is used to control the heaters to a specified temperature, and the thermocouple is used to measure the resulting temperature profile. Insulation reduces free convection off the outer surface of the pipe.

2.1.2.1 Steady State Method

An initial approach to flow estimation is to use the step response of the compressed air line to determine the flow. A pipe is heated to a reference temperature, after which the power to the heater is stopped and the response observed. The pipe will

cool back to the initial ambient temperature, and the rate of that cooling will depend on the airflow. This is illustrated in Figure 5. The Biot number of the pipe is sufficiently small to use lumped capacitance in the radial direction, thereby enabling the use of the outer surface temperature as an estimate of the inner surface temperature. A faster cooling rate expedites the temperature decay rate. A correlation between the estimated cool rate or time required to cool a specified temperature can be made to the flow rate. This method is basic in approach and relies predominantly on Newton's law of cooling.

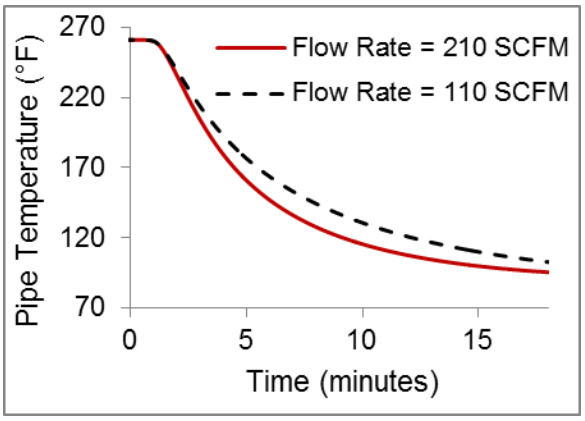

Figure 5. Proposed methodology to differentiate flow rates based on cooling of a compressed air pipe. The temperature drops faster with higher flow due to the higher convection coefficient on the inner pipe surface.

2.1.2.2 Dynamic Method

The dynamic method uses periodic input signals to derive the flow rate inside the pipe. The dynamic method can potentially extract more accurate flow information or shorten the time of data collection relative to the steady state method. After transients have settled out, the output profile will be periodic with the same input frequency but

with a phase shift and a different amplitude. A sketch of this is shown in Figure 6. Since the output gain and phase shift will be affected by the flow rate inside the pipe, an experimental correlation can be developed to estimate the flow rate.

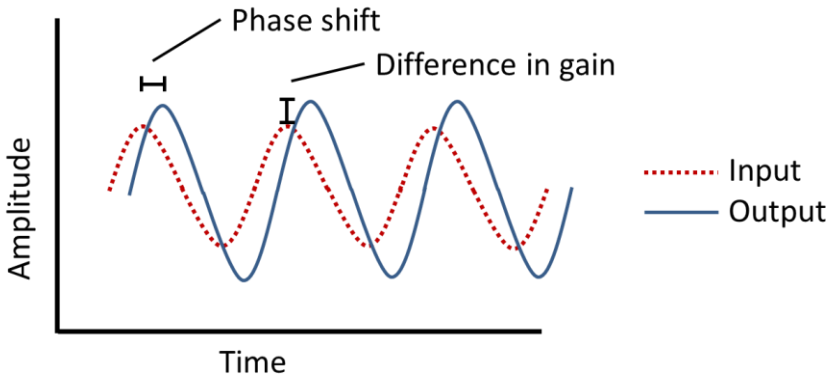

Figure 6. The dynamic method uses the phase shift and difference in gain between the input and output signals to estimate flow rate inside the compressed air line.

2.1.2.3 Advantages of Design

There are several advantages of the proposed design over other methods for non-intrusive measurement of flow rate. Each of these advantages is described in the following sections along with its significance.

2.1.2.3.1 Minimal setup and required expertise

The proposed method simplifies data collection to a single point of measurement with minimal setup and calibration. The device does not require extensive calibration or specialized expertise, making it suitable to undergraduate student technicians in the IAC program.

2.1.2.3.2 Simple dynamics

Fundamentally, the dynamics of the system are governed by Newton's law of cooling and Fourier's law of heat conduction. Each of these equations are straightforward and simple. Resulting correlations are derived from these dynamics, which overall simplifies the data analysis. As a single measurement takes several minutes to complete, the method will also inherently average out small variations in flow or noise in the data.

2.1.2.3.3 Low cost

The required materials to build and operate the system consist of thermal heaters, thermocouples or similar temperature measuring devices, a microcontroller, and insulation. Each of these items is readily-available and low cost, which results in an overall low cost to build and operate the system.

2.1.2.3.4 Independent of temperature

The resulting temperature profile is measured relative to the ambient temperature so that the initial conditions of the system are zero. Since all measurements are measured relative to the ambient temperature and there is a linear relationship between the temperature difference and heat transfer, the system can operate at different ambient temperatures with negligible change in the response.

2.1.2.3.5 Independent of pressure

The measured flow rate inside the pipe using the thermal response is independent of the pressure inside the pipe. This point is non-intuitive, especially considering how methods such as the ultrasonic flow meter directly measure the speed in the pipe and rely

on knowledge of system pressure. The governing equation for the thermal method is Newton's law of cooling (Equation (1)). The pipe will cool down from conduction within the pipe and convection on the inner surface of the pipe. The heat rate due to conduction depends on the convection coefficient and the temperature difference between the pipe and the flowing air:

(1) $$\dot{Q} = hA\left(T_{pipe} - T_{air}\right)$$

The convection coefficient in the heat transfer equation is a function of Reynolds number, the Prandtl number, the thermal conductivity of air, and pipe dimensions:

(2) $$h = f(Re, Pr, k_{air}, D)$$

The Reynolds number depends on the air density, speed, dynamic viscosity and the pipe diameter.

(3) $$Re = \frac{\rho_a v_a D}{\mu}$$

When measuring the flow rate, the flow rate at standard conditions is the desired measurement. The flow rate at standard conditions is related to the flow rate at other pressures using the ratio of densities, as seen in Equation (4).

$$\dot{V}_s = \dot{V}_a \left(\frac{\rho_a}{\rho_s}\right) = v_a \left(\frac{\pi D^2}{4}\right)\left(\frac{\rho_a}{\rho_s}\right) \tag{4}$$

Rewriting the equation and substituting into Equation (3), the Reynolds number can be rewritten and simplified as:

$$Re = \frac{4\rho_s \dot{V}_s}{\mu \pi D} \tag{5}$$

Using this form of the Reynolds number equation, the convection coefficient is a function of the density at standard conditions, the dynamic viscosity, diameter, Prandtl number, thermal conductivity of air, and the volumetric flow rate. Using tabulated information for air, the differences in dynamic viscosity, Prandtl number, and thermal conductivity of air with respect to pressure are small. Using the values at 0 psig and 100 psi, the convection coefficient is estimated using the Dittus-Boelter equation for both a 2" and 3" line at various flow rates. Figure 7 shows the estimated flow rate under these conditions. As seen in the figure, the dependence of the heat transfer coefficient is small over the wide range of operating pressures. Most plant pressure setpoints will fall in a

much smaller range, highlighting the independence of flow rate measurement due to system pressure.

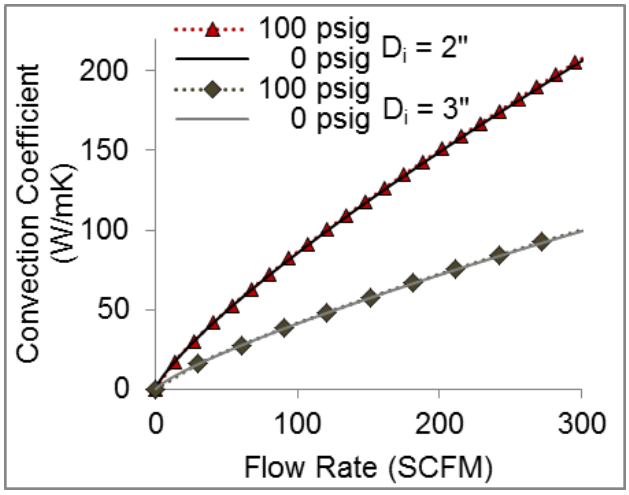

Figure 7. The heat transfer on the inner surface of the compressed air line is independent of the pressure inside the compressed air line

2.1.2.4 Method Limitations

There are several challenges to the proposed system that are addressed in this section. These challenges will cause the estimation of flow to deviate from the true value, degrading the accuracy and repeatability of measurements. Each of the challenges is discussed individually along with the approach to minimize or remove the effect.

2.1.2.4.1 Free Convection on the Outer Pipe Surface.

The inner pipe surface temperature is estimated by measuring the outer surface temperature. The outer surface of the pipe will be subject to free convection based on the ambient air temperature, resulting in increased cooling. Additionally, as the thermocouple is located on the outer surface, free convection from the thermocouple will

result in sensor error. To reduce the effect of free convection in both these instances, an insulation blanket is wrapped around the device and surface of the pipe.

2.1.2.4.2 Axial Conduction along the Pipe.

Lumped capacitance can be used to estimate the inner pipe surface temperature from the outer pipe surface temperature because of the high thermal conductivity of the pipe material. However, the conductivity in the pipe results in large heat transfer along the axial direction of the pipe. When the heaters are turned off, axial thermal conduction will increase the rate of cooling, resulting in overestimation of flow rates. In the proposed method, multiple approaches are used to mitigate the effects of axial conduction in the pipe. The first way is by minimizing conduction through control of the heaters. Once the pipe is initially heated during a test, the pipe is maintained at the elevated temperature to minimize the effects of axial conduction when cooling. Conduction heat transfer is determined by the temperature gradient in the material. By maintaining the pipe at a constant temperature prior to shutting off the heaters, the temperature gradient at the thermocouple is reduced, thereby reducing the effects of axial conduction.

The second approach to reduce the effect of axial conduction is through the testing procedure. During heating, the temperature is held for the same length of time on each test prior to turning the heaters off. Any remaining axial conduction will be incorporated into the experimental correlations because each test will be affected similarly by axial conduction.

2.1.2.4.3 Variability in Pipe Materials.

Although the dimensions and material properties can nominally be determined from manufacturing specifications, variations in these properties are inevitable. Differences in conductivity, specific heat, density, and pipe thickness result in discrepancies between the actual and estimated flow rate. No approach is taken in the proposed method in this chapter to address this challenge; the work in Chapter 3 will account for variation in pipe materials.

2.1.3 Evaluation of Design in Simulation

The proposed method was tested in two separate simulation programs. The first program used was Finite Element Heat Transfer, a basic finite element program. To test the dynamic method of flow estimation, a separate numeric model was developed in Matlab. Both of these models are built as a 2-dimensional axisymmetric model. The setup of each of these simulations are described in the following sections.

2.1.3.1 Setup of Simulation

2.1.3.1.1 Finite Element Modeling

A finite-element simulation software, Finite Element Heat Transfer (FEHT), was used to verify the initial concept and develop initial correlations between heat transfer coefficients and flow rates. The convection coefficient is directly correlated with the speed of the air and can be used to derive air flow rates. In simulating the system, the dimensions and material properties were used from an experimental compressed air line. The outer surface was specified as adiabatic. The heat transfer from the heaters to the pipe was simulated so that the pipe would have no axial conduction, thereby

guaranteeing all of the heat transfer occurred at the inner surface of the pipe. Gnielinski's correlation was used to relate the convection term to an estimated flow rate [54]. Convection coefficients were chosen based on common air flow leakage rates found in IAC assessments.

2.1.3.1.2 Dynamic State Space Representation

A state space representation was used to simulate the thermal response of the pipe to a periodic or transient heat input. The numeric model used to analyze frequency responses was created using the model setup displayed in Figure 8. This two-dimensional model allows for both axial and radial conduction with heat transfer on the inner surface of the pipe. Multiple inputs and outputs are allowed in the model, but only a single input and output was used in the simulation in testing the dynamic response. The input to the model is the signal given to the heater, while the output is the temperature on the outer surface on the pipe at the desired location.

The individual state equations can be quite involved when considering the heat transfer effects from each neighboring element. As such, the following section presents the heat transfer equation for a state on the inside surface of the pipe. Equations for internal states, states on the exterior of the pipe, and states for the corners of the pipes are all included in the Appendix A.

The model also captures the dynamics of air heating up in the pipe and the resulting asymmetric temperature profile around the heater. The derivation for the equations that include air differ significantly from the equations involving pipe material. As such, the derivation of these equations are presented in the following section.

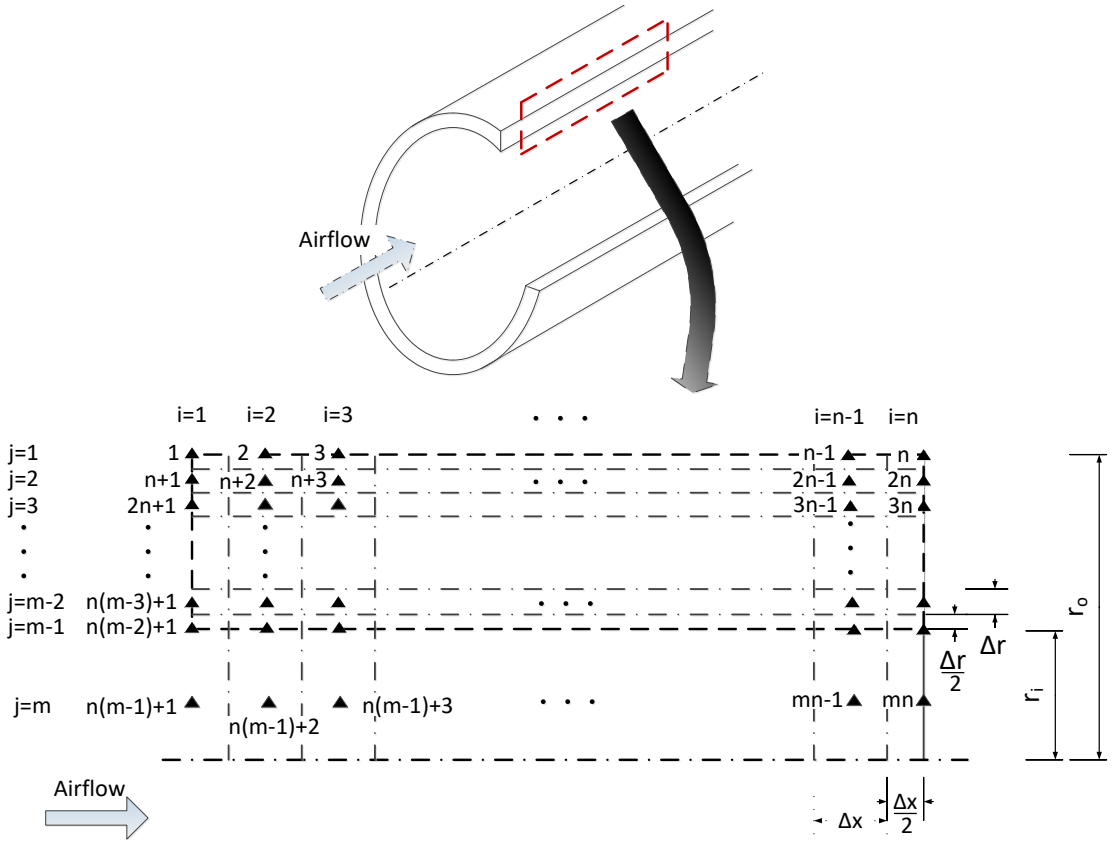

Figure 8. Outline of the model used for analyzing the dynamic response of the pipe.

2.1.3.1.2.1 Derivation of Interior Nodes Control Volume Equation

The individual state equations are derived based on first principles model of the energy balance. In general terms, the energy balance states that the change in energy of a system is the sum of the energy interactions of the system with the surroundings as shown in Equation (6).

(6) $$\frac{dE}{dt} = \dot{Q}_{x-1,r} + \dot{Q}_{x+1,r} + \dot{Q}_{x,r+1} + \dot{Q}_{x,r-1}$$

For the interior states, the energy interactions include conduction from neighboring pipe nodes and convection from the neighboring air node. Using the indexing format displayed in Figure 8, Equation (7) displays the energy balance of the interior states.

(7)
$$\rho c\left(\pi r_I \Delta r \Delta x\left(1+\frac{\Delta r}{4r_I}\right)\right)\frac{\partial T_{(m-2)n+i}}{\partial t} = \\ = k\pi r_I \Delta r\left(1+\frac{1}{4r_I}\right)\frac{\left(T_{(m-2)n+i-1}-T_{(m-2)n+i}\right)}{\Delta x} \\ +k\pi r_I \Delta r\left(1+\frac{1}{4r_I}\right)\frac{\left(T_{(m-2)n+i+1}-T_{(m-2)n+i}\right)}{\Delta x} \\ +\left(r_I+\frac{\Delta r}{2}\right)\Delta x\,\frac{\left(T_{(m-3)n+i}-T_{(m-2)n+i}\right)}{\Delta r} \\ +h2\pi r_I \Delta x\left(T_{(m-1)n+i}-T_{(m-2)n+i}\right)$$

Grouping like terms and simplifying the equation, the equation can be written in a form compatible for a state space representation. Equation (8) shows the rewritten equation that is used in the generation of the state space representation.

(8)
$$\frac{\partial T_{(m-2)n+i}}{\partial t} = -\left[\frac{2\alpha}{(\Delta x)^2}+\frac{2\alpha}{(\Delta r)^2}\frac{\left(r_I+\frac{\Delta r}{2}\right)}{\left(r_I+\frac{\Delta r}{4}\right)}+\frac{2h}{\left(1+\frac{\Delta r}{4r_I}\right)\rho c\Delta r}\right]T_{(m-2)n+i} \\ +\frac{\alpha}{(\Delta x)^2}\left(T_{(m-2)n+i+1}+T_{(m-2)n+i-1}\right) \\ +\frac{2\alpha}{(\Delta r)^2}\frac{\left(r_I+\frac{\Delta r}{2}\right)}{\left(r_I+\frac{\Delta r}{4}\right)}T_{(m-3)n+i}+\frac{2h}{\left(1+\frac{\Delta r}{4r_I}\right)\rho c\Delta r}T_{(m-1)n+i}$$

2.1.3.1.2.2 Incorporation of air heating up in the model

In adding the effects of air heating up inside the test section of the pipe, the following assumptions were made:

- Between time steps, the temperature of the air at each node remains constant.
- The effects of conduction of the air between nodes upstream and downstream are negligible.
- At each node, air is well mixed.

For each time step, the temperature of control volume k at time $t + \Delta t$ is equal to the temperature of the control volume that is n_s control volumes earlier plus the temperature increase as the volume of air passes through the pipe. Visually, this is displayed in Figure 9. Between time steps, the control volume travels n_s time steps and heats up during that section. For the final ending node, the volume of air may start between two nodes, and this is accounted for in the derivation of the equations.

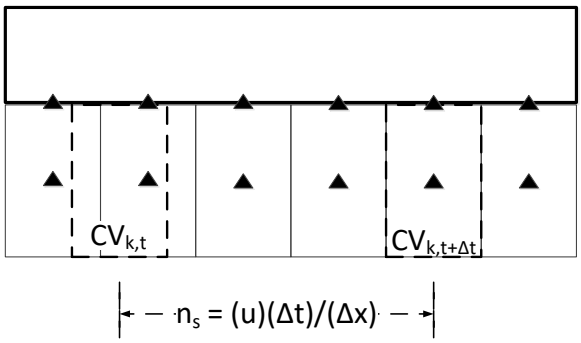

Figure 9. Illustration of how the control volumes are selected for the air nodes in the pipe model.

Note that although the previous equations are derived in continuous state space, the actual implementation and simulation occurs in discrete time, so the method here can be applied. For the control volume (CV) k of air, the number of nodes the CV travels in one time step is equal to:

(9) $$\mathrm{n_s} = \frac{\mathrm{u}\,\Delta\mathrm{t}}{\Delta\mathrm{x}}$$

The amount of time that the CV k spends at each node is equal to:

(10) $$\mathrm{t_s} = \frac{\mathrm{n_s}}{\Delta\mathrm{t}} = \frac{\mathrm{u}}{\Delta\mathrm{x}}$$

Let f_s be the fractional part of $\mathrm{n_s}$, $\mathrm{fl}(\mathrm{n_s})$ be the integer part of $\mathrm{n_s}$, and $\mathrm{cl}(\mathrm{n_s})$ be the ceiling of n_s. Also, let $c = \mathrm{t_s}\frac{\mathrm{hA_i}}{\mathrm{p_a c_{p,a} V}}$. Then the temperature of the CV k which is located at node $n(m-1)+i$ at the time $t+\Delta t$ is given by:

(11) $$\begin{aligned}\mathrm{T_{n(m-1)+i,t+\Delta t}} = {} & (1-\mathrm{c})^{\mathrm{fl(n_s)}}\left(1-\mathrm{c\,fr(n_s)}\right)\mathrm{T_{n(m-1)+i-n_s,t}} \\ & +(1-\mathrm{c})^{\mathrm{fl(n_s)}}\mathrm{fr(n_s)c\,T_{n(m-2)+i-cl(n_s),t}} \\ & +\sum_{\mathrm{q=1}}^{\mathrm{fl(n_s)}}(1-\mathrm{c})^{\mathrm{fl(n_s)-q}}\mathrm{c\,T_{n(m-2)+i-fl(n_s)+q-1,t}}\end{aligned}$$

Note that the summation is only used when $fl(n_s) \geq 1$.

In this equation, the temperature at the starting location of the control volume is given by the following.

(12) $$T_{n(m-1)+i-n_s,t} = (f_s)\left(T_{n(m-1)+i-cl(n_s),t}\right) + (1-f_s)\left(T_{n(m-1)+i-fl(n_s),t}\right)$$

Since the state space form is set up in terms of the temperature derivative at the specific location, the temperature relation is rearranged.

(13) $$\begin{aligned}\frac{T_{n(m-1)+I,t+\Delta t}-T_{n(m-1)+I,t}}{\Delta t} &= -\frac{1}{\Delta t}T_{n(m-1)+I,t} \\ &+\frac{1}{\Delta t}(1-c)^{fl(n_s)}\left(1-c\,fr(n_s)\right)(f_s)\left(T_{n(m-1)+i-cl(n_s),t}\right) \\ &+\frac{1}{\Delta t}(1-c)^{fl(n_s)}\left(1-c\,fr(n_s)\right)(1-f_s)\left(T_{n(m-1)+i-fl(n_s),t}\right) \\ &+\frac{1}{\Delta t}(1-c)^{fl(n_s)}fr(n_s)c\,T_{n(m-2)+i-cl(n_s),t} \\ &+\frac{1}{\Delta t}\sum_{q=1}^{fl(n_s)}(1-c)^{fl(n_s)-q}c\,T_{n(m-2)+i-fl(n_s)+q-1,t}\end{aligned}$$

2.1.3.1.3 Model Reduction

In order to accurately capture the dynamics of the temperature response, sufficient states are necessary in both the axial and radial directions. As additional states are added, computational time increases. Model reduction can be done to reduce model complexity while maintaining the model integrity. The numeric model, as derived in the previous section, can be written in the state space form:

(14) $\dot{z} = Fz + Gu \qquad y = Hx$

For model order reduction, a similarity transformation of the state space system is used to find a balanced realization of the system [55].

(15) $\dot{x} = Ax + Bu \qquad y = Cx$

where,

(16) $A = TFT^{-1} \qquad B = TG \qquad C = HT^{-1}$

The Hankel singular values are used to find an appropriate order of the system. Figure 10 shows the Hankel singular values for a 357-state system. As seen in the figure, the Hankel singular values quickly drop off. After approximately 4 states, the values are near zero, indicating that the additional states may be removed with minimal effect on the accuracy of the estimated model.

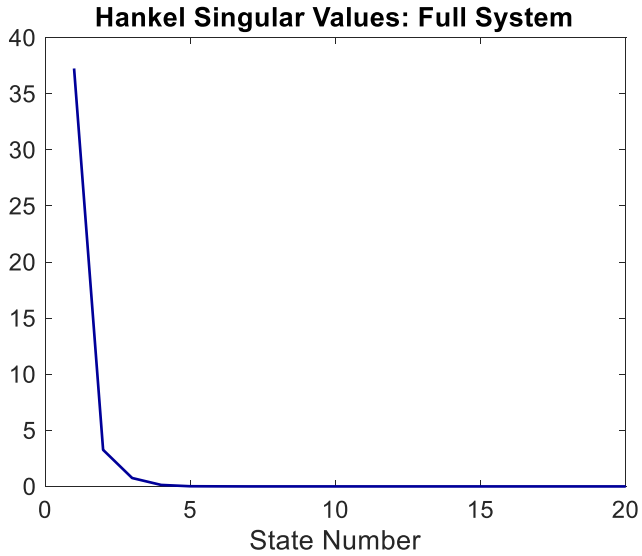

Figure 10. Hankel singular values for the numeric 357-state system. Only the first 20 states are displayed on the figure.

These states are eliminated from the model. The reduced order model is shown in Equation (17), where x_1 indicates the new states in the system.

(17)
$$\dot{x}_1 = A_{11}x_1 + B_1 u \qquad y = C_1 x_1$$

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \qquad B = \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} \qquad C = [C_1 \quad C_2]$$

In this form, x_1 is of dimension $[k\ x\ 1]$, where k is the reduced order of the model and all other matrices are of appropriate dimensions. Figure 11 shows the response for the full state model along with reduced model with 4 total states for a constant heat input. The models compare favorably to each other with minimal error. Figure 12 compares the models for an impulse response, where faster dynamics are

observed. Once again, the models compare favorably with minimal error. The reduced order model is used in simulation to reduce computation time.

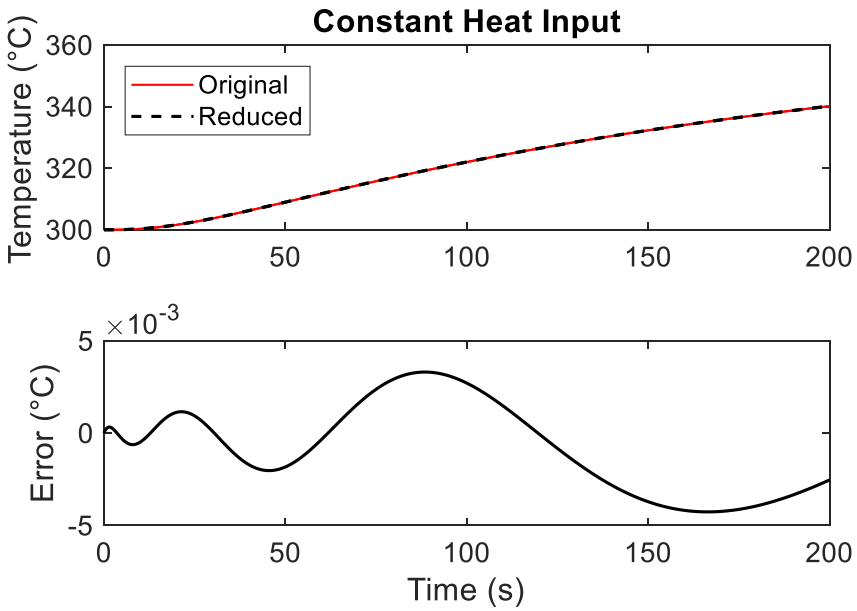

Figure 11. Comparison of the full- and reduced-order models with a constant heat input source.

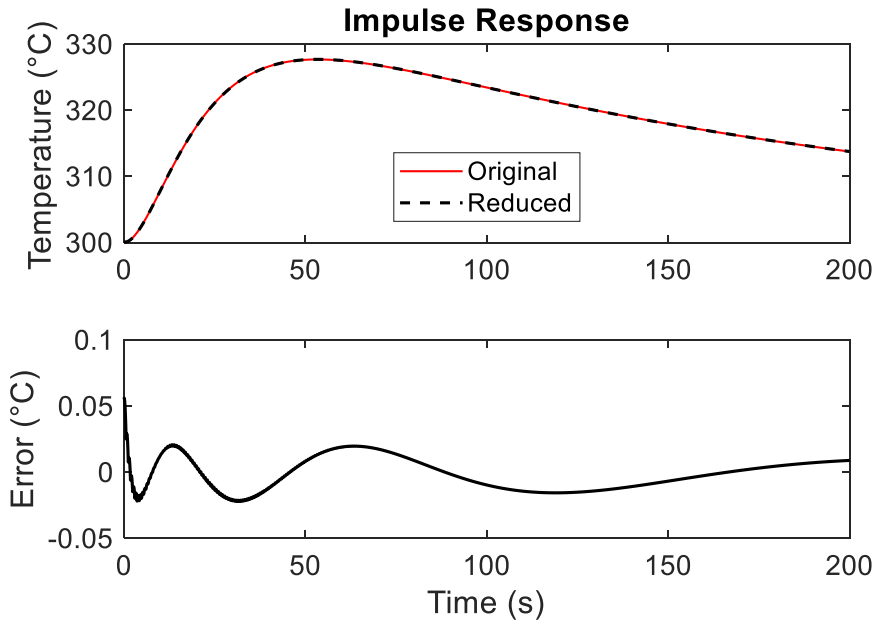

Figure 12. Comparison of the full- and reduced-order models for an impulse.

2.1.3.1.3 Comparison of FEHT and Finite Difference Models

Before experimental verification of the models, the two different models were simulated under the same conditions to compare the results of the two models. The results of a constant heat input are shown in Figure 13. This figure demonstrates that both models demonstrate a similar response with negligible difference in error. The comparison verifies the numeric model in Matlab against FEHT, a commercially-available software program for heat transfer simulation. The numeric model in Matlab has the additional capability of quickly simulating dynamic input signals, thereby enabling the dynamic analysis as described previously.

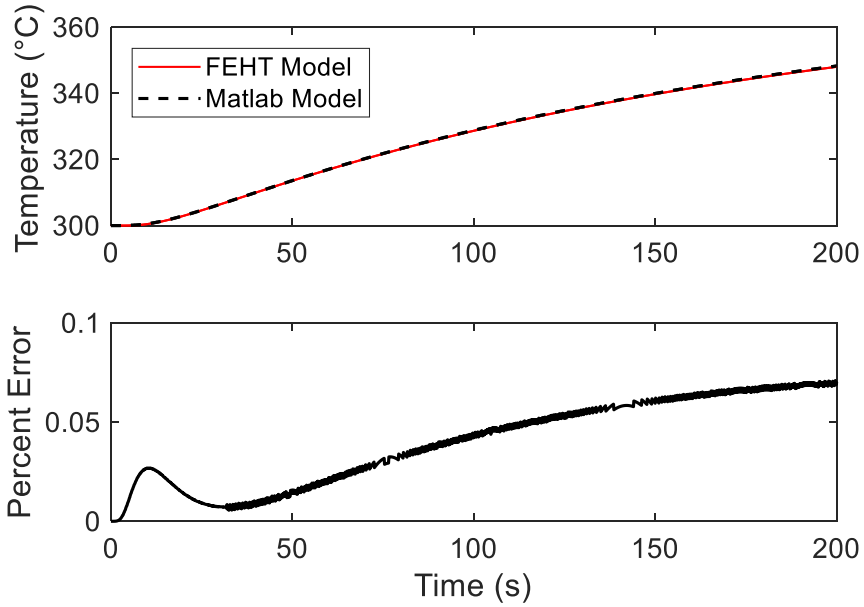

Figure 13. Comparison of the Finite Element Heat Transfer and Matlab models using the same parameters.

2.1.3.2 Simulation Results for Steady State

The ability to measure flow based on a steady state temperature response was tested in simulation using both the FEHT and Matlab numeric models. The two models gave similar responses, so only the Matlab model responses are displayed in the dissertation. In these simulations, the temperatures of all pipe nodes were set to an elevated temperature above ambient. In the experimental design, the heaters are set to maintain the pipe temperature for a specified period of time. The simulations capture the limiting case of maintaining the heater until the pipe is at uniform temperature. At the start of the simulation, the inner pipe wall is exposed to the air at ambient temperature, resulting in a step response of the pipe. The simulation results are displayed in Figure 14. The first graph displays the resulting temperature profiles for the different flow rates. As

expected, the instances with larger flow have a higher cooling rate. After the initial few seconds, the cooling rate of each test is approximately exponential.

Different metrics from the temperature profile can be used in deriving correlations between flow rate and the temperature response. As Newton's Law of Cooling predicts that the rate of temperature change is proportional to the temperature difference, the overall cooling rate is defined by exponential decay. The rate of cooling can be estimated by measuring the elapsed time at different percentages of the time constant. The second graph in Figure 14 shows the elapsed time for the temperature to reach 40% and 70% of one time constant. As can be seen in the graph, a clear trend indicates the relationship between the elapsed time and the flow rate. Additionally, the rate of exponential decay can be estimated by fitting an exponential curve to the measured temperature profile. The third graph shows the correlation of flow rate based on the estimated rate of exponential decay.

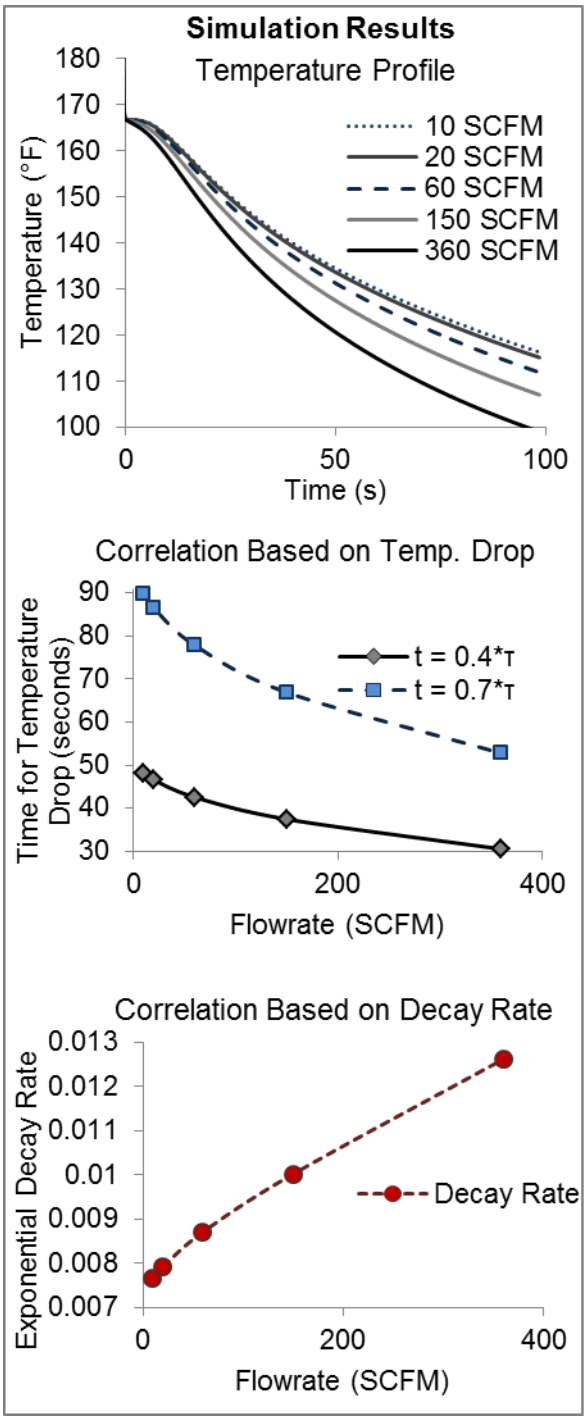

Figure 14. Temperature profiles and resulting flow rate correlations based on a finite-element simulation. Two different types of correlations were used, one based on measuring the time for the temperature to drop and the other based on the exponential decay rate.

2.1.3.3 Simulation Results for Dynamic Analysis

To determine the feasibility of using dynamic analysis in deriving compressed air flow, the Bode diagrams are contrasted for different flow rates of air. Significant differences in phase angle or magnitude across different flow rates can enable flow measurement through these frequency characteristics. To compare the Bode diagrams, the phase angle and magnitude were evaluated on the Bode plot for varying flow rates inside the pipe. Figure 15 shows the differences in phase angle and Bode magnitude for different frequencies as a function of flow rate.

This analysis reveals that using frequency response offers no practical benefit beyond a steady state approach. There are appreciable differences in phase angle for a certain range of frequencies. However, these frequencies are unreasonably long to achieve during data collection, especially because multiple time periods are often required for the transient response to die out. For very low frequencies (large time periods), including a step input, there is minimal difference in phase angle but a considerable difference in the magnitude of the output.

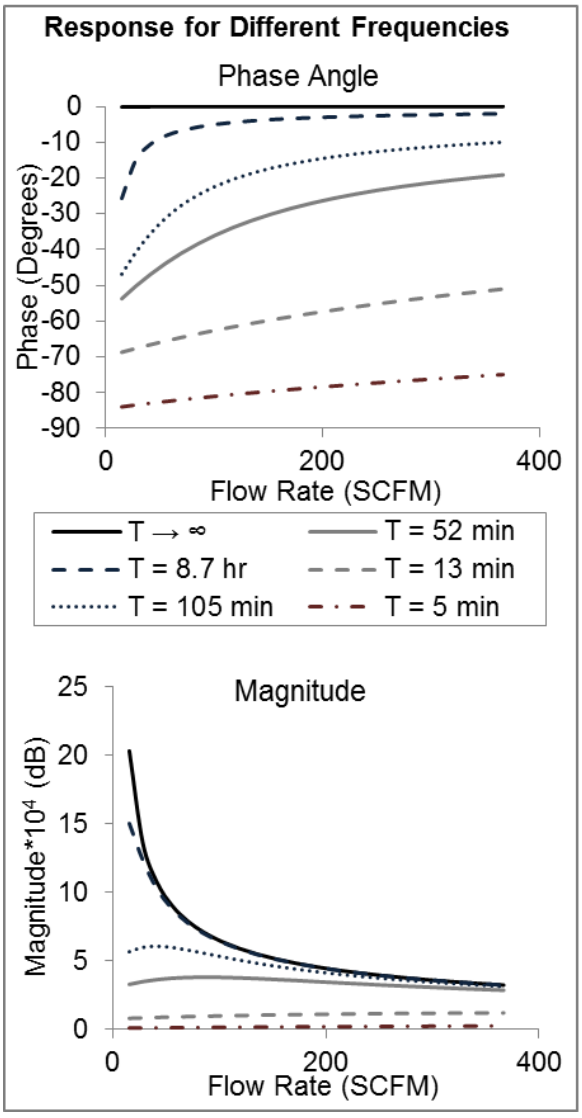

Figure 15. Use of frequency analysis to derive compressed air flow offers no practical advantage over a steady-state or step response.

2.1.4 Experimental Verification

2.1.4.1 Experimental Setup

An experimental setup was created to validate the simulation results and evaluate the effect of noise on the estimation of flow. Figure 16 shows the experimental setup. Two heater bands were wrapped around the pipe with the test thermocouple between the

two heaters. Additional thermocouples were attached on either side of the heaters for additional data collection. These additional thermocouples were used to visualize the general temperature profile of the pipe but not to estimate flow. An insulation blanket, seen in the figure behind the thermocouples and heaters, was wrapped around the pipe surface. The test thermocouple and the heaters were connected to a control box. The control box was equipped with relays, a PID controller, and a fuse. All of the thermocouples were attached to a data acquisition module for data collection. As the correlation between flow rate and heat transfer is independent of pressure, a fan with a variable transformer was employed to generate flow through the pipe. A handheld anemometer was used at the end of the open pipe to measure actual flow through the pipe.

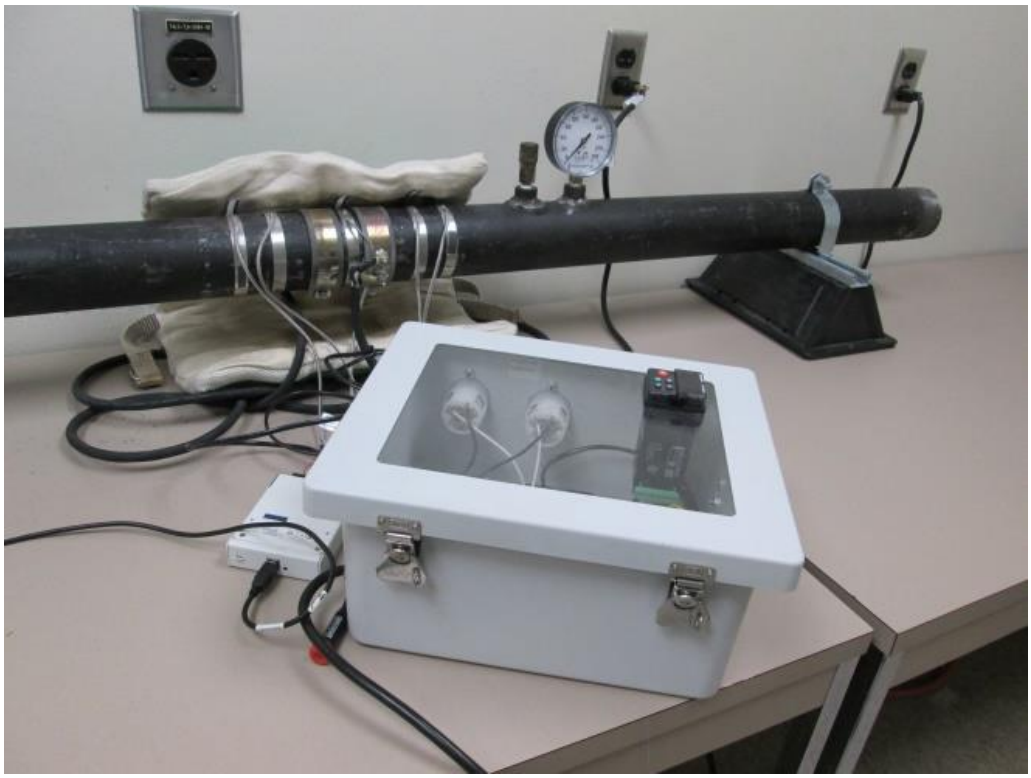

Figure 16. The physical test setup used in experimental work. Pipe heaters and thermocouples are attached to the compressed air line with an insulation blanket wrapped around the outside.

2.1.4.1.1 Experimental Test Procedure

Tests were run for a variety of flow rates between the minimum and maximum flow settings of the fan on a 2-inch and 3-inch pipe. Both pipes were schedule 40 black carbon steel. For each test, the fan was started a minimum of five minutes prior to activating the heaters. Once the heaters achieved the desired setpoint of 250 °F, the heaters maintained the temperature for ten minutes. Following these ten minutes, the heaters were turned off and data collection continued for a minimum of fifteen additional minutes. The temperature profile for three complete tests on the 2-inch pipe is shown in Figure 17. As seen, the controller is underdamped and results in decaying oscillations

around the setpoint of 250 °F. These oscillations settle out prior to turning off the heaters. Similar to simulation, the resulting temperature drop depends on the flow rate inside the pipe, with higher flow rates resulting in faster temperature drops.

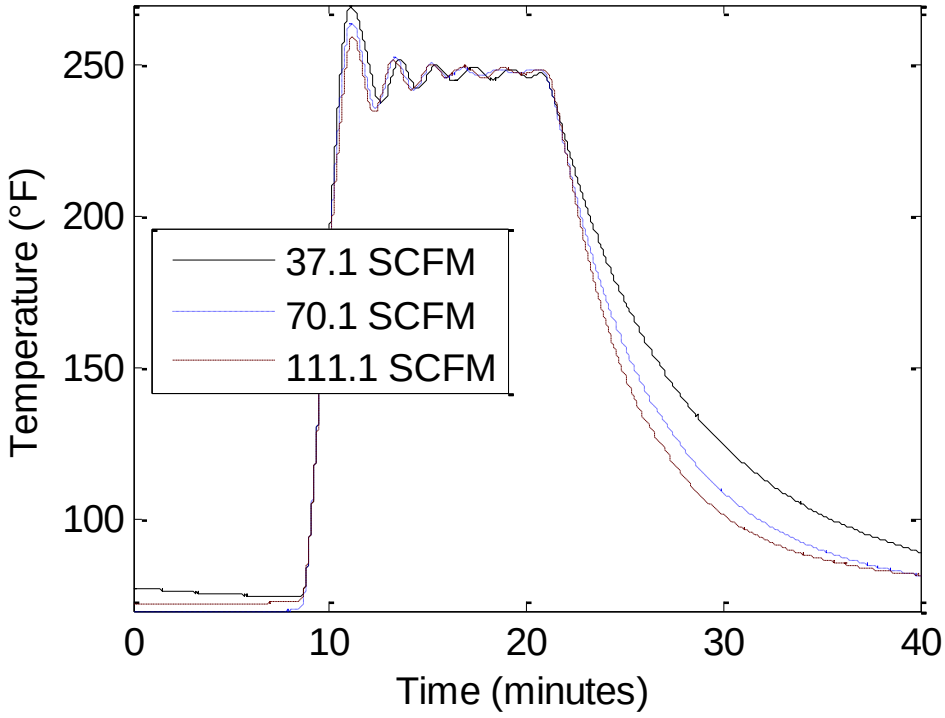

Figure 17. Complete temperature profile for three tests on the two-inch compressed air pipe. The decay rate of each test depends on the internal flow rate.

2.1.4.2 Experimental Test Results

In the same manner as simulation, correlations for flow rates were established based on the time for the temperature to reach a certain percentage of the time constant as well as the estimated exponential decay rate. The experimental results are displayed in Figure 18 and Figure 19 for the two pipes. The figures demonstrate the clear correlation

between flow rate and the different metrics from the temperature profile. Compared to simulation, the experimental results show very similar trends in the correlation, although the magnitude of each trend differs. These similar trends indicate that the simulations captured the underlying dynamics of the compressed air line, while differing magnitude values likely result from errors in heat transfer correlations, material parameters, or physical dimensions.

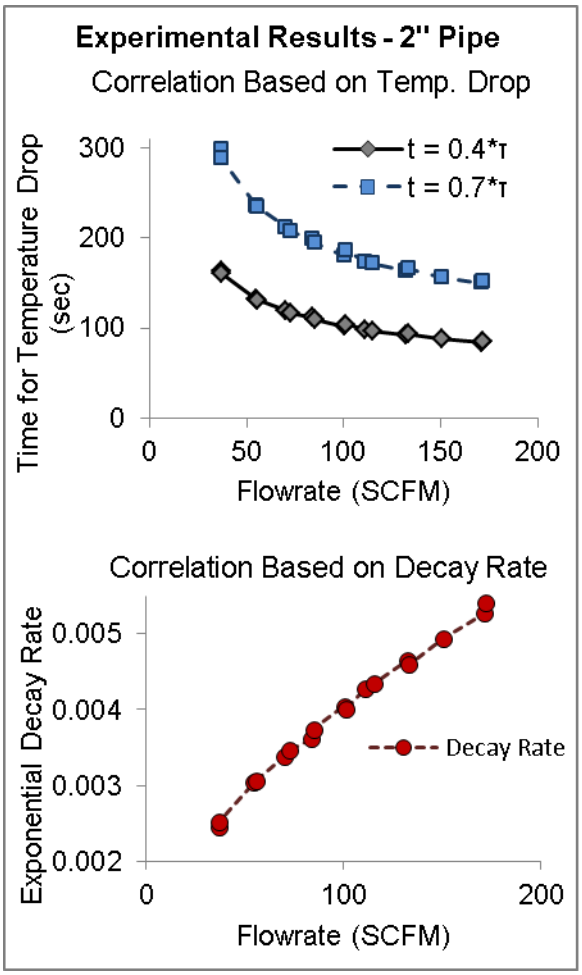

Figure 18. Experimental results for the two-inch pipe.

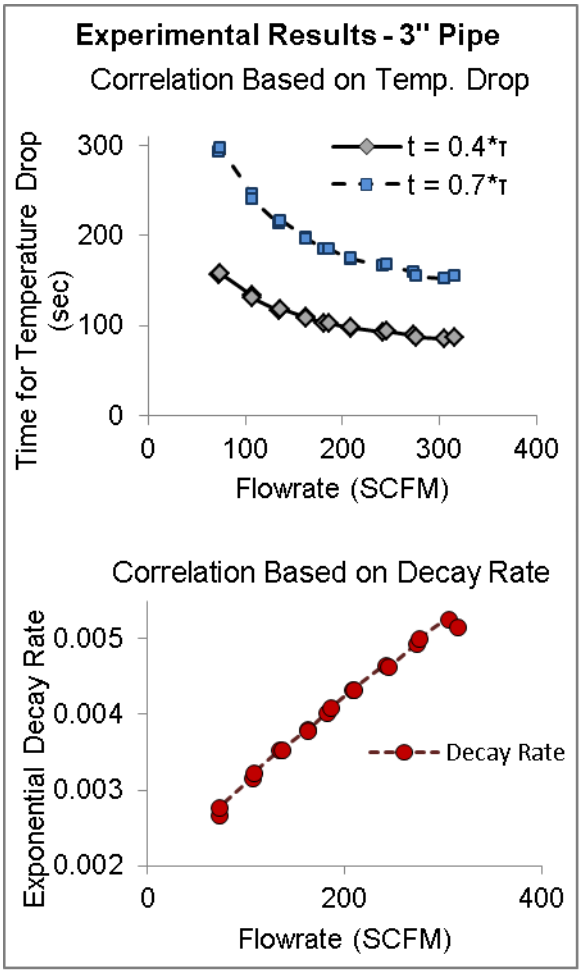

Figure 19. Experimental results for the three-inch pipe.

2.1.5 Field Test

Having verified the initial concept experimentally, the system was tested in an industrial setting to compare the test with other standard flow rate estimation methods. The compressed air system used for testing is shown in Figure 20. The system is run by a 15 hp, reciprocating compressor. A receiver tank is located directly under the compressor. The main header for the compressed air system is a two-inch copper compressed air line, with 70 feet of compressed air line in the compressor room.

Although the remaining compressed air lines in the facility could be measured, there is a minimum additional 350 feet of line in order to reach all areas of compressed air service. Although no additional receiver tanks were discovered in the facility, there is the possibility of an additional tank for air storage.

Figure 20. Compressed air system used in field test.

In the test case, the flow meter was installed on a straight section of 2” line. There was approximately 6 feet of straight line upstream of the installed meter and 30 feet of straight line downstream. Figure 21 shows the installed flowmeter on the line for the field test. The test was performed at a time with no facility usage, so the only compressed air flow resulted from leaks. During the test, a leak dropdown test was performed to measure air flow. Additionally, using the on-off cycles of the reciprocating

compressor, the duty factor was measured and related to compressor nameplate data to estimate the average flow rate.

Figure 21. Installed flowmeter on the 2" main header line.

Figure 22 shows the temperature of the thermocouple during the test. Similar to the calibration tests, the pipe was heated to 121 °C (250 °F), and the temperature was held for 10 minutes prior to turning off the heater. The temperature response was then observed as the pipe cooled. The decay rate was used to estimate the flow rate to account for the different pipe material and dimensions.

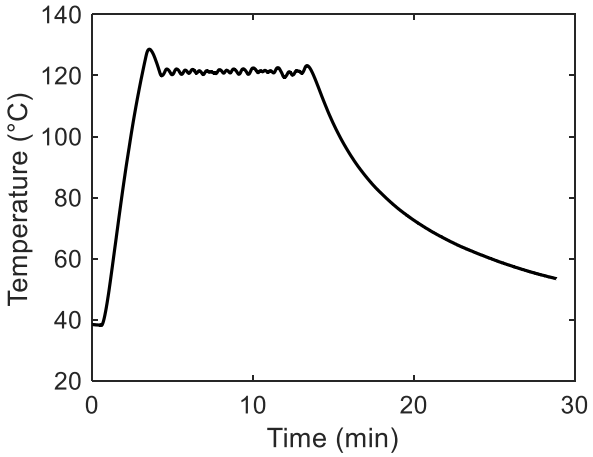

Figure 22. Original Temperature Profile

The original calibration tests were performed on black iron pipe, but the copper pipe has different material parameters is slightly thinner wall thickness. These parameters can be accounted for by using Newton's law of cooling. As discussed earlier in the chapter, linear conduction is assumed negligible through the testing procedure of maintaining the pipe temperature prior to turning off the heater. Under the eliminated conduction, the temperature profile is given by Equation (18):

$$T(t) = (T_o - T_\infty)e^{-\alpha t} + T_\infty \tag{18}$$

The decay rate, α, is defined in this equation according to Newton's law of cooling, as seen in Equation (19):

$$\alpha = z * h = \frac{4D_i}{\rho c\left(D_o^2 - D_i^2\right)} * h \tag{19}$$

The estimated decay rate is adjusted to account for the different parameters by looking at the ratio of the parameter, *z*, for each material according to Equation (20). The parameter values used for both the black iron and copper pipes are listed in Table 3.

$$\alpha = \frac{z_{\text{black iron}}}{z_{\text{copper}}}\alpha_{\text{measured}} \tag{20}$$

Table 3: Pipe Material Parameters

Parameter	Units	Black Iron	Copper
ρ	kg/m^3	7854	8960
c	kJ/kgK	434	385
D_i	m	0.0525	0.0498
D_o	m	0.0603	0.0540

A correlation between flow rate and the decay rate from the initial tests on the black iron pipe was developed based on a power law. This power law relationship was then used to estimate the flow rate of the field test. The measured decay rate of the test, the scaled decay rate, and the corresponding estimated flow rate are displayed in Table 4. Additionally, the estimated flow rate from the test case is displayed together with the tests on the black iron pipe for the 2" pipe in Figure 23. The test case flow rate actually falls outside the lowest estimated flow rate from the initial tests. Since extrapolation can lead to extensive errors, there is uncertainty in the exact estimated flow rate using the previous tests on the black iron pipe. However, the estimated flow rate will confidently fall below the lowest measured flow rate of 37 SCFM from the initial experimental tests.

Table 4. Test Case Results

Parameter	Units	Value
Measured Decay Rate	s^{-1}	0.00328
Scaled Decay Rate	s^{-1}	0.00145
Estimated Flow Rate	SCFM	12.7

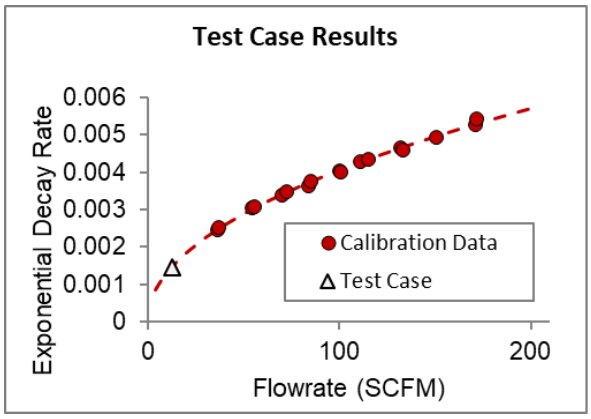

Figure 23. Estimated flow rate of the test case.

2.1.5.1 Comparison to Other Flow Estimation Methods

2.1.5.1.1 Duty Factor of Compressor

The cycle time of the compressor was measured to estimate the flow rate based on the duty factor. The duty factor is measured as the fraction of time the compressor is running during one loading cycle. The compressor specific efficiency, which estimates the power required to produce one SCFM, can be utilized with the power of the compressor to estimate average system flow rate over one cycle. The measured duty cycle, compressor efficiency, motor power and the estimated flow rate are shown in Table 5. While the specific efficiency is dependent on the condition of the compressor, operating pressure, and other environmental conditions, the uncertainty due to these parameters is limited.

Table 5. Estimated Flow Rate Using Compressor Duty Cycle

Parameter	Units	Value
Duty Factor	N/A	0.166
Motor Power	hp	15
Compressor Efficiency	SCFM/hp	3.53
Estimated Flow Rate	SCFM	8.8

2.1.5.1.2 Drop-down Test

A common method to estimate flow rate in the Industrial Assessment Center is to perform a drop-down test, where the volume of the system is correlated to the time required for a specified pressure to drop off the system. The correlation used for estimated flow rate is given in Equation (21):

$$\dot{V} = \frac{(P_{drop})(V_{sys})}{(P_{atm})(t_{drop})} \quad (21)$$

The pressure drop was recorded and timed at the main receiver tank for the system. The calculated volume was estimated based on the receiver tank, visible lines in the compressor room, and a header line that runs through the facility. The header line is assumed to be two inches, but the lines are inaccessible and couldn't be verified. Uncertainty in compressed air system volume represents one of the primary limitations in a leak drop down test. Table 6 shows the calculated parameters and estimated flow rate using the drop-down test method.

Table 6. Estimated Flow Rate Using the Drop-down Test

Parameter	Units	Value
Pressure Drop	psi	6
Time for Drop	min	4.92
System Volume	ft^3	23.5
Atmospheric Pressure	psi	14.7
Estimated Flow Rate	SCFM	2.0

2.1.5.2 Summary of Results

All the methods for flow estimation compare favorably for the test case. With typical flow rates between 0-200 SCFM for a 2" line, all three methods consistently measure a relatively low flow rate. As the duty factor estimate of flow rate is likely to have small variations, the drop-down test illustrates the difficulty of measurement if not all of the compressed air lines can be measured in the facility. The drop-down tests significantly underpredicts the flow rate compared to the duty factor method and the proposed flow meter in this chapter. The most like cause of this underprediction is that not all of the volume in the system was measured for the calculation, highlighting one of the disadvantages of this method.

2.1.6 Sensitivity of Method to Parameter Uncertainty

One challenge with the method presented in this chapter is the sensitivity of flow measurement to other parameters. Uncertainty in pipe material parameters and dimensions affect the estimate of flow inside of the pipe. Manufacturing processes and tolerances affect these parameters, which are not immediately known during testing. Flow measurement sensitivity to parameters occurs in other non-intrusive flow meters; this issue is further explored and addressed in Chapter 3.

2.2 Summary

This chapter proposed a non-intrusive flow meter that uses the thermal response on the exterior surface of a compressed air line to estimate the flow inside the pipe. A pipe is heated to a specified temperature and maintained for a period of time. Once the

heater is turned off, the rate at which the temperature drops is observed. Since the cooling rate depends on flow inside the pipe, a correlation is established between cooling rate and flow.

The method was tested in both simulation and experiment, and both methods demonstrated the ability to measure flow. Extension work includes verification of the method at high pressures and in actual industrial environments. Additional metrics of the temperature profile may be explored, including using power consumption to maintain a steady temperature or determining the steady-state temperature that is approached from a constant heat flux.

CHAPTER 3

ENHANCING ADVANCED SENSORS THROUGH DATA EXPLOITATION

The previous chapter presented how physical relationships can be used to create virtual sensors for parameters that are difficult to measure directly. Use of virtual sensors can drive down cost of measurement and increase which parameters can be measured. These simple relationships at times fail to incorporate all the physical effects on the measurement, resulting in measurement error or confounding factors. This chapter presents a case study on how processing the input data can improve overall parameter measurement. The work presented in this chapter extends the research in the previous chapter on the non-intrusive flow meter by exploiting the measurement data to improve the overall flow estimation. This is accomplished by simultaneously and dynamically determining pipe material parameters in addition to the flow measurement.

The outline for the chapter is as follows. In this case study, previous research and potential application areas of the flow meter is summarized. The mathematical and numerical model used for dynamic parameter estimation is then presented. The gradient method used for parameter estimation is discussed, along with the approach to determine initial estimates of parameter values. The effectiveness of the dynamic parameter estimation and air flow measurement is then presented for both simulation and experimental systems. The chapter concludes by discussing two test cases in industrial environments.

3.1 Case Study: Dynamic Parameter Estimation for Improved Non-Intrusive Compressed Air Flow Measurement using Thermal Signatures

Fluid flow in pipes is a key component in many industrial systems, including process heating & cooling, building space conditioning, oil & gas production, and steam systems. When evaluating these systems for potential energy savings or health diagnostics, one of the critical measurements is the flow rate of the working fluid. Current methods for non-intrusive, non-invasive methods are expensive and often require time-consuming calibration. Research in the literature on potential low-cost flow rate measurement relies on knowledge of pipe parameters, pipe dimensions and fluid properties.

This case study presents work on using dynamic parameter estimation using the gradient descent method with the thermal response of a pipe to measure flow. Using a constant heat flux and derived analytical model, the thermal diffusivity, heat capacity term, and internal flow rate are estimated online during data collection. These parameters are identified simultaneously to isolate the estimation of flow inside the pipe. Initial parameter estimates are dynamically determined to improve convergence by leveraging unique relationships in the thermal response. Simulation and experimental results show the potential of the method, and an additional two tests cases in industrial environments are presented.

3.1.1 Background and Introduction

In diagnosing or analyzing process systems, flow measurement inside of pipes is a necessary measurement that is difficult to obtain. There are many applications in which

one-time spot measurements are desirable to assess the health of the system, such as steam, process hot and chilled water or compressed air systems [47, 48, 53]. Other process systems, such as production water in the gas industry or brine in the geothermal industry contain corrosive fluids that quickly destroy inline sensors [46, 51, 52]. For both one-time measurements and corrosive fluids, use of non-intrusive, non-invasive flow meters are highly desirable.

The use of non-intrusive meters for flow measurement has historically been limited. The primary commercially-available meter is an ultrasonic flow meter, which uses ultrasonic waves to ascertain the fluid speed inside the pipe. This method comes with several limitations that restrict its widespread adaptation in many industries. Beyond its high cost, the meter requires extensive calibration and placement in specific pipe sections to ensure accurate flow measurement. The calibration requires knowledge of several important parameters, including the speed of sound in the fluid of interest, pipe thickness, pipe material parameters, and pressure inside the pipe. These values can be difficult to obtain, and the parameter uncertainties contribute to total measurement uncertainty.

Other types of non-intrusive flow meters have been proposed and studied in the literature. Several researchers have examined the use of turbulent-induced pipe vibration or acoustics to measure flow [51, 53, 56, 57]. They found that the standard deviation of the vibration is proportional to the velocity squared of the fluid. Early experimental results demosntrated the validity of the method. However, the method requires knowledge of the pipe materials, fluid properties, pipe support structure, and the pipe

dimensions [58]. Uncertainties in these parameters and disturbance vibrations from pumping or compressor equipment limit the widespread use of this method [52]. Others have explored using the thermal response of the pipe to estimate fluid flow inside the pipe [59]. They too note the strong dependence of flow on external parameters. Since the commercially-available ultrasonic meter and other researched non-intrusive flow meters in the literature depend on pipe parameters, improved knowledge of these parameters will decrease uncertainty in flow measurement. Many of these parameters are difficult to determine independently once a pipe is installed and in use.

There has been limited work in the literature on parameter estimation for improvement in fluid flow measurement [60]. In the referenced paper, the researchers used a genetic algorithm with numerical simulations to match the simulated response to the experimental data. Prior to using parameter estimation, the researchers found high uncertainty in the initial correlation between the temperature response and fluid flow. The parameter estimation improved the overall estimation and helped insulate the uncertainty against these parameter values.

3.1.2 Effect of Parameter Uncertainty on Flow Estimation

Previous work by the authors used the temperature response on the exterior surface of the pipe to estimate gas flow rate inside a compressed air line. This work showed that there is a strong level of correlation between flow rate and a first order temperature step response. However, accurate estimation of the flow rate depends on knowledge of multiple pipe parameters, several of which are not readily available. The relationship between various parameters and their effect on flow measurement was

explored in simulation for the flow meter presented in Chapter 2. Figure 24 shows the effect on flow measurement for a 1% change in the indicated parameters. In the previous chapter, the two metrics used to estimated flow were the decay rate (α) and the time required to drop a specified temperature (t_1). As seen in the figure, changes in the heat capacity, thermal conductivity, and pipe thickness all affect the measured flow rate. For other transient methods of flow estimation, thermal conductivity will also have a meaningful effect on the estimated flow.

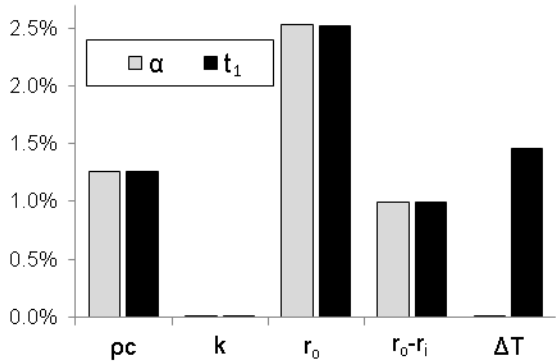

Figure 24. Percent change in flow estimation to a 1% change in the specified parameters.

In addition to variations in material parameters, long measurement times and power requirements limit the effectiveness of the proposed method in the previous chapter. Although longer measurement times are often tolerable for one-time measurements, the proposed setup requires continual heating of the pipe for approximately 10-15 minutes, following which the pipe must cool down sufficiently

long before the decay rate can be determined. The long heating time results in large power requirements for the two heaters, limiting portability of the proposed method.

3.1.3 Parameter Estimation

3.1.3.1 Introduction

The nonintrusive flow meters discussed in the previous section all rely on knowledge of multiple parameter for accurate estimation of flow. To accurately account for this, extensive testing is required to build calibration curves that account for changes in each of these properties. By simultaneously measuring these parameters along with the flow term, the effect of these terms can be separated from the flow measurement.

Thermal parameters, such as the thermal diffusivity, are not typically measured directly for a material. Instead, many methods to measure thermal parameters use energy signals and the resulting temperature profiles of a sample. A common way to determine the thermal conductivity is with the laser flash method, which uses a laser pulse and the resulting temperature profile to derive the thermal diffusivity [61]. Knowledge of the heat capacity with the measured thermal diffusivity allows the ability to calculate the thermal conductivity, a difficult parameter to measure. Improvements and adaptations of this method have resulted in multiple ways to determine the thermal diffusivity, heat capacity, and related thermal conductivity. These methods rely on an induced temperature response with mathematical relations to determine the parameters of interest.

3.1.3.2 Alteration to the Initial Proposed Steady State Method

The method proposed in the previous chapter works by heating a pipe with two heaters on either side of a thermocouple. The pipe temperature is maintained at a specified temperature for a period of time, after which the heaters are turned off. The period with elevated temperature is maintained long enough to reduce the effect of axial conduction along the pipe. When the heater is turned off, the temperature response of the pipe is analyzed. This is a first order system with a step response, and so the decay rate can be determined from the exponential decay. Figure 17 shows the complete temperature profiles for three experimental tests, with a clear distinction between the decay rate and different flow rates. The convection term is independent of pipe pressure, increasing the robustness of measurement.

The proposed system in this chapter relies on dynamic parameter estimation during the initial warmup stage of the previously-described method. The parameters are determined online and flow rate calculated within approximately five minutes, reducing both measurement time and power requirements. Additionally, the method only requires a single heater and multiple thermocouples, thereby further reducing power requirements for the system. A diagram of the proposed system is displayed in Figure 25. The thermocouples are placed upstream of the heater, and the power supply delivers a constant current to the heater. As the heater warms the pipe, the temperature profiles from the thermocouples are dynamically analyzed. In this work, only the initial heating of the pipe is considered. Potentially improved estimation of flow could be achieved by

also recording the temperature profile of the thermocouples as the pipe cools down once the heater is turned off.

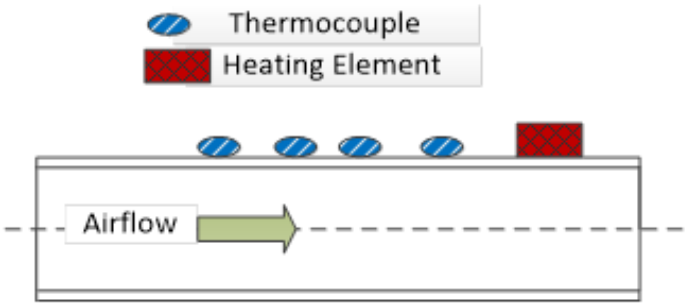

Figure 25. Proposed setup for parameter estimation method of measuring flow.

3.1.3.3 Analytical Model Derivation

The analytical model used for the dynamic parameter estimation is derived from heat transfer first principles. The pipe is modelled as a one-dimensional pipe with internal heat transfer due to the fluid flow in the pipe. The system is assumed to be axisymmetric and lumped capacitance is used in the radial direction in order to use a one-dimensional model. The following assumptions are made in the development of the model:

- Due to lumped capacitance in the radial direction, the temperature at any location along the pipe is uniform in the radial direction of the pipe. In other words, the pipe is sufficiently thin and conductivity is sufficiently high to assume that the pipe temperature is the same for all radii at any given location along the pipe.

- The pipe is sufficiently long that it is assumed to be infinite in length. This is a valid assumption if the effects from the boundary conditions at the end of the pipe are negligible.
- The boundary condition at x=0 is modeled by a heat flux into the pipe. For the initial modeling, a constant heat flux is assumed during heating, but any function of time can be used. Note that when no heat is added into the pipe, this heat flux is zero (an adiabatic condition). Due to symmetry in the problem, this assumption can still be used with no heat flux.
- The flow inside of the pipe is steady. The constraint on the physical system from this assumption is the need for relatively constant flow during data collection.
- The pipe is initially at a uniform temperature.
- Specific heat, density, and conductivity of the material is constant with both space and time.
- The outside of the pipe is insulated.
- The flow through the pipe remains a constant temperature.

Using these assumptions, the model is derived by analyzing the heat terms in a small section of pipe. The setup of the pipe is shown in Fig. 5. Using the first law of thermodynamics, the change in energy of the pipe is equal to the sum of the heat added and removed from the system as seen in Equation (22).

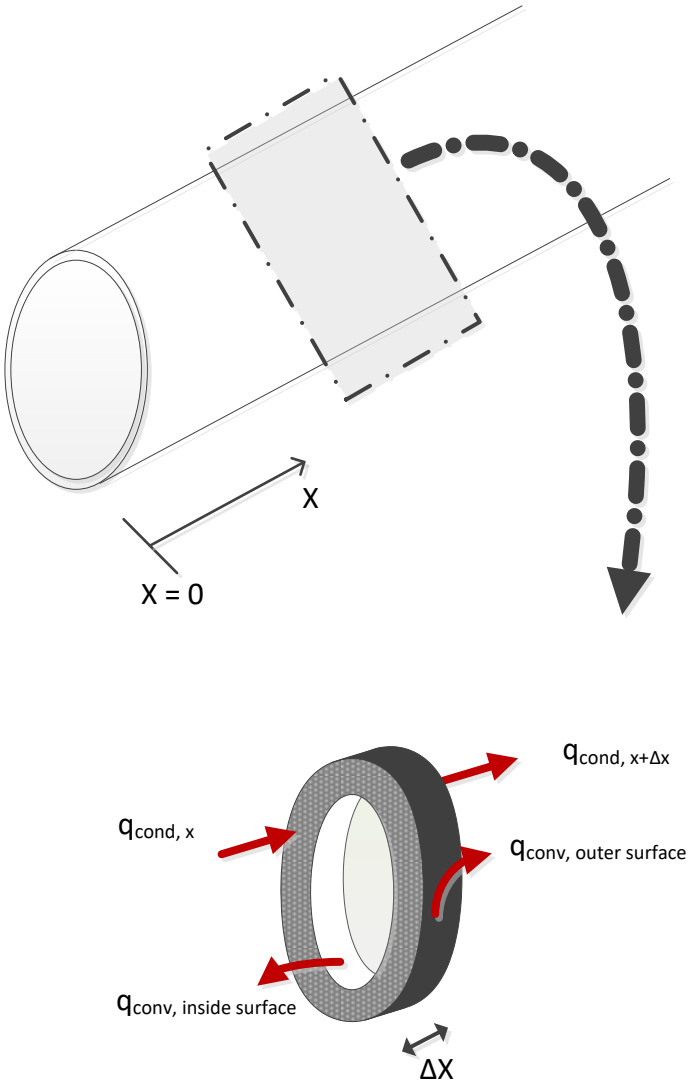

Figure 26. Energy balance used for analytical model.

(22) $$\frac{\partial E}{\partial t} = q_{cond,x} - q_{cond,x+\Delta x} - q_{conv,o} - q_{conv,i}$$

Insulation is added to the exterior surface of the pipe, so the exterior surface is approximately adiabatic. Newton's law of cooling is applied for the interior surface of the pipe, and Fourier's law of heat conduction is use in the axial direction. The density,

specific heat, and pipe dimensions are assumed to be constant with space and time. After applying these laws and simplifying the equation consistent with constant material parameters, the resulting partial differential equation is given by Equation (23).

$$\frac{\partial\left(\rho c A_c \Delta x T\left(x+\frac{\Delta x}{2},t\right)\right)}{\partial t} = -\frac{k A_c \, \partial T(x,t)}{\partial x} + \frac{k A_c \, \partial T(x+\Delta x,t)}{\partial x} - h P_i \Delta x (T(x,t) - T_\infty), \tag{23}$$

After applying these laws and simplifying the equation consistent with constant material parameters, the resulting partial differential equation is given by Equation (24).

$$\frac{\partial u(x,t)}{\partial t} = \alpha \cdot \frac{\partial^2 u(x,t)}{\partial x^2} - b \cdot u(x,t) \tag{24}$$

In this equation, the lumped parameters are defined in the nomenclature. In physical terms, α is the thermal diffusivity and the b term contains the convection coefficient. The boundary condition at the heater $(x=0)$ is of the second type, which is physically described as a heat flux at the boundary. No physical condition is specified for the second boundary condition. Instead, the boundary condition is formed by recognizing that as $x \rightarrow \infty$, the temperature must remain bounded. The initial temperature of the pipe is assumed to be uniform throughout at the same temperature. Since the temperature scales can be arbitrarily shifted, all temperatures are shifted by the ambient temperature so that the initial condition is equal to zero throughout. These boundary

conditions and initial conditions are formalized in Equation (25) and Equation (26) . Note that positive Q_c indicates heat transferred into the system.

Boundary Conditions

(25)

$$1.\ \frac{\partial u(0,t)}{\partial x} = \frac{-q_{in}}{k} = -Q_c$$

$$2.\ u(x \to \infty, t) < \infty$$

3.1.3.3.1 Initial Condition

(26)

$$1.\ u(x,0) = \phi(x,0) = 0$$

The solution to this partial differential equation can be solved by the method of separation of variables by assuming and verifying that the time component and space component of the solution are independent of each other. The solution is given by Polyanin [62] as seen in Equation (27).

(27)

$$u(x,t) = \int_0^\infty G(x,\sigma,t) \cdot f(\sigma)\, d\sigma - \alpha \int_0^t g(\tau) \cdot G(x,0,t-\tau) d\tau$$

where Equation (28) formalizes the function to in the integral.

(28) $$G(x,\sigma,t) = \frac{e^{-bt}}{2\sqrt{\pi\alpha t}}\left\{e^{\frac{-(x-\sigma)^2}{4\alpha t}} + e^{\frac{-(x+\sigma)^2}{4\alpha t}}\right\}$$

When the zero initial conditions are accounted for and the heat flux is assumed to be constant, the first term in the equation equals zero. The remaining simplified term is given by Equation (29).

(29) $$u(x,t) = Q_c\sqrt{\frac{\alpha}{\pi}}\int_0^t \frac{e^{-b(t-\tau)}}{\sqrt{(t-\tau)}} e^{\frac{-x^2}{4\alpha(t-\tau)}} d\tau$$

The integral is simplified by introducing a substitution variable, $y = \sqrt{t-\tau}$, and the temperature profile is solved in [63]. After applying the limits of integration, the temperature profile is given by Equation (30).

(30) $$u(x,t) = \frac{1}{2}Q_c\sqrt{\frac{\alpha}{b}}\left[\begin{array}{c} e^{x\sqrt{\frac{b}{\alpha}}}\left\{\operatorname{erf}\left(\sqrt{bt}+\sqrt{\frac{x^2}{4\alpha t}}\right)-1\right\} \\ +e^{-x\sqrt{\frac{b}{\alpha}}}\left\{\operatorname{erf}\left(\sqrt{bt}-\sqrt{\frac{x^2}{4\alpha t}}\right)+1\right\}\end{array}\right]$$

In the equation, the time and space terms both appear in the error function terms. As the convection term approaches zero, the term on the left side of the equation approaches infinity while the error function initially starts at zero. But as time increases, the absolute value term of the function increases but remains finite. Therefore, this

model is only valid for $b > 0$. A separate model must be used when comparing the temperature profile for a no-convection case. Appendix B shows the model derivation and solution for the no-convection case.

By introducing the non-dimensional parameters $f=\sqrt{x^2/4\alpha t}$, $g=\sqrt{bt}$, and the parameter $w=Q_c x/4$, the temperature profile reduces to the Equation (31). This analytical solution is used in parameter estimation and flow measurement. Figure 27 shows the temperature profile as given by Equation (31).

$$u(x,t) = \frac{w}{fg}\left[e^{2fg}\{\mathrm{erf}(g+f)-1\}+e^{-2fg}\{\mathrm{erf}(g-f)+1\}\right] \tag{31}$$

The temperature solution is completely defined by these three parameters. Estimation algorithms are used to determine these parameters, from which physical parameters may be derived.

3.1.3.3.2 Temperature Derivative Derivation

The slow dynamics of thermal systems enables the ability to captures the time derivative of the temperature in experimental tests with reasonable accuracy. The shape of the time derivative profile is more contoured than the temperature profile. As such, the time derivative time profile is also used in the parameter estimation algorithm. The complete time derivative of the temperature profile is given by Equation (32).

$$\frac{du(x,t)}{dt} = \frac{\partial u}{\partial g}\frac{dg}{dt}+\frac{\partial u}{\partial f}\frac{df}{dt}+\frac{\partial u}{\partial w}\frac{dw}{dt} \tag{32}$$

where the partial derivatives of the non-dimensional parameters are given by Equation (33).

(33) $$\frac{\partial g}{\partial t} = \frac{g}{2t} \qquad \frac{\partial f}{\partial t} = -\frac{f}{2t} \qquad \frac{\partial w}{\partial t} = 0$$

The resulting time derivative, after all substitutions, is given by Equation (34).

(34) $$\frac{du(x,t)}{dt} = -\left(\frac{f}{2t}\right)\left(e^{2fg}[\mathrm{erf}(g+f)-1]\left[\frac{2w}{f}-\frac{w}{f^2 g}\right] - e^{-2fg}[\mathrm{erf}(g-f)+1]\left[\frac{2w}{f}+\frac{w}{f^2 g}\right]\right) + \left(\frac{g}{2t}\right)\left(\begin{array}{c} e^{2fg}[\mathrm{erf}(g+f)-1]\left[\frac{2w}{g}-\frac{w}{fg^2}\right] - e^{-2fg}[\mathrm{erf}(g-f)+1]\left[\frac{2w}{g}+\frac{w}{fg^2}\right] \\ +\frac{2w}{fg}\frac{2}{\sqrt{\pi}}e^{-g^2-f^2} \end{array}\right)$$

This equation can be simplified, and the simplified time derivative of the temperature profile is given by Equation (35).

(35) $$z = \frac{du(x,t)}{dt} = \left(\frac{2w}{ft\sqrt{\pi}}e^{-g^2-f^2}\right)$$

3.1.3.4 Gradient Descent Method Derivation

The proposed method estimates parameters using the gradient descent method during data collection to match the actual temperature profile with the analytically-

derived temperature profile. This method has been widely integrated in industry for many years across many applications [64]. The three parameters that are dynamically estimated are those found in Equation (31). Once these are determined, the thermal parameters can be derived from the non-dimensional terms. Specifically, the unknown parameters to be estimated are the thermal diffusivity of the pipe material, the lumped heat input term of the pipe, and the lumped convection term. By determining these values, other individual terms can be derived if desired. Additionally, another parameter of interest is the distance from the heater to the first thermocouple. Although a nominal distance can be measured, the physical edge of the heater may not correspond to the actual end of heat input. The distance between thermocouples is known with a high degree of accuracy, and so the offset from the physical edge of the heater to the thermal edge of the heater becomes a fourth parameter to estimate. The vector of estimated parameters is shown in Equation (36).

(36) $$\theta=[\alpha \;\; b \;\; Q_c \;\; x_o]$$

The parameter estimation problem is framed by using a cost function based on the error of the temperature profile as a function of the parameter values in Equation (37).

(37) $$e(\theta)=u(\theta) - u_m$$

The objective function, as seen in Equation (38),is a function of the squared error and is minimized to converge on parameter values.

(38) $$J(\theta)=e(\theta)^2/2$$

For each iteration, the update law for the parameters is based on the gradient of the objective function and is given in Equation (39).

(39) $$\theta_{k+1}=\theta_k+\lambda\nabla J(\theta_k)=\theta_k+\lambda e(\theta_k)\nabla u(\theta_k)$$

For the analytical solution seen in Equation (31), the gradient of the temperature profile with respect to the parameters is shown in Equation (40).

(40) $$\nabla u(\theta)=\left[\frac{\partial u}{\partial f}\frac{\partial f}{\partial \alpha} \quad \frac{\partial u}{\partial g}\frac{\partial g}{\partial b} \quad \frac{\partial u}{\partial w}\frac{\partial w}{\partial Q_c} \quad \frac{\partial u}{\partial f}\frac{\partial f}{\partial x_o}+\frac{\partial u}{\partial w}\frac{\partial w}{\partial x_o}\right]$$

3.1.3.4.1 Partial Derivatives of All Equations of Interest

When using estimation techniques to derive the parameters, the relationship of how varying the different parameters affect the temperature profile needs to be known. In order to derive the relationship, each partial of the three dimensionless parameters in relation to the temperature profile is derived. These partial derivatives are shown in Equation (41). Individual steps in the solution are removed for brevity.

3.1.3.4.1.1 Partial Derivatives with Respect to Non-Dimensional Parameters

$$\frac{\partial u}{\partial f} = e^{2fg}[\mathrm{erf}(g+f)-1]\left[\frac{2w}{f}-\frac{w}{f^2 g}\right] - e^{-2fg}[\mathrm{erf}(g-f)+1]\left[\frac{2w}{f}+\frac{w}{f^2 g}\right] \tag{41}$$

$$\frac{\partial u}{\partial g} = e^{2fg}[\mathrm{erf}(g+f)-1]\left[\frac{2w}{g}-\frac{w}{fg^2}\right] - e^{-2fg}[\mathrm{erf}(g-f)+1]\left[\frac{2w}{g}+\frac{w}{fg^2}\right] + \frac{2w}{fg}\frac{2}{\sqrt{\pi}}e^{-g^2-f^2}$$

$$\frac{\partial u}{\partial w} = \frac{1}{fg}\left[e^{2fg}\{\mathrm{erf}(g+f)-1\} + e^{-2fg}\{\mathrm{erf}(g-f)+1\}\right]$$

3.1.3.4.1.2 Partial Derivatives with Respect to Estimated Parameters

The derivatives of the partials of non-dimensional parameters are taken with respect to the different parameters of interest. Only non-zero derivatives are shown in Equation (42). Any derivative that is not shown is equal to zero.

$$\frac{df}{d\alpha} = -\frac{1}{2}\sqrt{\frac{x^2}{4\alpha^3 t}} \qquad \frac{df}{dy_o} = \frac{1}{\sqrt{4\alpha t}} \qquad \frac{dg}{db} = \frac{1}{2}\sqrt{\frac{t}{b}} \tag{42}$$

$$\frac{dw}{dQ_C} = \frac{x}{4} \qquad \frac{dw}{dx_o} = \frac{Q_c}{4}$$

The partial derivatives in this section are used in Equation (40) for the gradient descent algorithm.

3.1.3.4.1.3 Partial Derivatives of the Temperature Derivative Equation

Similarly, to use the parameter estimation algorithm on the profile of the time derivative of the temperature, the partial derivatives of Equation (35) with respect to the parameters of interest much be known. These partial derivatives are shown below with respect to the estimated parameters. The intermediate step of taking the partials with respect to the non-dimensional parameters and the partials of the non-dimensional parameters with respect to the estimated parameters is omitted for brevity.

(43)
$$\frac{dz}{d\alpha} = \frac{w}{t\alpha\sqrt{\pi}}\left(\frac{1}{f} + 2f\right)e^{-g^2-f^2}$$

$$\frac{dz}{db} = -\frac{2wg^2}{fbt\sqrt{\pi}}e^{-g^2-f^2}$$

$$\frac{dz}{dQ_C} = \frac{x}{2ft\sqrt{\pi}}e^{-g^2-f^2}$$

$$\frac{dz}{dy_o} = \left(\frac{Q_C}{2ft\sqrt{\pi}} - \frac{w}{t\sqrt{\pi\alpha t}}\left(\frac{1}{f^2} + 2\right)\right)e^{-g^2-f^2}$$

3.1.3.5 Initial Parameter Guesses

The gradient descent algorithm only works for sets in which the problem is convex for the correct parameter values. In simulations of parameter estimation using (6), if initial parameter values are chosen too far away from the actual parameter values, they will converge to non-physical, i.e. negative, values. Initial parameter guesses need to be selected close enough to the actual parameter values to ensure convergence to correct values. Thermal material parameters vary significantly between metals, as is seen

in steel, aluminum, or copper piping. Even between pipes of the same nominal material, wide manufacturing tolerances result in considerable differences in thermal parameters.

The unique temperature profile in this problem is leveraged in order to find the initial parameter estimates online during data collection. The temperature profile has a distinct inflection point, resulting in a global maximum in the derivative of the temperature profile. This inflection point in the temperature profile and maximum point in its derivative are shown in Figure 29.

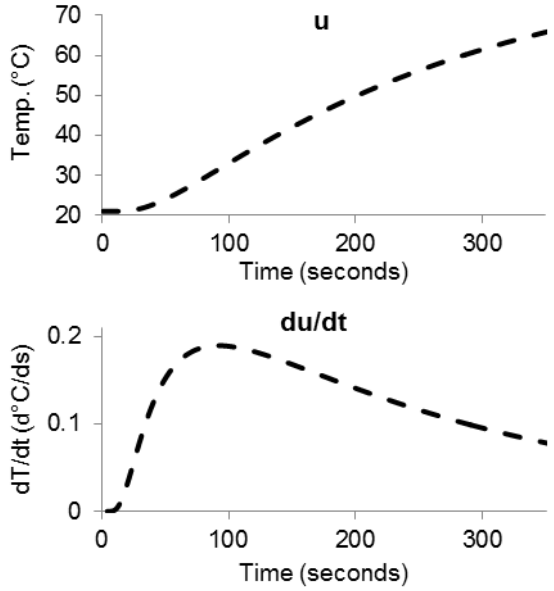

Figure 27. Example of the temperature profile and its derivative.

3.1.3.5.1 Complete Second Derivative of the Temperature Solution

The inflection point occurs when the time derivative of the temperature reaches a maximum value or when the second time derivative of the temperature profile crosses

zero. The first time derivative of the profile is shown in Equation (35), and the second time derivative is given by Equation (44).

(44) $$\frac{dz(x,t)}{dt} = \frac{d^2u(x,t)}{dt^2} = \frac{\partial z}{\partial t}\frac{dt}{dt} + \frac{\partial z}{\partial f}\frac{df}{dt} + \frac{\partial z}{\partial g}\frac{dg}{dt}$$

Substitution of all the appropriate partial derivatives results in Equation (45).

(45) $$\frac{dz(x,t)}{dt} = \left(\frac{-2w}{ft^2\sqrt{\pi}}e^{-g^2-f^2}\right) + \left(\frac{-2w}{t\sqrt{\pi}}\left(\frac{1}{f^2}+2\right)e^{-g^2-f^2}\right)\left(-\frac{f}{2t}\right) + \left(\frac{-4wg}{ft\sqrt{\pi}}e^{-g^2-f^2}\right)\left(\frac{g}{2t}\right)$$

This equation can be simplified to its final form in Equation (46).

(46) $$\frac{dz(x,t)}{dt} = \frac{2w}{\sqrt{\pi}}e^{-g^2-f^2}\left(-\frac{1}{2ft^2}+\frac{f}{t^2}-\frac{g^2}{ft^2}\right)$$

When looking for a zero crossing, the exponential term and scaling term do not affect the zero-crossing location. The zero-crossing location is determined exclusively by the last term in the equation. This term is set to zero, as seen in Equation (47).

(47) $$\left(-\frac{1}{2ft^2}+\frac{f}{t^2}-\frac{g^2}{ft^2}\right) = 0$$

Rewriting this equation in terms of the parameters of interest (α, b, and Qc) and simplifying the equation, the time of the inflection point is given by Equation (48).

(48) $$t = \frac{1}{4b}\left(-1 + \sqrt{1 + \frac{4bx^2}{\alpha}}\right)$$

The time location of the peak in the temperature derivative is independent of heat added, and is dependent on the thermal diffusivity, lumped convection term, and the location of the temperature measurement. The Q_C term only appears in w, which is a scaling factor and does not affect the peak time location. This equation can be re-written in matrix form in order to solve for α and b in the relation as shown in Equation (49). The peak locations from two temperature profiles are required to determine the parameters α and b. Once these initial parameter guesses are estimated during data collection, the initial value of Q_c is estimated by scaling the temperature derivative equation so that the analytical peak value matches the measured peak value. Although stable across a wide range of values, the starting x0 value is determined by stepping through a small (10 points) vector of possible x0 values. The matrix given in Equation (49) is solved for each iteration and the value of x0 that results in the lowest total profile error is selected at the initial parameter value.

(49) $$\begin{bmatrix} 4{t_1}^2 & -{x_1}^2 \\ 4{t_2}^2 & -{x_2}^2 \end{bmatrix} \begin{bmatrix} b \\ 1/\alpha \end{bmatrix} = \begin{bmatrix} -2t_1 \\ -2t_2 \end{bmatrix}$$

3.1.4 Simulation Results

A numeric model was developed to validate the dynamic parameter estimation methodology as described in previous sections. The model setup is displayed in Fig. 7. This two-dimensional model allows for both axial and radial conduction with heat transfer on the inner surface of the pipe. The model was transformed into a state space representation. Multiple inputs and outputs are allowed in the model, but only a single input was used with multiple outputs in the simulation. The input to the model is the signal given to the heater, while the outputs are temperatures on the outer surface on the pipe.

The simulation was set up for the properties of a 3" Schedule 40 black steel pipe with air flowing in the pipe. The input was a constant heat flux for six minutes. Four locations on the exterior surface of the pipe were set up as outputs, and the location of the outputs matched similar locations to the experimental setup discussed in the following section. Simulations were run for flow rates between 30 SCFM and 300 SCFM, which represents realistic flow rates seen in industrial compressed air systems for 3-Inch Schedule 40 pipe. For each simulation, the four temperature outputs were used with the parameter estimation to determine the convection term b, the thermal diffusivity, and the heat flux term. The methodology in the previous section was used for the initial parameter guesses. The results of the final estimated parameter values from the gradient descent method are displayed in Figure 28. for the simulations at the different flow rates. As seen in the figure, the estimated final values correspond closely to the actual parameter values specified in the simulation. This result represents an

idealization of an actual system as it does not include disturbances of sensor noise or unmodeled dynamics as found in actual systems. The simulation results demonstrate the possibility of using parameter estimation for flow measurement and provide motivation for performing experiments on a physical system.

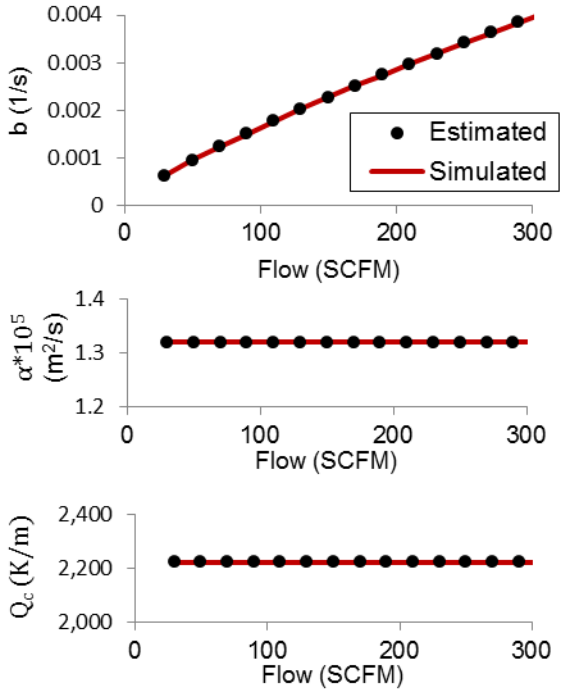

Figure 28. Simulation results showing estimated parameters including the flow term.

3.1.5 Experimental Results

An experimental setup was created to validate the simulation results and evaluate the effect of noise on the estimation of flow. Figure 29 shows the experimental setup. One heater band was wrapped around the pipe with the test thermocouples spaced along the pipe. Although only two thermocouples are necessary for the flow estimation,

additional thermocouples were attached to improve overall robustness by examining the error in estimation for all four thermocouples. An insulation blanket, seen in the figure behind the thermocouples and heaters, was wrapped around the pipe surface. The test thermocouples and the heater were connected to a control box. The control box was equipped with relays, a microcontroller, and a fuse. All of the thermocouples were attached to a data acquisition module for data collection. As the correlation between flow rate and heat transfer is independent of pressure, a fan with a variable transformer was employed to generate flow through the pipe with a handheld anemometer at the end of the pipe to measure actual flow.

Figure 29. Experimental setup for parameter estimation.

A total of 18 tests were run. Nine of these tests were performed on a 2-inch Schedule 40 black iron pipe, and the other 9 tests were performed on a 3-inch line. The outer diameter of the pipe was measured at the test section, and the inner diameter of the pipe was measured at both open ends of the pipe. During each test, the fan was operated for a minimum of 20 minutes prior to activating the heater in order for the temperature of the pipe to equilibrate with the increased air temperature due to the fan. The heater

maintained constant power for six minutes during data collection. Following the six minutes, the heater was turned off and data were collected from the thermocouples as the pipe cooled for future data analysis. During processing, a simple low pass filter was applied to the temperature profile and its time derivative to improve the estimation on the actual peak in the derivative profile.

Figure 30 shows the parameter estimation results from four tests. The flow rates ranged between 60 SCFM and 270 SCFM on the 3-inch Schedule 40 black steel pipe. The figure shows the final estimated values from the parameter estimation algorithm along with the initial starting guesses for each of these tests. For the parameter b, the initial parameter values are visibly not accurate enough to predict potential flow rates, as there is no correlation to flow from the initial estimates. However, the values of initial estimates for all three parameters are sensible and close to the final parameter values.

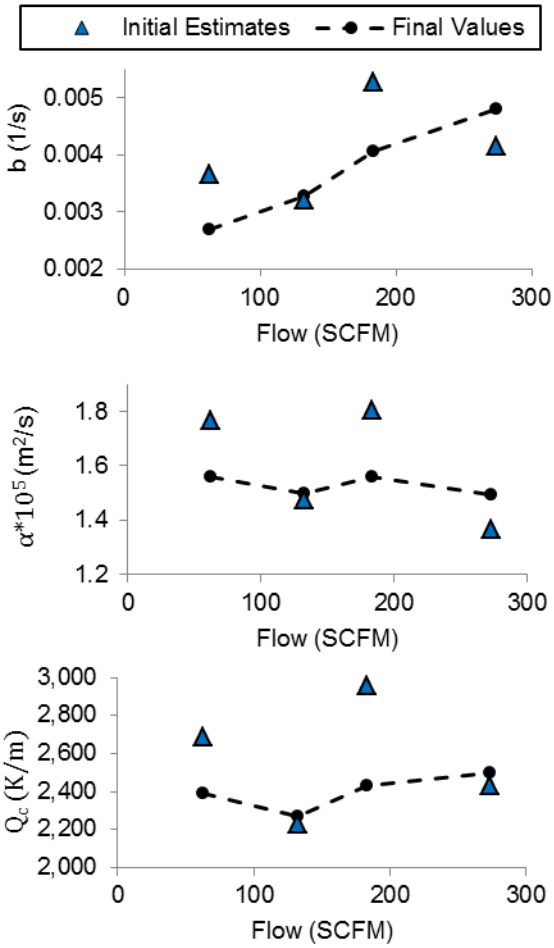

Figure 30. The initial estimated parameters and final estimated parameters for the three experiments. This figure shows the correlation for Qc, b, and alpha for the four experiments.

The temperature profiles and their derivatives for all four thermocouples for one of the experimental runs is shown in Figure 31. Also shown in the figure is the estimated analytical profile for each of the temperature profiles after performing the parameter estimation for the experimental test. The estimated profile is closely correlated to the experimental profile for all thermocouples as seen in the figure.

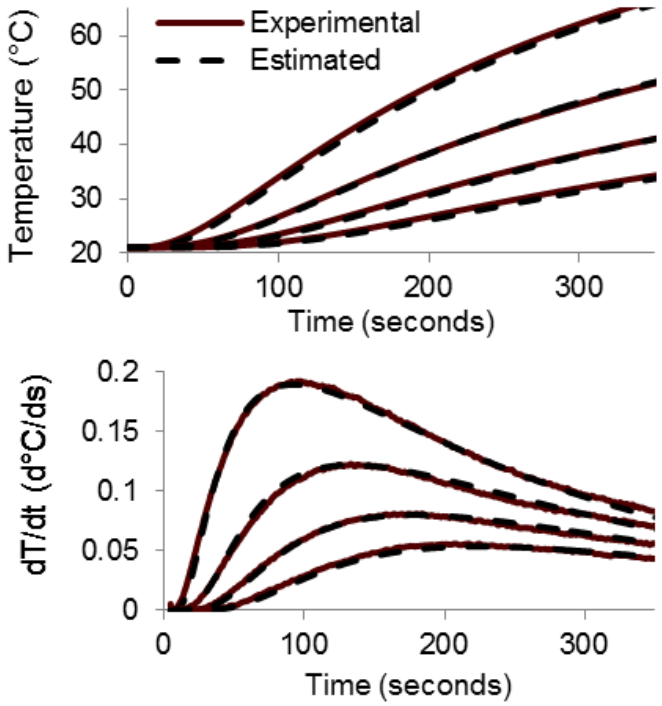

Figure 31. The final estimated temperature profile and derivative curve for all thermocouples on one experiment.

The final experimental results for the 2-inch line are shown in Figure 32. The results from the 3-inch line are shown in Figure 33. In both sets of tests, each experimental predicts a relatively constant value for x0, α, and Qc. Additionally, there is a clear increasing trend of the parameter, b, with increasing flow rate. The final estimated values for the thermal diffusivity fall within the range of expected values for black steel pipe, and the other values are on the same order of magnitude as the values used in simulation. Note that the predicted values, although similar to the values seen in Figure 28, are not expected to match each other. The values taken for simulation were reasonable values that could be expected for black steel pipe, but they do not attempt to estimate the actual values of the thermal diffusivity, heat input into the pipe, or heat

capacity. Overall, the preliminary experimental results confirm results from simulation and show promise in the future ability to separate flow measurement from other pipe material parameters.

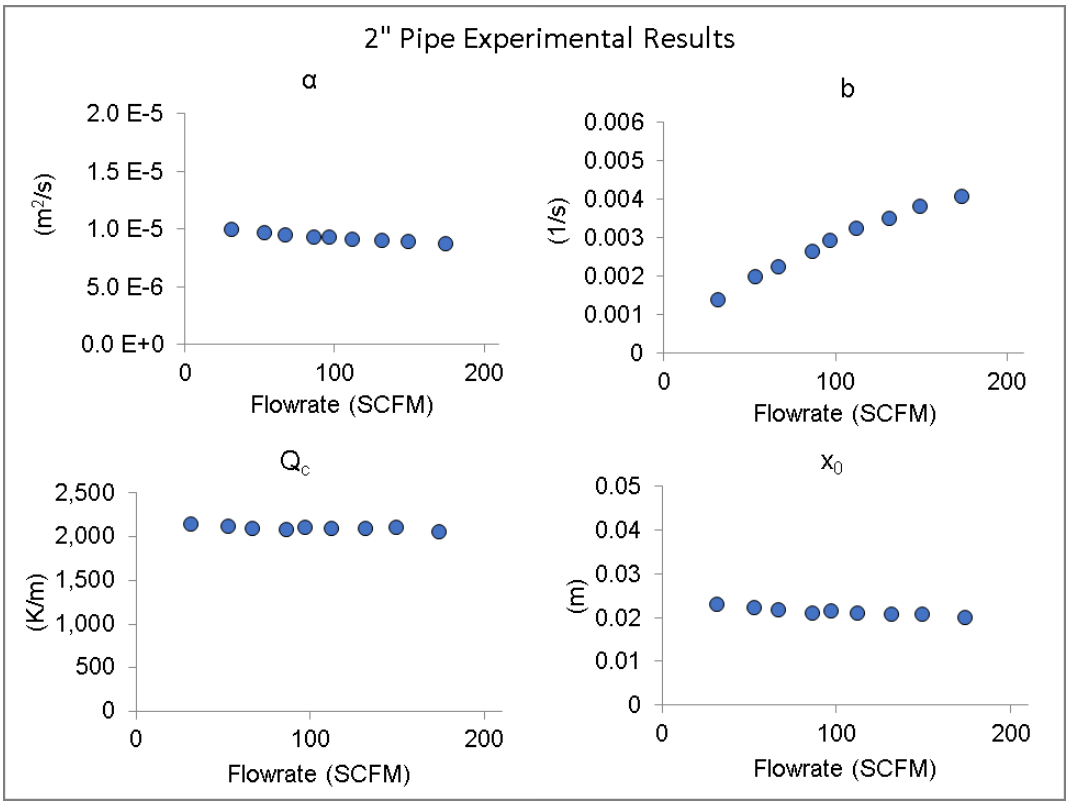

Figure 32. Experimental results for the 2-inch pipe.

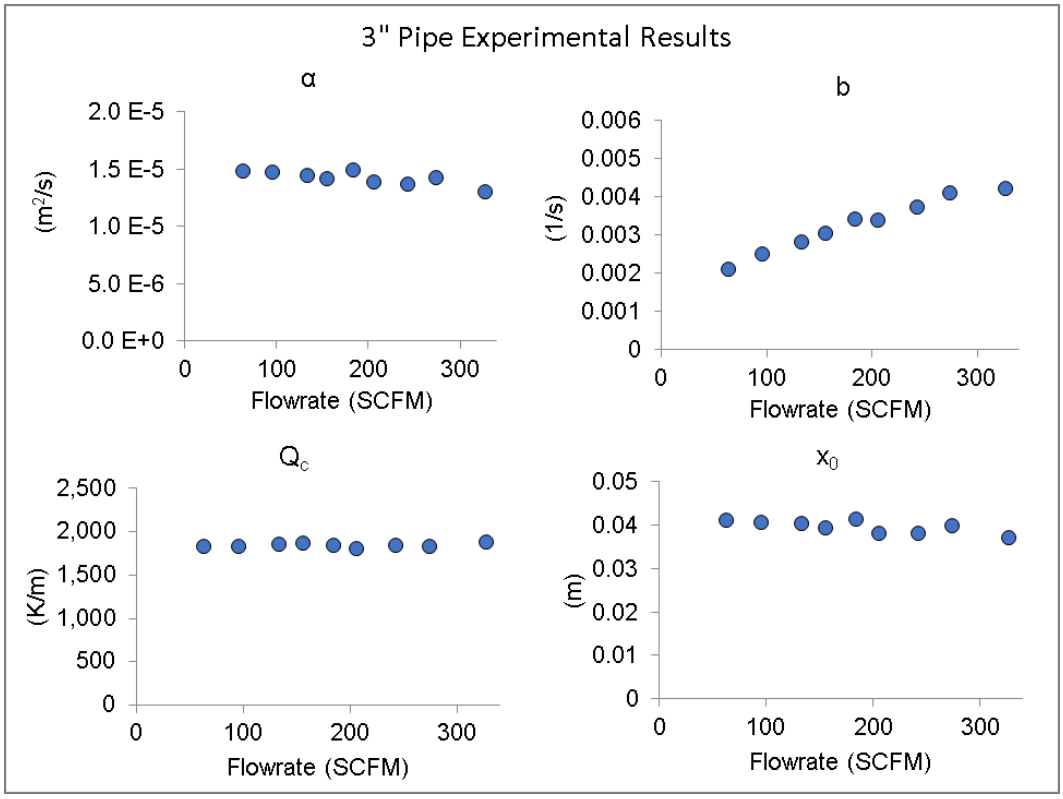

Figure 33. Experimental Results for the 3-inch pipe.

While the values for the thermal diffusivity and x-offset are relatively constant over the tests for each pipe, there is a slight decreasing trend with flow for both the 2-inch and 3-inch line. As both of these parameters are only affected by material parameters or dimensions of the pipe, there should be no discernible trend correlated to flow. These correlations point to either unaccounted phenomena in the model that affect the estimation or insufficient estimation. This topic is discussed in further detail below with an exhaustive parameter estimation algorithm and the discussion of the test case.

3.1.5.1 Exhaustive Search Comparison

To gauge the effectiveness of the parameter estimation using the gradient descent method, an exhaustive search over the parameter space was explored. The length of

computation using an exhaustive search method is prohibitive for a real-time field test, but it is useful as a one-time comparison to the estimated parameters using the gradient descent method.

Due to the computational cost in searching the entire parameter space, an exhaustive search was modified so that the entire parameter space was not searched in detail. This method was used because, when looking at the error profile of a sparse grid of parameter points, the error profile was smooth with no likely candidates of local minima. The minimum error obtained will likewise be the global minimum.

The stacked searching method is illustrated in Figure 34. An initial sparse grid of parameter values over the entire parameter space is searched exhaustively. The minimum error of this grid is identified as the center of the new searching grid, from which a new finer search grid is developed that extends slightly beyond the neighboring parameter values in the original grid. The process is iterated of a user-identified number of iterations.

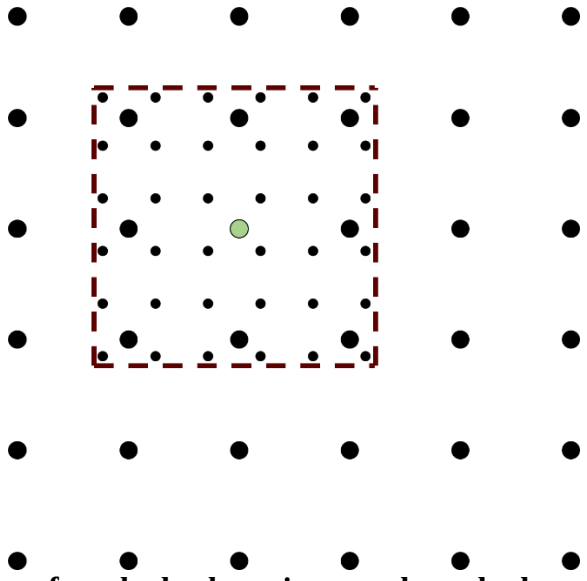

Figure 34. Illustration of stacked exhaustive search method.

Table 7 shows the initial range of parameters searched in the sparse grid. Included in the table are the observed final values from the gradient descent method to illustrate the wide parameter space that was searched. This value incorporates the entire space of feasible parameter values due to physical constraints.

Table 7. Searched Parameter Space in the Exhaustive Search Method

Parameter	Units	Min Value	Max Value	Range of Observed Values in Gradient Descent
Qc	K/m	100	5000	2,010 to 2,930
b	1/s	1e-5	0.02	0.001 to 0.005
α	m^2/s	1e-7	1e-4	9.0e-6 to 1.8e-5
x0	m	0	0.5	0.03 to 0.06

Figure 35 compares esimated parameters from the gradient descent and exhaustive search parameter estimation method. This figure shows the tests for the 3-

inch pipe. The tests for the 2-inch pipe showed similar results. There is close agreement between the two methods, indicating that the gradient descent method of estimation converges on the minimum error in the parameter space.

The slight correlation of α and x0 with flow appear in the exhuastive search method. This inidicates that the same correlation observed in the gradient descent method does not result from limiations in the parameter estimation method. Instead, the correlation likely results from limitaitons in the chosen model. Likely, unmodelled or unconsidered dynamics affect the experimental profile that the analytical model doesn't capture.

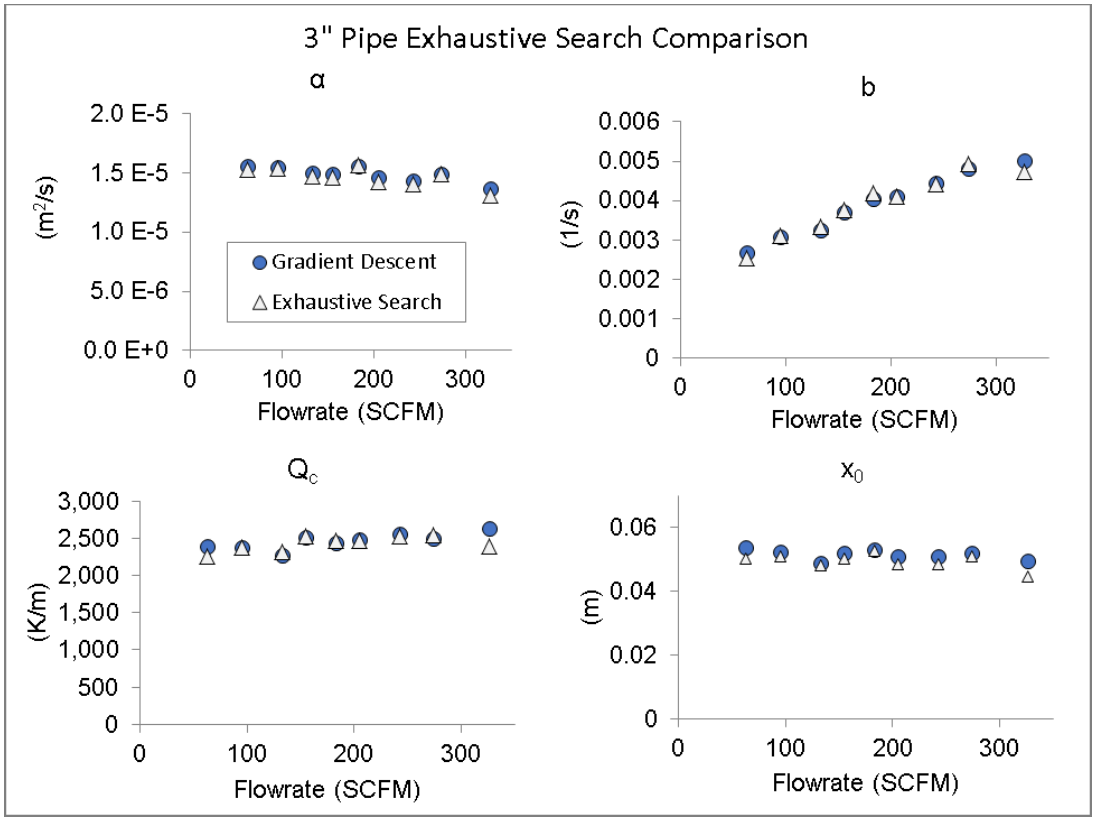

Figure 35. Comparison of parameters values between the exhaustive search and gradient descent parameter estimation methods.

Figure 36 shows the average profile error between the estimation and measured profiles for each of the tests. The exhaustive search shows a slightly improvement of error over the gradient descent method for all tests.

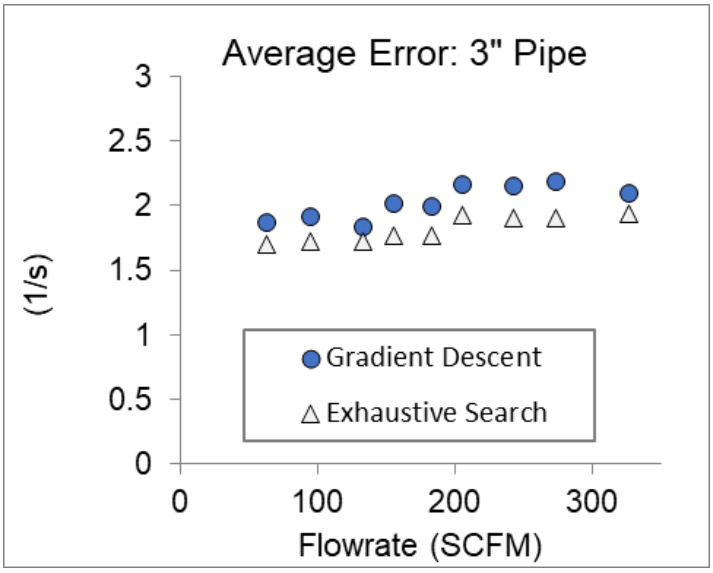

Figure 36. Comparison of error between the exhaustive search and gradient descent parameter estimation methods.

3.1.6 First Field Test

Similar to the steady state method presented in the previous chapter, a field test was conducted for the dynamic parameter estimation. The same compressed air system from the previous chapter was used to perform the test. The heater and thermocouples were attached to the 2" copper header line, and the system is serviced by a 15 hp reciprocating compressor. The flowmeter was installed in the same section of line, with a straight run extending 6 feet prior to and 30 feet after the test section. Figure 37 shows the installed flowmeter prior to testing.

Figure 37. Picture of setup for the first field test.

In performing the test, the pipe was heated continuously at a constant rate for 5 minutes. Temperature data were recorded as the pipe cooled, but only the time with heating is used in the analysis. Figure 38 shows the temperature response of the 4 thermocouples installed on the pipe. The temperature trends are similar to those seen in Figure 27 for previous experimental tests.

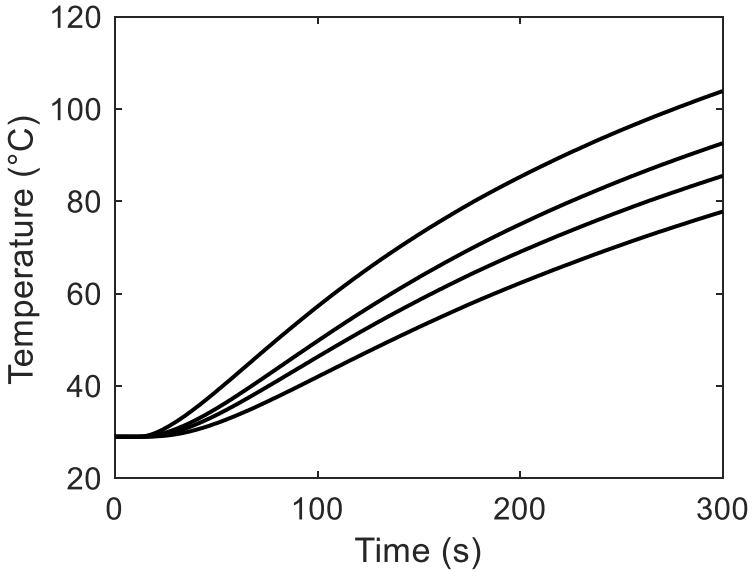

Figure 38. Temperature profile of the first field test.

Since the pipe parameters were slightly different for the copper pipe compared to the black iron pipe used in other experimental tests, these differences were corrected for by using the definition of the lumped parameter, b. The correct parameter b is given by Equation (50).

(50) $$b = \frac{(\rho * c * D_i)_{copper}}{(\rho * c * D_i)_{black\ iron}} b_{measured}$$

The same values used in Table 3 in the previous chapter were used in this analysis. Due to the material properties, the scale factor for correction of the parameter, *b*, was calculated as 0.97. This results in only nominal correction of the *b* parameter. The other material parameters, Qc and α, were left uncorrected due to uncertainty in the thermal conductivity. While the thermal conductivity of the copper can be readily

obtained, the thermal conductivity of black iron varies over a considerable range, introducing significant uncertainty on the appropriate scale factor. As knowledge of these parameters is not necessary for flow estimation, this correction was omitted for the field test.

Figure 39 shows the estimated and measured temperature values for each of the four thermocouples from the test case. The results are notably worse than those obtained from the experimental tests on black iron. In the gradient descent method, a limitation is placed on the maximum axial offset, x0, that can be achieved through the estimation. This ensures consistency with the experimental tests, where the value of the offset accounts for the difference between the physical edge of the heater and the location at which heat is effectively added. Left unrestrained, the gradient descent algorithm provides a much better fit, but the offset value exceeds a reasonable value. This likely results from unmodeled heater dynamics, which is discussed in detail below.

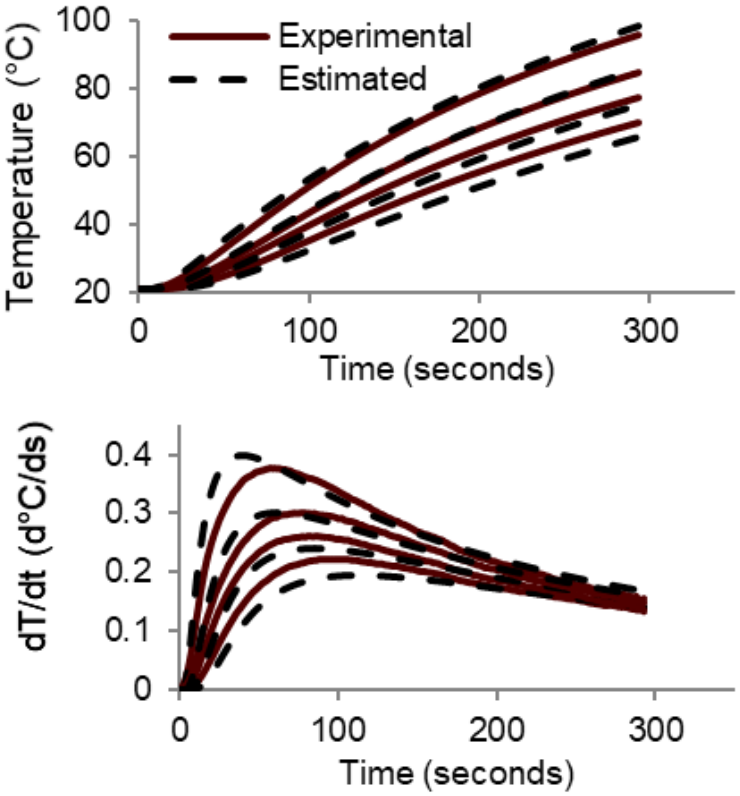

Figure 39. Comparison of measured and estimated temperature profiles for the first field test.

The estimated parameters based on the parameter estimation are displayed in Figure 40. They are also shown in Table 4. As seen from the results, the estimated flow rate is lower than the lowest calibration flow rate, similar to what was seen during the steady state field test. While the thermal diffusivity, α, is lower than the actual value for copper, the estimated value is higher than the diffusivity for the tests on black iron. This higher diffusivity value is expected due to the higher thermal conductivity of copper. Similarly, the high conductivity of copper results in a lower value of the parameter, Qc, compared to the tests on black iron.

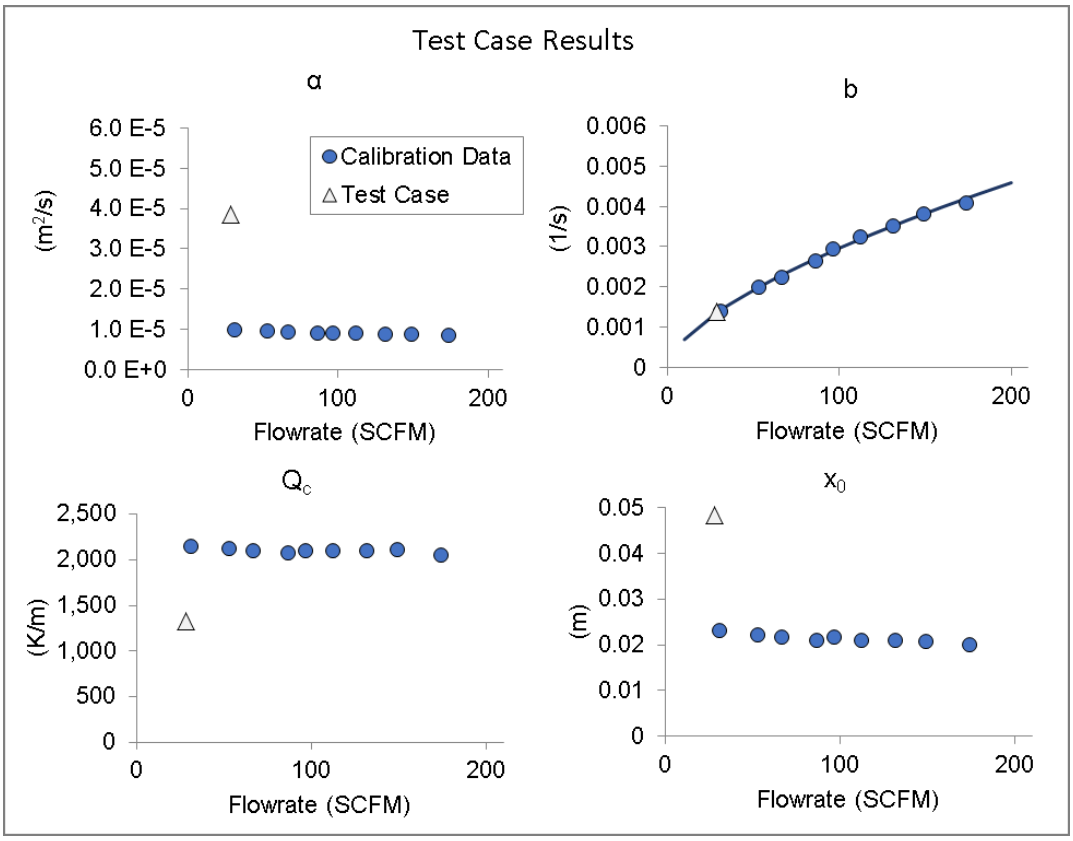

Figure 40. Parameter estimation comparison for the first field test.

Table 8. Parameter Estimation Test Case Results

Parameter	Symbol	Units	Value
Measured Lumped Convection Term	$b_{measured}$	s-1	0.00139
Scaled Lumped Convection Term	b	s-1	0.00135
Measured Thermal Diffusivity	α	m^2/s	3.85e-5
Measured Heat Flux Boundary Condition	Qc	K/m	1,330
Measured Axial Offset	x0	m	0.00483
Estimated Flow Rate	-	SCFM	28.6

The flow rate was estimated from the parameter, *b*, based on a power law relationship for the initial experimental tests for black iron. While the overall estimated value for the flow rate is low, this value of 28.6 SCFM is slightly higher than the

estimated flow rates for both the drop-down and duty factor methods of flow estimation. The estimated flow values using these two methods are displayed in Table 5 and Table 6 for reference.

This discrepancy, along with the error in the estimated value of α, likely come from unmodeled dynamics in the test setup. Earlier in the experimental tests for black iron, there was a slightly trend of estimated values of α relative to the flow rate. The thermal diffusivity, α, is a material parameter and independent of flow. As all tests were performed on the same pipe, this trend results from a discrepancy between the modelled and actual temperature profile.

The most likely unmodeled dynamic is that of the heater. As electricity is applied to the heater, the temperature of the heater rises, causes a heat addition into the pipe. The model assumes a uniform heat addition at the 'edge' of the pipe. The physical heater has finite width, spreading the heat over an area on the exterior surface of the pipe. Additionally, although uniform heat is assumed, the heater's thermal inertia will result in delayed and non-uniform heat into the pipe. The dynamics of the heater potentially have a smaller effect when the thermal conductivity of the pipe material is lower, such as was seen in the black iron pipe used for calibration. With the much higher thermal conductivity of the copper pipe, the thermal dynamics of the heater are potentially exacerbated.

There are several future approaches that can be used in future work to address this issue. For example, a first approach is to use the thermocouple closest to the heater as the input to the model. The other thermocouples on the pipe record the response to

pipe, where the model input is the temperate at the first thermocouple. An analytical solution to this model would likely be infeasible, as the input to the model would be complex, and would therefore require estimation of a state space model. A model suitable for estimation of temperature was presented earlier in Chapter 2. Model reduction could potentially be achieved in a way to still estimate the thermal diffusivity and the lumped flow term. A second approach would be to incorporate the dynamics of the heater into the model. As the heater's dynamics will be consistent across different pipe materials, the effect from the heater will be consistent across different tests.

3.1.7 Second Field Test

A second field tests was performed in another industrial environment. The main header lines were only 1-inch black iron pipes. The available heaters for this field test could only tightly wrap around lines at least 2 inches in diameter. The heaters were secured as tightly as possible, but only about half of the circumference was in physical contact with the heater. Two tests were run; one test had near zero flow and the flow rate of the other tests was around 30 SCFM. The load/unload cycles of the reciprocating compressor along with nameplate information was used to estimate these flow rates.

Due to the different pipe sizes from earlier experimental tests and limited physical contact with the pipe and heater, this section will present two illustrative considerations for future reference and study. The pipe consisted of the same material as previous experimental tests on the 2-inch and 3-inch pipes.

The flowmeter was located on a line close to the compressor immediately after a receiver tank. During the first test with very low flow, the compressor remained off for

the duration of the test and the resulting temperature profile looks very similar to Figure 38. On the second test, with the higher flow rate, the compressor cycled on and off multiple times during the experiment. Figure 41 shows the temperature profiles of the four thermocouples over the duration of the test. The figure also highlights times when the compressor cycled on and off. The temperature profiles are clearly dependent on the compressor cycling.

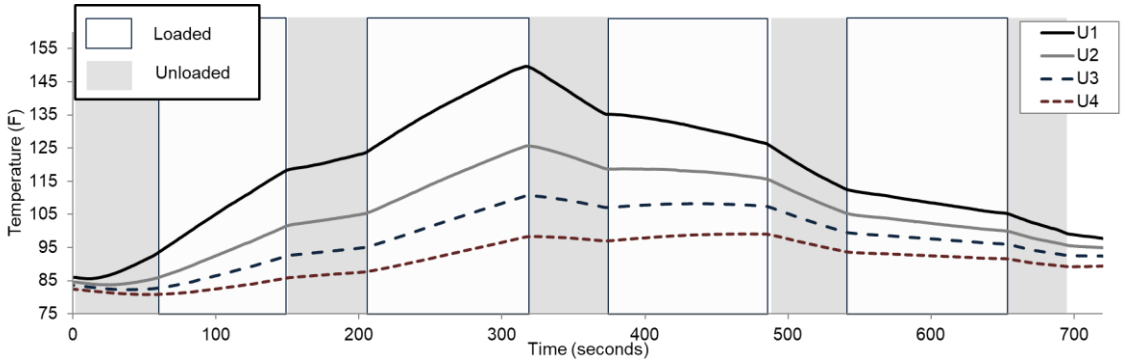

Figure 41. Temperature profile of second field test with compressor cycling.

The first observation from this field test is that the position of the flowmeter is a consideration due to compressor cycling. While the exact cause of the profile shape is uncertain, the likely reason comes from heated air at the outlet of the compressor. The air experiences a meaningful temperature increase through the compressor, and this higher temperature will inhibit heat transfer at the flowmeter. This theory is consistent with the observed profile, where the temperature rises faster in the heated section of the test. The inhibited heat transfer will result in lower heat losses, causing a greater rise in temperature. Likewise, the pipe cools more slowly when the compressor is cycled on.

The second observation from the field test comes from observing the estimated parameters from the gradient descent method. Figure 42 shows the estimated parameters for both field tests along with all the experimental tests for both the 2-inch and 3-inch lines. The tests on the different lines are not differentiated as the field test was performed on a 1-inch line, and this figure illustrates the overall results of the parameter estimation. As seen in the figure, the parameter estimates are overall consistent with the results seen in previous tests. While the accuracy values of these field tests are unknown, the limited contact of the heater with the pipe appears to have a limited effect on the parameter estimation.

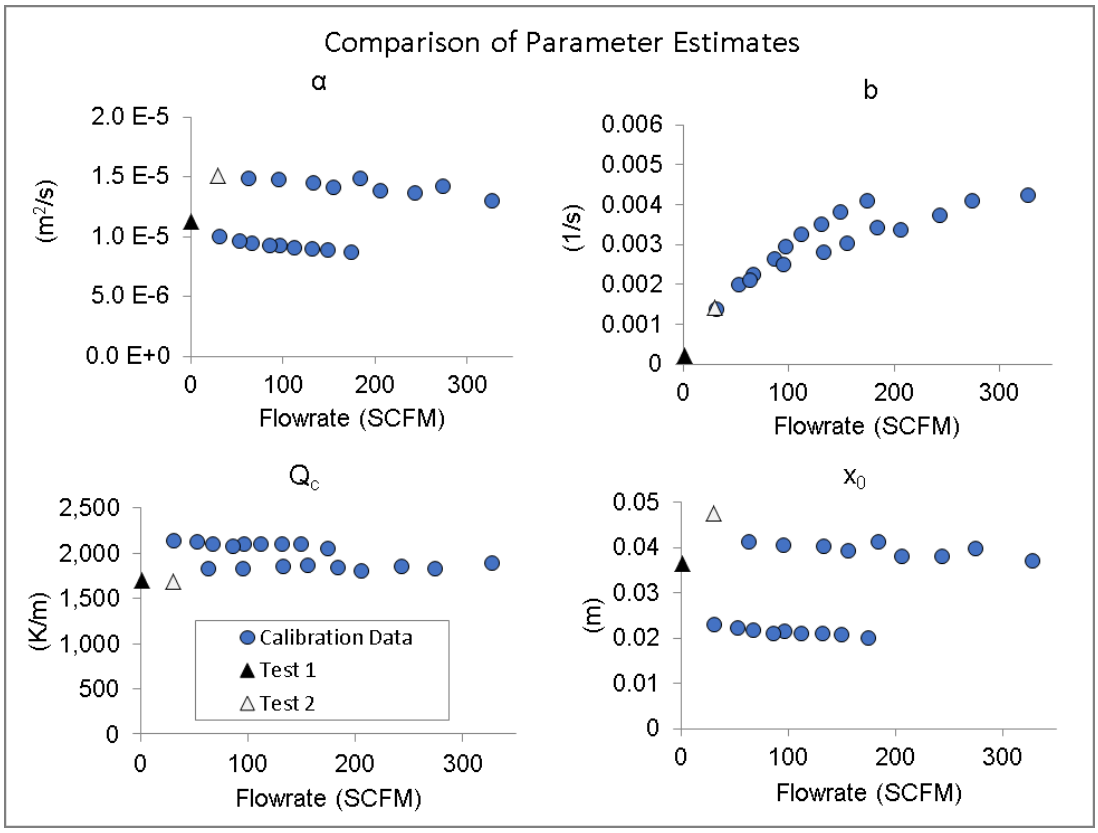

Figure 42. Comparison of estimation parameter from second field test.

3.1.8 Future Directions of Research

There are several open challenges and directions for future research on this topic. One of the primary topics to address is improvement of the estimation. By removing the observed correlation of α and x0 with flow, the estimated parameter b will be updated and reflect a more accurate estimate of the flow. Another possible direction of research is the potential measurement of pipe thickness using an approach similar to the approach used in this chapter. Thickness estimation is an open field of research for corrosion detection and health monitoring of pipes. Several possible directions of research are

presented in this section. A preliminary analysis has been performed for some of these methods, and this analysis is included in the section.

3.1.8.1 Use of First Thermocouple as System Input

This method, which was mentioned briefly previously in the chapter, alters the boundary condition of the model. The constant heat input assumption is dropped, and the pipe is given defined temperature profile as the boundary condition. The temperature profile is the recorded temperature profile of the thermocouple closest to the heater. As there are no other heat sources downstream of this thermocouple and the pipe starts at uniform temperature, the resulting temperature profiles of the other thermocouples will be exclusively dependent on the temperature at the boundary. By using the approach, the dynamics and contact resistance of the heater are removed. This approach will likely improve the estimation of the parameter estimation as the dynamics are only dependent on the pipe material, dimensions, and flow inside the pipe.

There are two challenges with this approach. As mentioned previously in the chapter, the temperature profile has sufficient complexity to limit the ability to use an analytical solution in the estimation. A state space approach with appropriate model reduction will likely enable estimation from the more complex model input. The second limitation is the reduced number of parameters to estimate. The x0 offset is no longer a parameter to estimate, and the equivalent Qc parameter can no longer be estimated. For the work presented in this chapter, there are no immediate drawbacks. However, physical parameters, such as the thermal conductivity, total heat capacity, and the convection coefficient, can currently be estimated with knowledge of the power input to

the heater. By removing the parameter, Qc, only the thermal diffusivity and lumped flow term can be measured without the ability to estimate the physical parameters.

3.1.8.2 Dynamic Input Signals for Improved Estimation

Some methods in the literature suggest that periodic or other non-steady-state signals may be used to estimate the heat flux in the determination of heat capacity or thermal conductivity. This section explores the addition of a sinusoidal input signal to determine the parameters. The analytical model for a sinusoidal input is derived in this section, as well as the practicality of using this to determine the parameters. Potential improved sensitivity is achieved by using a dynamic signal instead of a constant input.

The initial model is derived similar to the constant heat input, with the governing equation given by Equation (22). The first boundary condition is formed similar to the earlier model by recognizing that as $x \rightarrow \infty$, the temperature must remain bounded. The initial condition is a uniform temperature throughout the pipe. The second boundary condition is a sinusoidal input, given by Equation (51).

(51) $$\frac{\partial u(0,t)}{\partial x} = -\frac{q_{in}}{k}\sin(\omega t) = -Q_c \sin(\omega t)$$

The solution is given by Equation (27), and the equation simplifies to Equation (52).

(52) $$u(x,t) = Q_c\sqrt{\frac{\alpha}{\pi}}\int_0^t \sin(\omega\tau)\frac{e^{-b(t-\tau)}}{\sqrt{(t-\tau)}}e^{\frac{-x^2}{4\alpha(t-\tau)}}d\tau$$

As before, a substitution variable, $y = \sqrt{t - \tau}$, is introduced, and the resulting integral is given by Equation (53).

(53) $$u(x,t) = -2Q_c\sqrt{\frac{\alpha}{\pi}}\int_{\sqrt{t}}^{0}\sin(\omega t - \omega y^2)\, e^{-by^2 - \frac{x^2}{4\alpha y^2}}dy$$

The integral is solvable. The result of the integral, as computed by an online integral solver, is shown in Equation (54) [65].

(54) $$\int \sin\left(w\left(t - x^2\right)\right) e^{bx^2 - \frac{c}{x^2}}dx =$$

$$\frac{1}{8}i\sqrt{\pi}e^{-itw}\left[\begin{array}{c} -\dfrac{e^{2itw-2\sqrt{c}\sqrt{-b+iw}}\,\mathrm{erf}\left(-\dfrac{\sqrt{c}}{x}+x\sqrt{-b+iw}\right)}{\sqrt{-b+iw}} \\ -\dfrac{e^{2\sqrt{c}\sqrt{-b+iw}+2itw}\,\mathrm{erf}\left(\dfrac{\sqrt{c}}{x}+x\sqrt{-b+iw}\right)}{\sqrt{-b+iw}} \\ +\dfrac{e^{-2\sqrt{c}\sqrt{-b-iw}}\,\mathrm{erf}\left(-\dfrac{\sqrt{c}}{x}+x\sqrt{-b-iw}\right)}{\sqrt{-b-iw}} \\ +\dfrac{e^{2\sqrt{c}\sqrt{-b-iw}}\,\mathrm{erf}\left(\dfrac{\sqrt{c}}{x}+x\sqrt{-b-iw}\right)}{\sqrt{-b-iw}} \\ -\dfrac{e^{2itw-2\sqrt{c}\sqrt{-b+iw}}}{\sqrt{-b+iw}}+\dfrac{e^{2itw+2\sqrt{c}\sqrt{-b+iw}}}{\sqrt{-b+iw}} \\ +\dfrac{e^{-2\sqrt{c}\sqrt{-b-iw}}}{\sqrt{-b-iw}}-\dfrac{e^{2\sqrt{c}\sqrt{-b-iw}}}{\sqrt{-b-iw}} \end{array}\right]$$

While no attempts are made in this preliminary analysis to reduce this integral to a manageable form, the limits of integration can be applied. In solving the solution in Equation (29), many terms were grouped together and the solution simplified considerably. This likely will also occur for the sinusoidal heat input, especially considering the output will include a sinusoidal term. Preliminary further analysis suggests the solution is an infinite solution, but later terms can be dropped with negligible error. Equation (55) demonstrates that a solution exists, and future research can reduce the solution and apply it to the parameter estimation.

If future work includes parameter estimation in the state space domain, sinusoidal or other periodic signal can be utilized over a constant heat input. This would reduce complexity in solving the analytical solution, at the expense of more difficult parameter estimation. Use of periodic inputs would be appropriate when using the first thermocouple measurement as the model boundary condition.

3.1.8.3 Asymmetric Temperature Profiles for Improved Flow or Thickness Estimation

When performing tests on the method presented in Chapter 2, thermocouples were placed upstream and downstream of the heaters for additional measurements. While these measurements were not used in the estimation of flow, patterns observed from these temperature profiles can potentially improve flow estimation or enable measurement of pipe thickness.

When tests were performed with no flow inside the pipe, the upstream and downstream temperature profiles were symmetric. However, even with low flow rates,

there was a notable difference in the upstream and downstream temperature profiles. As the air passes through the pipe, the air heats up, inhibiting heat transfer with the smaller difference between the pipe and air temperatures. While only a small increase in air temperature, it produced measurable differences in the temperature response of the pipe. Preliminary analysis of the upstream of the analytical and numerical models indicate that this temperature difference is a function of the airflow inside the pipe and the thickness of the pipe. Currently, which parameter produces the dominant effect on the temperature difference is unknown. But both parameters are of interest, and using this temperature difference can result in either improved flow estimation or measurement of pipe thickness.

3.1.8.4 Other Smaller Improvements on Analysis to Improve Flow Estimation

- Instead of a wrap-around heat, a heat impulse could be sent to the pipe, and the temperature response circumferentially around the pipe could be measurement. With the assumption of a sufficiently thin pipe, the problem could be framed as a flat plate with appropriate boundary conditions at the interface. There are similar examples using thin circular rings in partial differential equations references [66].
- Current models presented in this chapter only consider the heating portion of the temperature response. Including additional analysis on the cooling portion of the temperature profile could potentially improve flow estimation.
- In the method presented in Chapter 2, there was an expected temperature profile axially along the pipe at the end of the constant temperature section. This

temperature profile is likely to be dependent on the thermal diffusivity and flow term.

- When heat is added to the pipe, there is a time delay between the initial heater input and the observed temperature response. There is a distinct upturn in temperature after this delay period. Analysis in both simulation and experiment show that the time of this term is independent of the flow term. As such, it can potentially be used to determine the thermal diffusivity in the estimation.
- The steady state value when applying constant heat is independent on the volumetric heat capacity (ρ*c) of the pipe. The volumetric heat capacity only affects the rate at which the steady state value is reached. As such the steady state value can be approximated and used in the estimate of α and b.

3.2 Summary

Advancements on the non-intrusive flowmeter presented in Chapter 2 enabled on-line parameter estimation for improved flow measurement. An analytical model was derived, and a gradient descent method was used to derive multiple parameters. Simulations and experimental tests were performed to validate the method, and the results showed a clear relationship between the estimated flow term and the measured flow. A field test demonstrated the viability of the method on a different pipe material, but it also revealed the potential advantage in incorporate the heater dynamics into the model. A second field tests revealed a potential limitation in using the flowmeter close to the operating compressor. Future work includes various methods to improve the flow

meter either through improved parameter estimation or incorporating observed phenomena during experimental tests.

CHAPTER 4

AUTONOMOUS DATA ANALYSIS*

One challenge to energy efficiency in industrial facilities is the cost to identify and analyze potential retrofit opportunities. For both the external energy auditor and internal energy manager, the cost to evaluate potential energy-saving opportunities reduces the economic viability of such opportunities. In general, the increasing amount of data that are available in many fields requires additional and automated data analysis techniques. This chapter presents s case study on automated data analysis can be used to identifying and evaluate energy-saving opportunities in industrial facilities.

Specifically, this chapter presents a demonstration of a method to automatically assess buildings for energy retrofit opportunities, with a focus on the data collection and analysis of lighting systems. The generalized procedure for analyzing these systems is presented. The method of data collection and analysis to identify opportunities and quantify potential savings is then detailed. The lighting analysis is improved through spatial information during navigation and exploration. The potential for additional analysis beyond traditional lighting assessments through simulation is then discussed. The navigation strategies and exploration algorithms, which are interconnected with the

* Parts of the content in this chapter are reprinted with permission from "Autonomous lighting assessments in buildings; part 2 – light identification and energy analysis" by Trevor Terrill, Christopher Bay, and Bryan Rasmussen PhD, 2017, *Advances in Building Energy Research*, 11:2, pg 1-18.

lighting analysis, are not the focus of this work and are discussed in other published papers [67-69].

4.1 Case Study: Autonomous Robotic Assessments of Energy (AuRAE)

Reducing building energy consumption in both commercial and industrial facilities has been a recent area of focus due to their large percentage of total U.S. energy consumption [70]. Codes and standards are in place to ensure energy efficient construction of new buildings. With long operating lifespans of buildings, a large focus is also placed on improving energy use in existing buildings. However, a limited number of energy assessments have been performed on industrial and commercial buildings, primarily due to the expense of the assessment. With additional data available through advanced sensing, the labor expense to analyze these data climb. Autonomous robotic assessments of energy (AuRAE) are proposed as a method to reduce the cost of performing retro-commissioning on buildings, thereby expanding implementation and realizing greater energy reductions in buildings. More extensive and faster data collection over a human assessor is possible with the AuRAE, thereby enabling additional extended opportunities for analysis and visualization. This automated analysis is applicable to many energy systems in industrial buildings, but the following sections concentrate on the viability of automated analysis in building lighting systems.

4.1.1 Energy Use in Buildings

In addressing the task of reducing carbon emissions, buildings present a large opportunity for savings. Buildings constitute approximately 40% of all United States energy use [70]. Despite the relatively stagnant appearance, buildings are dynamic

systems that change over the course of days, seasons, and years. Additionally, efficiency of components in energy systems degrades over time while new technology offers performance improvements. New buildings are constructed with energy efficiency measures in insulation, lighting, and air conditioning systems [71]. However, significant cost savings can be realized by analyzing energy efficiency opportunities in existing buildings and upgrading building systems. This process, termed retro-conditioning, has significant documented energy and cost savings for participants [72].

Retro-commissioning has the potential to make a significant impact on reducing energy use and carbon emissions. One of the primary challenges of retro-commissioning is the expertise required of the energy auditor [73]. The extensive required training, paired with the labor intensive process of an assessment, results in high cost of commissioning. Additionally, the quality of the auditing is highly dependent on the experience of the auditor.

In an effort to improve the implementation rate and efficacy of the audits, AuRAE have been proposed. A team of ground and aerial robots will navigate a building, performing measurements and analyzing energy efficiency opportunities. Initially, these teams of robots will augment the energy auditing process, eliminating simple, time-consuming tasks and providing additional depth of analysis not possible with a human auditor. In the future, entirely automated visits are envisioned with these robots.

The vehicles will navigate buildings and perform energy audits in reduced time with lower capital costs than current expert energy auditors. These audits, which have

shown measurable benefits in cost and energy savings, have the potential to be performed on a much larger scale through use of autonomous vehicles. They can produce significant and sustainable energy efficiency improvements, promoting a sustainable energy portfolio. These AuRAE offer several advantages over a traditional audit. They provide consistent, repeatable results, independent of the expertise of the auditor. They can take multiple measurements at a time, and they can store information at a faster rate than humans, taking necessary measurements in reduced time. Finally, the extended measurement capabilities enable increased depth of analysis. For example, spatially dependent measurements such as temperature can be taken throughout a room instead of a representative sample somewhere near the middle of the room. These data offer additional opportunities for savings, simulation, and visualization.

4.1.2 Lighting in Buildings

Energy usage in commercial and industrial buildings constitutes a dominant portion of the energy usage both inside and outside of the United States [70]. Over 30% usage in these buildings is derived from lighting and heating, ventilation, and air conditioning (HVAC) systems. Lighting in the United States accounted for roughly 275 billion kWh or 21% of electrical consumption in 2011 [72]. In global energy usage, lighting composed 19% of the total electricity consumption in 2005 [74]. Consequently, industrial lighting presents one of the greatest opportunities for energy efficiency improvements. With such widespread use, technology continually evolves to provide more energy efficient alternatives while maintaining appropriate lighting. Some of the biggest contributing factors to energy inefficiency are excess heat generation due to

inefficient types of lighting, incorrect level of lighting for a given task, ineffective use of available daylight, and lights being left on when a space is unoccupied.

New buildings are often equipped with the most cost effective lighting technology that is available. However, with rapidly changing technology in lighting and lighting controls, cost effective lighting retrofits become available multiple times over the operating lifetime of the building. Although there is an upfront capital cost with installation of the various lighting upgrades, the reduced operating cost from the lower energy consumption generates savings for the building operator.

The rapid evolution of lighting technology causes energy auditors to commonly look at lighting retrofits and upgrades first among energy assessment tasks. A history of lighting upgrades can be evaluated in industrial facilities through the Industrial Assessment Center (IAC) database [19]. The IAC program is targeted at performing energy assessments at manufacturing facilities across the United States. The program, run by the U.S. Department of Energy, has operated for over 40 years and recorded over 16,000 assessments. Anonymous assessment data are recorded from each visit and made publicly available for analysis. This database includes information on all recommendations made since the beginning of the program. Of the over 120,000 recommendations, over 20% of these recommendations pertain to lighting. This extensive use of building energy for lighting, coupled with consistent analysis of lighting by human assessors, makes lighting a logical first step in developing AuRAE.

An overview of common lighting recommendations from the IAC database is given in Table 9. These statistics include over 25,000 recommendations for lighting.

Many different types of recommendations are available to decrease energy consumption in lighting, with lighting upgrades being the most common recommendation.

Table 9. Statistics on Common Lighting Recommendations

Potential Measure	Percent of Visits Recommended	Percent of Clients Implementing	Average Cost Savings (US $)	Average Payback (years)	Possible with lighting type and location	Possible with additional lighting simulation	Possible with additional occupancy information
Lighting Upgrade	93%	54%	6,269	1.4	✓		
Occupancy Sensors	23%	35%	2,365	0.8			✓
Reduce Light Level	16%	49%	4,821	0.2		✓	
Turn Off Lights	6%	65%	3,868	0.1			✓
Install Photocells	5%	41%	3,664	0.5	✓		
Use Available Daylight	3%	43%	3,829	0.4		✓	
Install Skylights	2%	21%	6,243	1.5	✓		
Area Lighting Switches	1%	42%	5,939	0.8			✓
Lower light fixtures	1%	41%	17,564	0.7		✓	

The main objective of the lighting audit is to calculate current energy costs and to give recommendations for energy and cost savings. Once the quantification and classification of lighting is complete, sufficient data are present to make recommendations on lighting retrofits. These recommendations include potential energy consumption and demand savings, cost savings, and implementation costs.

4.1.3 Automated Analysis Approach

In analyzing lighting, an energy auditor will discriminate between types of lighting and count the number of lights of each type of lighting. With knowledge of current lighting used and usage patterns in the facility, total energy usage and associated monetary costs can be calculated. Recommendations can be given to save costs with lighting while still meeting desired light levels, such as updating outdated lamps with newer, higher-efficiency lamps. These energy and costs savings can be calculated along with appropriate implementation costs and payback period, providing sufficient information for a building manager to make informed decisions on lighting upgrades.

The energy efficiency analysis of the robots includes identifying the possible recommendations listed in Table 9 and quantifying the possible cost and energy savings from implementing the recommendation. The goal of the proposed AuRAE will be to analyze as many of these opportunities as possible. To accomplish this, the AuRAE will perform a lighting assessment by navigating and gathering sufficient data to make specific lighting recommendations. Once information is gathered, it is used to make recommendations that estimate energy savings, cost savings, cost to implement, and payback. With the sensor package presented in this paper, coupled with the autonomous

navigation and exploration algorithms presented in other work, the number and type of lights can be determined to recommend several recommendations. Additionally, these values can be used with lighting locations to simulate lighting levels, enabling additional recommendations in Table 9. If additional occupancy information can be acquired, all of the common recommendations listed in the table are possible.

4.1.3.1 Overview of Proposed System

In development of the sensor payload, the auditing process is divided into two main areas of focus: enumeration of lighting in a facility and classification of lighting type. In accomplishing these areas of focus, the following four sensors were selected and used: a spectrometer, an optical camera, a distance sensor, and an accelerometer. The sensors collect and record data as the aerial vehicle navigates through a facility. The automated process is summarized in Figure 43. An unmanned vehicle navigates in an unknown environment, avoiding obstacles and exploring the space. As the camera mounted on the vehicle captures images, the images are analyzed to determine the relative locations of the lights in the image. Individual images are merged to a global image with the aid of a laser range finder to find the global position of the lights. These locations are fed to the navigation program as target objects. In addition to exploring the space and mapping the space, the navigation system also targets each identified light to take a spectrometer reading. The spectrometer reading is then analyzed to classify the lighting type, which is used in conjunction with the locations of the lights to simulate the lighting in the space.

The autonomous vehicle navigates and explores the buildings using a Simultaneous Location and Mapping (SLAM) algorithm and a rapidly-exploring random tree (RRT) algorithm, respectively. The full details of the exploration and navigation are found in other work.

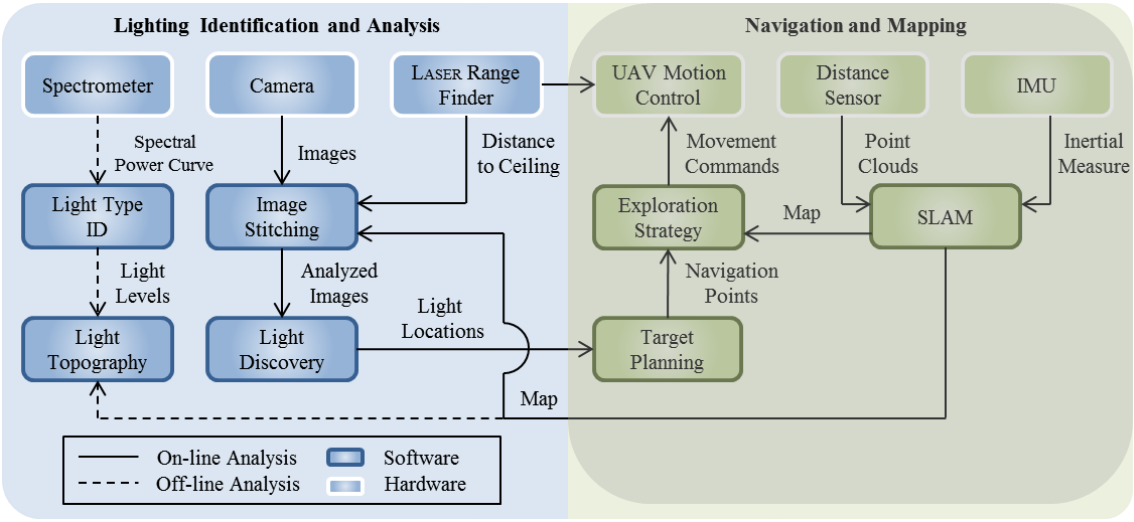

Figure 43. Overview of data collection and analysis system. The lighting identification and analysis, shown on the left half, is discussed in this paper.

Each of the following data collection and analysis steps are demonstrated experimentally in a three-room section of a commercial building. The layout of this wing of the building is shown in Figure 44. This building is equipped with 3-lamp T8 fixtures with a ceiling height of 2.7 m. In this wing of the building, the lower room is an office space with desks and computers. The middle area is a hallway, and the upper area is a laboratory space with an accompanying closet. This physical space and associated lighting will be used for demonstration of the identification and quantification of

lighting, image stitching, individual image processing, assessment recommendation calculations, and lighting simulation.

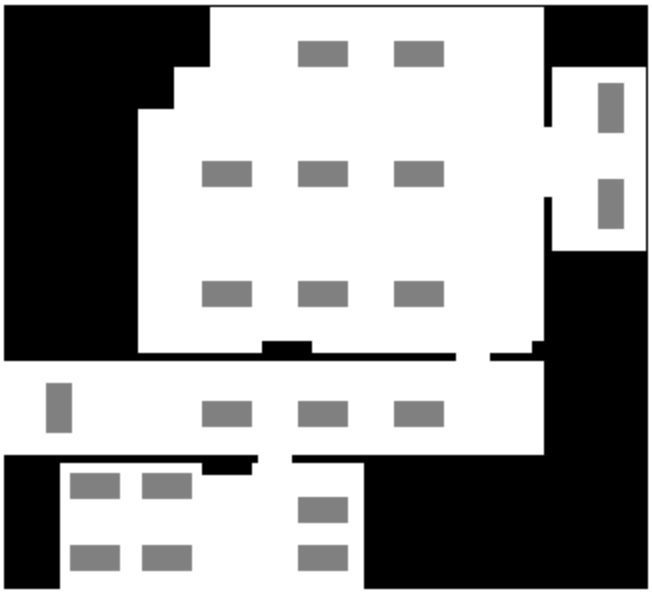

Figure 44. Example space used for simulation for the lighting simulation. This is the same layout used for the example report recommendation, navigation, exploration, and light simulation.

4.1.4 Lighting Quantification

4.1.4.1 Identification of Lights

As the vehicle navigates through the space, images are taken from an optical camera that is oriented towards the ceiling. Each individual image is analyzed to determine the presence of lighting. Camera parameters are fixed during data collection to provide consistent lighting conditions for analysis and improve robustness of lighting identification. Specifically, two different images are taken with different camera parameters. The first image is captured with a longer exposure setting to identify the

presence of lighting for use as targets in the navigation program. An example of the differences in exposure times is shown in Figure 45 on the left. With the longer exposure times, lights that are included in the image will saturate the camera and enable robust identification of lighting. An additional image with reduced exposure time is captured that is used to process each individual image by identifying the number of lamps in each light fixture and determining the shape of each lamp. An example of a ceiling image with the reduced exposure time is shown in Figure 45 on the right.

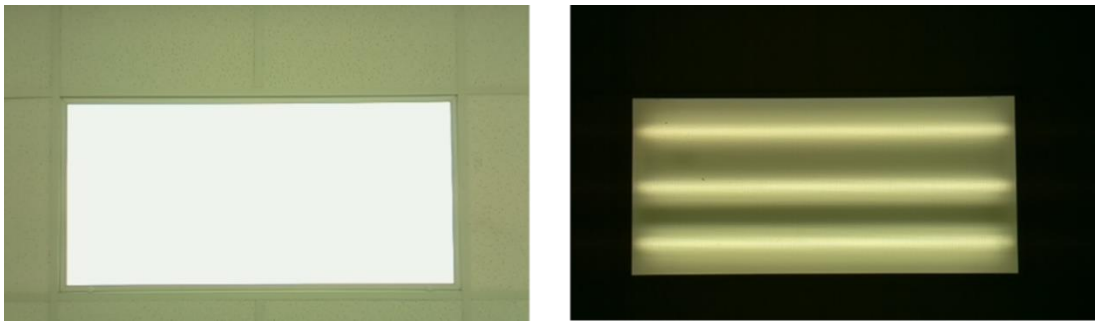

Figure 45. Example of optical images for two different camera parameters.

Images were taken in different locations, including the entire office area (lower area) of the space shown in Figure 44. In total, over 70 images were tested experimentally using the differences in exposure time to investigate the robustness of identifying lights. In these experimental tests, including the images used in the experimental test presented later in the chapter, the program successfully identified the presence of lights in every image with no false positives. Due to the high exposure time of the camera, the presence of a light consistently saturated the camera and allowed robust identification of the lighting. A threshold value of 220 out of 255 was used to

define saturation of the camera on the grayscale image. As an example of the robustness of the method in experimental tests, Figure 46 shows the outcome of two images. One had multiple lights on the edge of the perimeter and the other had no lights in the image. The stark contrast in brightness allows the successful differentiation between areas with and without lighting.

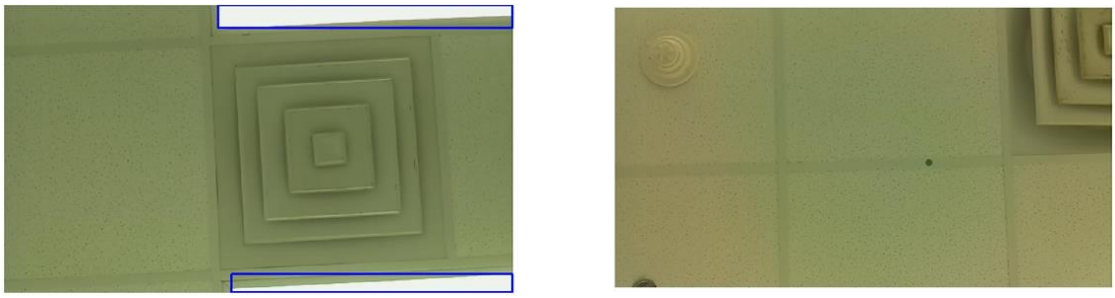

Figure 46. The light identification program successfully determines the presence of multiple lights (left) as well as the absence of lights (right).

4.1.4.2 Individual image processing

Individual images are processed by analysing grayscale images of lighting intensity. The images with longer exposure times (left side of Figure 45) are thresholded based on saturation of the camera. The purpose of this threshold is to quickly identify the presence of lights so that targets can be dynamically added to the navigation program real-time. After thresholding the image into a binary image, the center of each identified light and its accompanying bounding box are recorded and transformed according to vehicle location coordinates to compare with existing known locations of lights.

When processing images to determine the number of lamps per fixture, images with lower exposure times (right side of Figure 45) are first converted to grayscale

images. The bounding box from the images with longer exposure times for light identification is used to reduce the image size of the analyzed image. This reduced image is thresholded based on Otsu's autonomous thresholding method [75]. In thresholding the image into a binary image, Otsu's method calculates the optimal threshold to minimize the inter-class variance of the pixels. Morphological operations are then performed on the binary image to eliminate gaps or errant pixels, following which a connected component algorithm is used to count the number of lamps in the image.

An experimental example for the lights in the office area (lower area of Figure 44) is shown in Figure 47. This figure summarizes the process through each of the individual steps. The initial image with a longer exposure time is thresholded based on saturation and a bounding box is applied as seen in (b). This same bounding box is applied to the image with shorter exposure times, as seen in (d). This reduced image is thresholded using Otsu's method in (e), and morphological operations are performed on the image in (f). The total number of lamps in the fixture is then counted based on the number of unique objects.

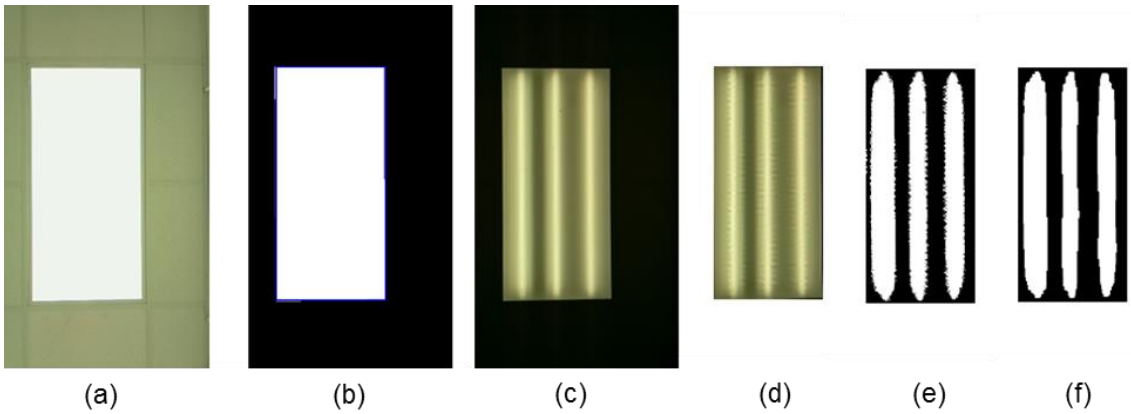

Figure 47. Step-by-step overview of individual image processing. The original image is thresholded based on saturation of the camera to find a bounding box for the darker image. This darker image is analyzed through thresholding and morphological operations to identify the number of fixtures and lamps in an image.

4.1.4.3 Image Stitching

As images are taken, they are stitched together into a composite image of the ceiling. The composite image of all lighting is used for two reasons. First, the composite image is used to prevent double counting of lights. As lights are identified in images and consequently stitched to the composite image, the locations of identified lights are merged with previously identified lights. Only newly identified lights are sent to the navigation program to be added as targets to avoid inefficiencies in visiting a single location multiple times. Second, the composite image is used to avoid errors in classification of lighting dimension or shape. The composite image is used to identify images that contain only partial segments of lights to avoid using only the partial segment when analysing the light.

4.1.4.3.1 Image Stitching using Keypoints: Methodology

Images can be stitched to form a composite image using a well-defined method from the literature [76]. In this method, keypoints in each individual image are identified. These keypoints are then matched with similar keypoints in other images, relating how the two images correlate to each other. A filter is applied to remove false matches, following which a homography matrix is calculated. An image transformation is used to warp one image, which aligns the keypoints from the two images. The images are then merged, and blending or feathering is used on the overlapping region of the images. Figure 48 summarizes the process used to merge images.

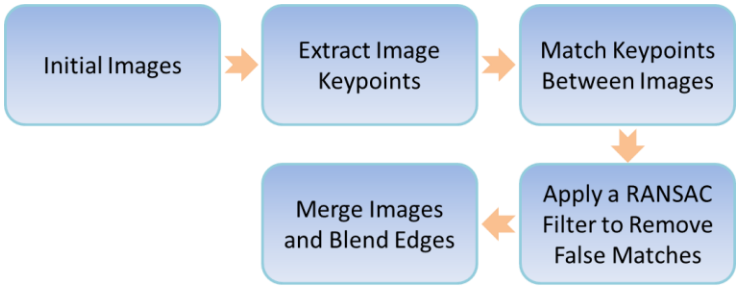

Figure 48. Overview of the image stitching process used to form a combined image map using image keypoints.

In merging images using this method, several constraints need to be met. One of the main constraints published in literature is that the origin of the camera is rigidly fixed with only one or two degrees of motion (rotation about its own axes) [77]. In this particular project, the vehicle is constantly moving; hence, the origin of the camera will move along with the vehicle. This issue is shown below in Figure 49. If the camera remains in position A, the read sphere will never be revealed for any orientation. But as

the camera moves position, the image camera will see different sides of some objects and some new objects entirely. If one attempts to merge two images and one object, (such as the red sphere) is only apparent in one of the images, the merged image will appear distorted. This problem is worse for objects in the foreground of images. In this project, ceilings and most objects suspended from them typically remain in the background, alleviating the distortion from moving the camera center. Additionally, Brown has noted that current methods are robust to small amounts of movement of the camera origin [77]. Therefore, in order to mitigate this distortion, images should be captured with large overlap and small distances between the camera's locations. In the current setup, images are taken as rapidly as reasonable with hardware constraints to allow for the best image merging possible.

The initial step to stitching images is the determination of keypoints in each image. In this project, keypoints were identified using a Scale Invariant Feature Transform (SIFT) [78]. These features are invariant to rotation, scaling and affine geometric transformations. The invariant nature of these points allows for similar keypoints to be matched between images despite differences in zoom, orientation, or camera center.

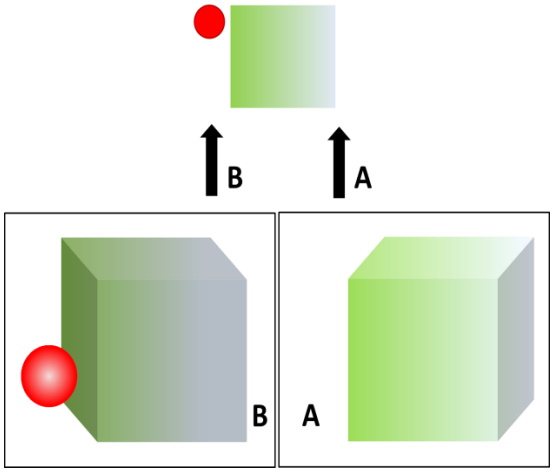

Figure 49. When the camera origin changes, additional features of some objects appear and some new objects appear. This creates distortion problems when merged two images such as seen in views A and B.

As images are taken during the flight of the UAV, overlap occurs between subsequent images. The overlapping portion of images can be used to map keypoint matches by looking at the Euclidean distances between the SIFT feature vectors. Once all matches have been determined, a Random Sample Consensus (RANSAC) algorithm is applied to eliminate false matches [79]. This algorithm takes a random set of matches and establishes a correlation between to the matches in the two images. Each additional set of matches is compared to this correlation, and if these matches fit the correlation within a defined threshold, they are considered inliers to the set. This method is completed for many iterations, with a new random set of matches forming the initial correlation in each iteration. The solution set with the most inliers is considered the best correlation between the images. This method has proven robust even when the number of inliers is not the majority of keypoint matches [79].

Using this correlation, an affine transformation is made to align the matching points in the two images. An affine transformation allows for scaling, rotation, shearing and squeezing of the image, similar to what is seen in Figure 50. After performing the affine transformation and merging the images, the amount of overlap is determined and excess overlap is trimmed from each image. The small amount of remaining overlap is linearly blended between the two images. The combined image forms the initial composite image and additional images are joined with this composite image until all captured images have been added. The final composite image can be used with the previously-developed program to quantify the bulbs in an image.

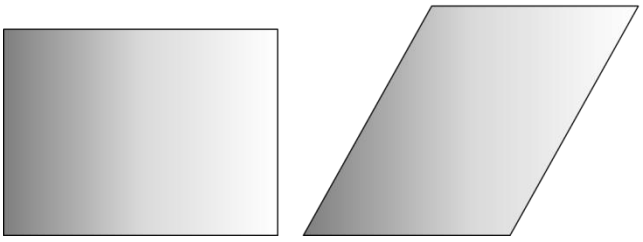

Figure 50. An affine transform. This transform allows for translation, rotation, shearing, and squeezing an image.

4.1.4.3.2 Image Stitching using Keypoints: Experimental Demonstration

An example of the complete image stitching process using the keypoint method is demonstrated for the images displayed in Figure 51. Although the automatic exposure adjustment of the camera altered the brightness of neighboring images, the chosen SIFT algorithm for identifying keypoints is invariant to these biases of brightness. Stepping

through each image individually, neighboring images were matched and merged onto the composite image.

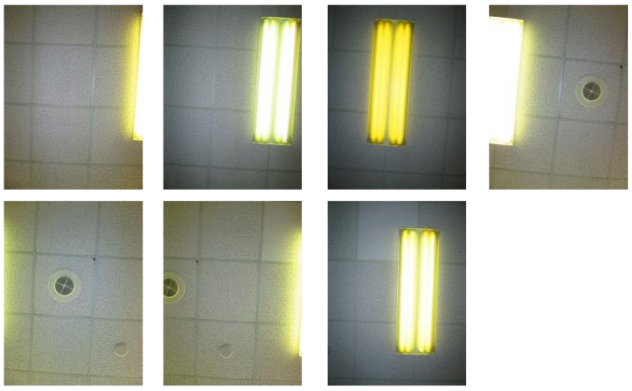

Figure 51. Images of a ceiling prior to merging. The intensity of the backgrounds varies depending on the automatic exposure settings of the camera. The images proceed from left to right, top to bottom.

After identifying keypoints in each image, the keypoints were matching for neighboring images, as seen in Figure 52. This figure depicts the two images to be merged placed side by side with lines connecting matching keypoints between the images. Prior to applying the RANSAC algorithm, there are a few false matches as seen on the image on top. After removing these false matches through the RANSAC algorithm, the final set of matches used to establish a homography is shown on the bottom image of Figure 52.

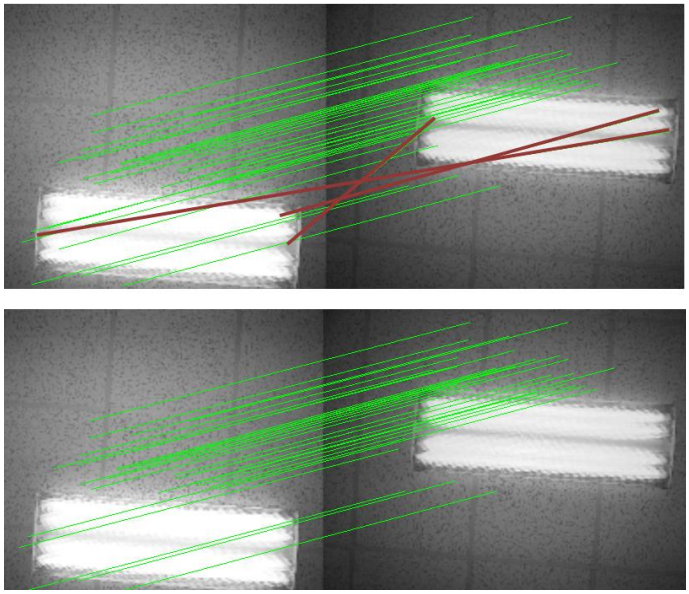

Figure 52. Matched keypoints between two images in the image set. The top set of images show the original images (with false matches emphasized), and the bottom set show the new set of images after the RANSAC filter is applied.

Once the homography matrix was established for each set of images, the images were transformed and then merged. Initially the areas of overlap were linearly blended so that each image contributed equally to the final merged image. The majority of overlap was then removed from each image, and the remaining overlapping region was linearly blended. The full overlap and subsequent reduced overlap is displayed on the images in Figure 53. Future work could include more advanced methods of feathering and blending to give a more seamless resultant image.

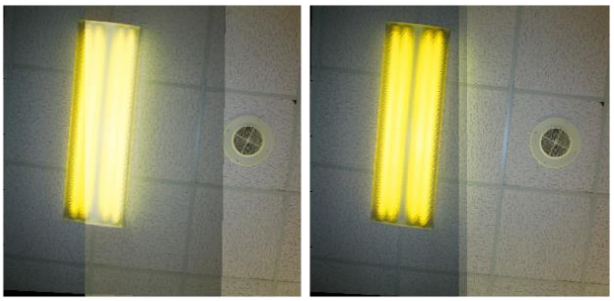

Figure 53. The merged images of the original 3rd and 4th images from the image set. The one on the left shows the images with all of its overlap and linear blending. The merged image on the right has much of the overlap removed.

These image merging steps were continued for each subsequent image until the entire image set was merged. Figure 54 shows the final merged image from the original 7 images. The differences in background intensity are very noticeable, but applying gain compensation and multi-band blending have been shown to effectively remove these discrepancies in other work [77].

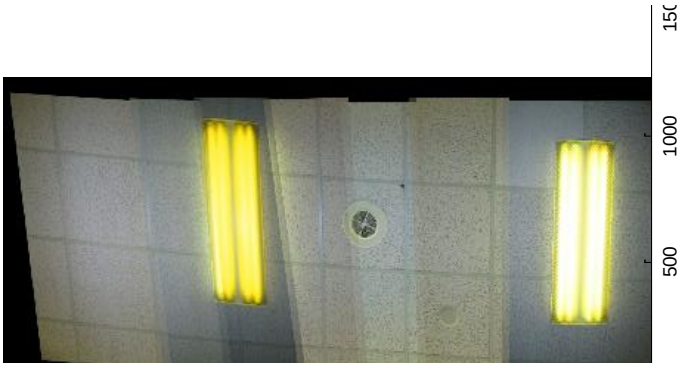

Figure 54. Final composite image of all 7 original images.

Since the purpose of combining the images is to assist the navigation system on the location of lights, real-time merging of images is important. Merging the images is computationally expensive. Since the image keypoints are the characteristic feature of

the images used to determine the necessary transformation to merge the images, only the keypoints and location of lights need to be retained in the program. By removing actual image merging, the global position of the identified lights in each image is determined by applying the image transformation. Retaining only the keypoints allows the images to be processed real-time. The composite image is only useful for visualization and validation purposes. Therefore, additional advanced methods of blending are not necessary for functionality in this paper.

4.1.4.3.3 Using Navigation Information

In the preceding section, images stitching was demonstrated by using the established keypoint method to identify matching keypoints between images. The demonstration of the image stitching method was performed on a ceiling with texture and sufficient detail. One of the primary challenges associated with this method is the potential lack of necessary keypoints. Keypoints are typically identified in areas with detail and contrast in light intensity [78]. For ceilings with smooth surfaces and little detail, the image stitching program breaks down as insufficient keypoints are available.

As the purpose of generating a composite image is to prevent double counting of lights and prevent errors in the analysis, the accuracy available through conventional image stitching isn't necessary. Wider errors are tolerable as individual images are used for individual light processing. Therefore, image stitching can be accomplished through the spatial information available through the navigation program. As navigation indoors relies on SLAM, the localization information is readily available and can be used to determine the necessary transformation to align the images. By using navigation

information, the robustness of the image merging algorithm is improved as the effectiveness of the algorithm no longer depends on the ceiling features of the building. The accuracy of the image merging algorithm then depends on the localization and orientation information from the navigation system. Although the image blur from stitching is increased because of uncertainty on the vehicle's location in the SLAM algorithm, the increased blur is acceptable as long as it achieves its purpose in preventing double counting of lights.

As each image is recorded, the global location of the vehicle is recorded by the image stitching algorithm. The images are transformed with an affine transformation using the recorded orientation of the vehicle and scaled using information from the laser range finder. There will still be potential perspective issues when merging images due to the center of focus moving, but these problems will be minimal because most of the images only contain objects in the background. New images are added to the composite image until the space has been completely explored. In each individual image, the local location of the lights relative to the vehicle is determined with the laser range finder, camera parameters, and image analysis. The positions of the lights are transformed according to the image transformation and compared to the existed list of light locations to determine if the light has previously been observed. Locations of newly discovered lights are sent to the navigation program as target locations.

This method was tested experimentally for the office area of the space described by Figure 44. Figure 55 shows the image stitching of the office area using the known location and orientation to transform the image. As seen in the figure, there are small

issues of misalignment in the lights, specifically the two lights in the upper left corner. However, these misregistration errors don't prevent the composite image from successfully identifying all of the lights in the explored area. In comparison to the office space presented in Figure 44, the composite image matches closely to the lower area that correlates to the office area. For purposes of navigation, there is no need for an actual composite image as only the coordinates of the lights need to be tracked. The composite image serves as a visual aid to evaluate the effectiveness of determining duplicate light locations.

Figure 55. Demonstration of image stitching using known locations. In the proposed system, the known locations come from navigation and exploration information.

4.1.5 Lighting Classification

Reliably determining the light type is necessary to make accurate estimations on potential savings for different lighting recommendations. In determining lighting type, a spectrometer is used to analyze the incoming light spectrum. A spectrometer specifically measures the intensity of incoming light over a spectrum of frequencies. This spectrum

is often called the spectral power curve. Different types of lighting exhibit unique spectral power curves and allow for differentiation of lighting type [80]. In the sensor package, the spectrometer is mounted onto the frame with a 400 micron optical fiber protruding out the top of the structure as the data collection point. Light entering the aperture of the optical fiber is collected and guided to a photodetector. Photons striking the photodetector are used to produce the spectral power curve. These curves are typically different in both relative magnitudes across various frequencies and location of intensity peaks. The spectral power curve for three different lighting types are displayed in Figure 56 and summarized in Table 10. Although similar, the spectral curves for incandescent and halogen bulbs have a peak value at different wavelengths and other subtle, distinguishing differences.

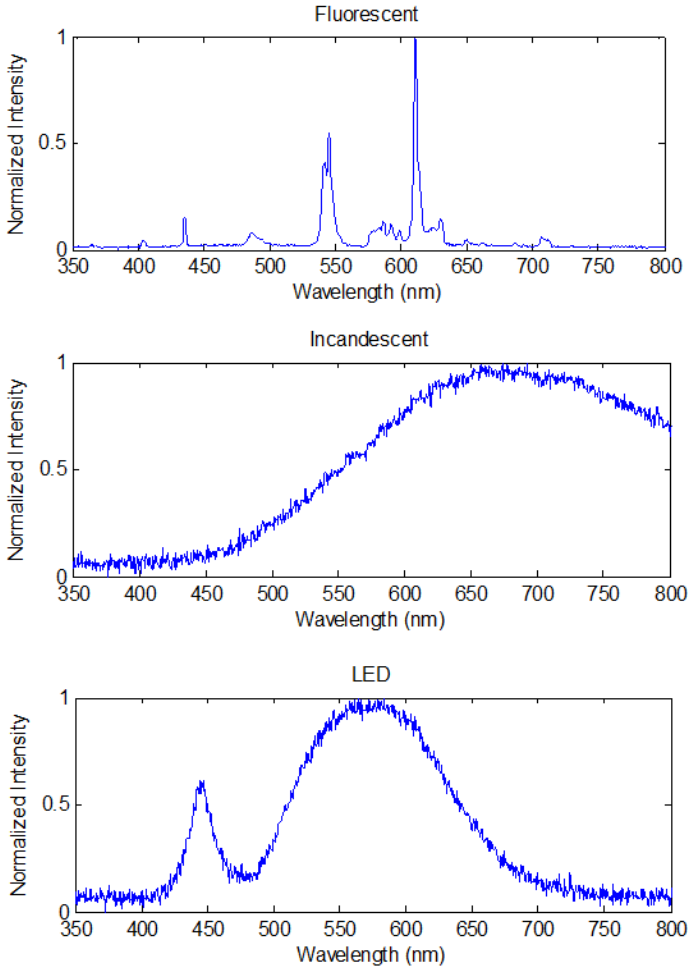

Figure 56. Spectral power curves for fluorescent, incandescent, and LED lighting, respectively. The spectral power curves are used to classify lighting.

Table 10. Description of Spectral Characteristics for Different Lighting Types

Light Type	Spectral Curve Descriptor	
	Wavelength Peaks [nm]	***Curve Type***
Incandescent	650-670	Smooth
Halogen	640-660	Smooth
Fluorescent	436, 546, 612	Spectrum Spikes
Metal Halide	819, 671	Spectrum Spikes
High Pressure Sodium	819,569	Spectrum Spikes
LED	456, 548-565	Smooth, Dual Peak

There are multiple ways to compare the peaks in determining the lighting type, and in this work a root-mean-squared (RMS) error approach is used. The incoming light spectrum is normalized to one according to the peak value among all frequencies. The whole light spectrum is then compared to the reference spectra for different lighting types by looking at the Euclidean distance between the measured and reference spectra. Specifically, Equation (55) shows the RMSE equation used when comparing each curve. If the incoming light spectrum matches any curve with a RMSE of less than 0.15 (average error of 15%), the light is classified according to the closest match.

(55) $$\mathrm{RMSE} = \frac{\sum_{i=1}^{n}\left(y(f_i)-r(f_i)\right)^2}{n}$$

where

n = number of frequencies
f_i = frequency
y = measured spectrum
r = reference spectrum

This method for classifying lights was experimentally tested under a variety of conditions to gauge the robustness of the method. Measurements were taken from distances ranging from 1.5 meters to 5 meters, which encompasses the range of expected distances in the field. The spectrometer was not shielded from ambient artificial and natural lighting, thereby creating a realistic field environment. A total of twenty tests for

each of the lighting types were tested, and from those tests, the algorithm correctly classified the lighting in 97% of the instances. The results of the tests are summarized in Table 11. In the instances of misclassification, the program incorrectly identified the incandescent bulbs as halogen lighting. This work compared the RMS error of the entire spectrum to the reference spectra. The classification accuracy could potentially be increased by using the specific peaks locations of different lighting types. These additional descriptors, as seen in Table 10, will be explored in future work to improve classification of lighting type in the facility.

Table 11. Light Identification with the Spectrometer at Varying Distances

Light Type	Number of Correct Classifications	Number of Tests
Incandescent	18	20
Compact Flourescent	20	20
LED	20	20
Halogen	20	20

Since the spectrometer can only measure the incoming spectrum over a small viewing angle, the vehicle must navigate directly under the light to obtain an accurate reading. For this reason, when a light is identified, the navigation algorithm incorporates these lighting locations as target locations. When the vehicle navigates to each target location, the spectrometer may obtain a reliable reading to accurately classify lighting

type. The incorporation of lighting targets into the navigation strategy is presented in other work and not the focus of this chapter.

4.1.6 Lighting Simulation

Correct lighting levels are an important part of creating a safe and productive workplace. Detailed tasks can require a much higher lighting level than general lighting to prevent stress on occupants. The illuminance, given in lux or footcandles, is a measure of the amount of light flux striking a surface per unit area. This is the measure used to determine the lighting level at a specific location [81]. The recommended light levels for different tasks vary greatly, from lower levels of approximately 150 lux for warehouses to over 5,000 lux for visual tasks of smaller size and prolonged periods of time. These recommendations are published by the Illumination Engineering Society (IES) in the IES Lighting Handbook [82]. The recommendations for various tasks and areas are only guidelines for recommended levels; consequently, these levels aren't always achieved during initial construction of a building or during remodelling. Under lit spaces can create an unsafe or taxing environment on occupants, while over lit areas present opportunities for cost savings by reducing light levels to recommended levels.

The light level in a space is heavily dependent on the location and direction of measurement. The light level from the same location but different orientations can vary by orders of magnitude. Additionally, the light level in different locations can vary significantly. Often, a representative light measurement will be taken to estimate overall lighting conditions, but with such a small sample size, accurate lighting conditions are difficult to ascertain. Estimation of lighting levels through simulation can provide

guidance on how lighting levels vary throughout a space, providing a means to quickly and accurately analyze potential energy-saving recommendations.

Light level measurements are taken in the working plane of the desired measurement. For specific task lighting measurements, the working plane is the table or other surface that is used for the task. The light meter is directed towards this working plane. For general lighting levels in non-specific tasks, such as in a warehouse or hallway, the working plane is defined horizontally. Light measurements are taken with the meter directed horizontally at eye level. Several measurements are typically taken in different directions to obtain an average light level.

Once the lighting simulation is created, several recommendations can be made using the additional information available from the simulation. First, in cases of over lighting, the correct number of lamps or fixtures to remove can be readily determined in achieving the proper light level. Second, in cases of industrial lighting, often the light fixtures are placed near the ceiling. With high ceilings, the light fixtures can be lowered, which creates delamping opportunities while still maintaining similar light levels. Lowering light fixtures typically will increase the light level at the working plane at the expense of decreased uniformity of light. The trade off and optimal distance to lower the fixtures can be determined in simulation. Finally, future work could explore adding natural lighting to the simulation. Use of available daylight instead of artificial light can be estimated with knowledge of window placement and area.

4.1.6.1 Point-by-Point Method

In general, simulation of vertical illumination is useful to evaluate task lighting by multiplying the vertical illumination with the reflectivity of the desired task surface. The horizontal illumination is useful for determining the general lighting characteristics in a space. Several methods are available to estimate both the vertical and horizontal illuminance in a space. The two generally accepted methods are the zone area method and the point-by-point method. The zone area method calculates the total amount of lighting in an area needed to provide adequate average illumination of a space based on the number of lighting fixtures, room cavity ratio, and other parameters. However, this method does not provide information on how light levels vary throughout the space. The point-by-point method finds the illuminance at specific locations in an area with a given set of luminaires and can be used to look at variation in lighting levels throughout a space.

The point-by-point method looks at the light level at every point of interest by summing the total light contribution of each light fixture at that point. The illuminance over an entire area or room can be determined by dividing the room into a grid of points. Although computationally expensive, this method gives high resolution of lighting levels throughout an area. The light level is a function of distance between the light fixture and the point of interest, the angle between the two points, and the light intensity of the fixture [81]. The governing equation guiding the point-by-point method is given by Equation (56):

(56) $$E_A = \frac{I_\theta \cos(\theta)}{d^2}$$

where E_A is the measurement illuminance at the point A, I_θ is the intensity of light from the light source for that specific angle, d is the distance between the light source and the point A, and θ is the angle between the two points, relative to the normal of the plane at A. Figure 57 illustrates these parameters. The intensity from the light source for the specific angle, I_θ, is a value that is documented by the manufacturer of the luminaire.

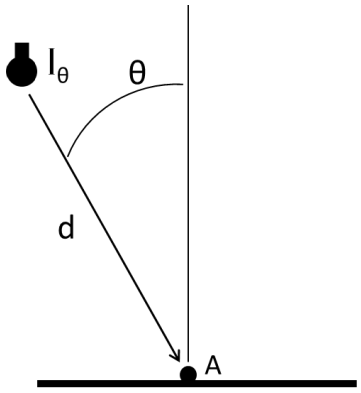

Figure 57. Parameters used in the point-by-point method of determining illuminance.

This method can be accurately used as long as the distance 'd' is at least five times the longest dimension of the light source [81]. For example, in a T8, 4-foot fluorescent tube, the longest dimension would be the length of the tube. This method could only be used if the distance to the point of measurement was at least 20 feet away. If such distances aren't possible, the lighting source can be segmented into sections with

the appropriately-scaled intensity values in order to satisfy these constraints. At any given point in the lighted area, the illuminance at that point of interest can be calculated by adding the individual contributions from each light source, or if the light source has been segmented, every segment of each light source.

In addition to these light sources, light reflected from the walls will also contribute to the measured illuminance. However, the computation burden drastically increases with reflection calculations; additionally, the dimensions of all the walls must be known relative to the positions of the luminaires in the space. In a large facility with many light fixtures, measurement of the distance from light sources to the walls can be difficult to obtain. When the primary objective of the lighting map is to identify areas of overlighting and find correct lighting levels based on the simulations, forgoing reflections will underestimate the true lighting level, thus providing a conservative estimate for delamping opportunities.

4.1.6.2 Light Intensity Variation

The intensity of a light source is given in candelas and is a measure of amount of light flux given in a specific solid angle [81]. The intensity of a light source is documented by the manufacturer of the light source and given in a special file type. These file types are regulated by IES and follow a specific format [82]. These files document the variation of intensity of light from all angles from the specific lamp and light fixture. They account for the reflected light off the fixture, allowing easier implementation of realistic lighting levels without increased complexity. Manufacturers provide these light intensity files not only for lamps, but also for entire fixtures. The

variation in lighting intensity can be significant for different angles and is important to accurately estimate light levels. An example of this significant variation is shown for a lighting intensity file for a metal halide fixture in Figure 58.

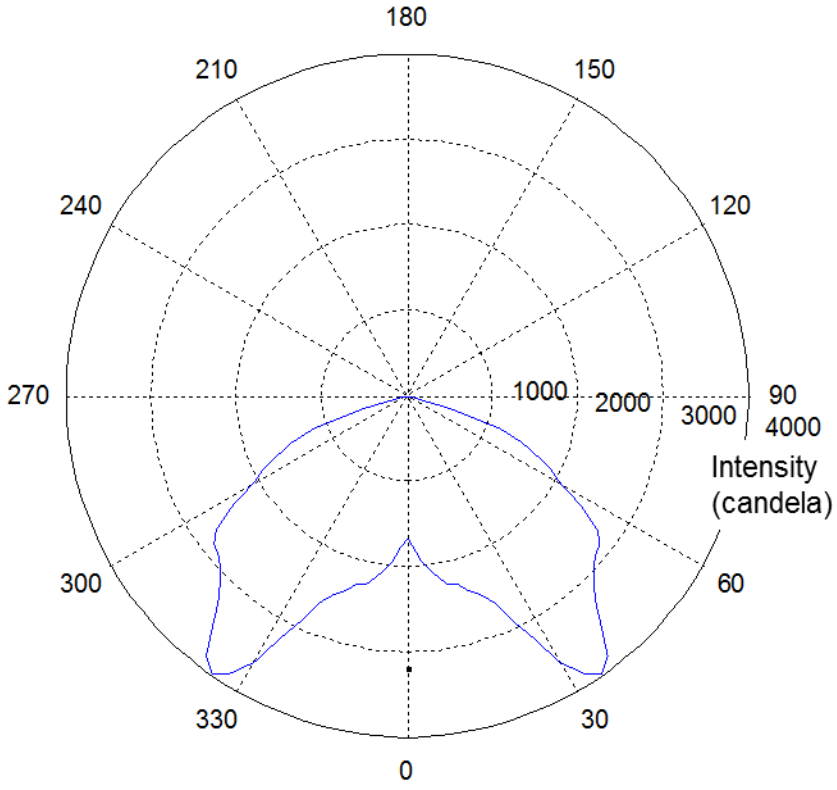

Figure 58. Example of a lighting intensity diagram for a metal halide fixture.

The cross section of the light intensity graph changes varies as rotated around the graph. Figure 59 displays a 3D version of the light intensity data provided for the metal halide fixture. As the fixture is rotated, the cross-sectional shape of the light intensity diagram changes considerably. In simulating the light levels, these values for light intensity are incorporated into the calculations to improve lighting level estimation.

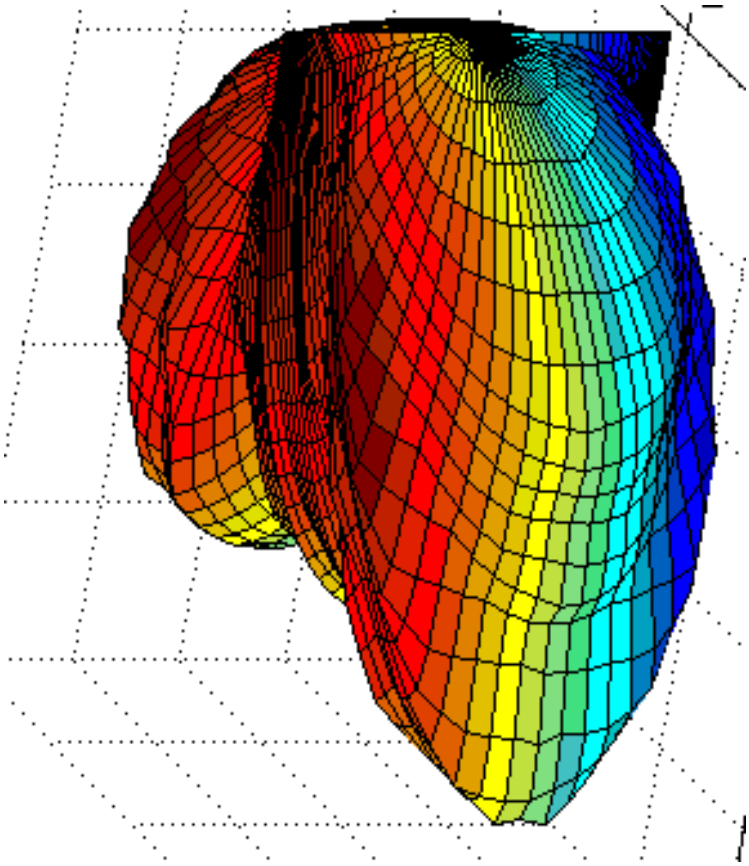

Figure 59: 3D visualization of light intensity data of a metal halide fixture.

4.1.6.3 Simulation Results

The layout of the room used for simulation is the same layout used for the exploration algorithm, example lighting calculations, and image stitching demonstration (Figure 44). This room consists of four different spaces, each equipped with fluorescent tube lighting fixtures. Each lighting fixture contains three, T8 fluorescent lamps. In the simulation, the light levels were simulated using Equation (56) and the appropriate light intensity diagrams from manufacturers of the different light types. Reflections off the walls are not accounted for in the simulations, and this will result in an underestimation of the actual light levels in the space.

The results for simulating the lighting for three different lighting types are given in Figure 60. In this figure, the results are displayed for the current T8 fluorescent lamps, as well as two other possible configurations of incandescent recessed lighting and LED replacement tube lamps for the T8 fixtures. Each simulation file used the actual light intensity data from manufacturers of the specific lighting type. The vertical illuminance shows the light intensity if a light meter was directed towards the ceiling. In order to determine general lighting characteristics in a space, the horizontal illuminance is also shown. Since walking or other activities are not necessarily performed in a single direction, the averaged light level from multiple directions is advantageous to better understand the representative light level. At each location, several readings can be recorded and averaged in multiple directions in order to evaluate the average light level. For the horizontal illuminance, the average horizontal illuminance was determined for four different directions, two in each direction of each axis. The average value from these four light levels was used for the average horizontal illuminance.

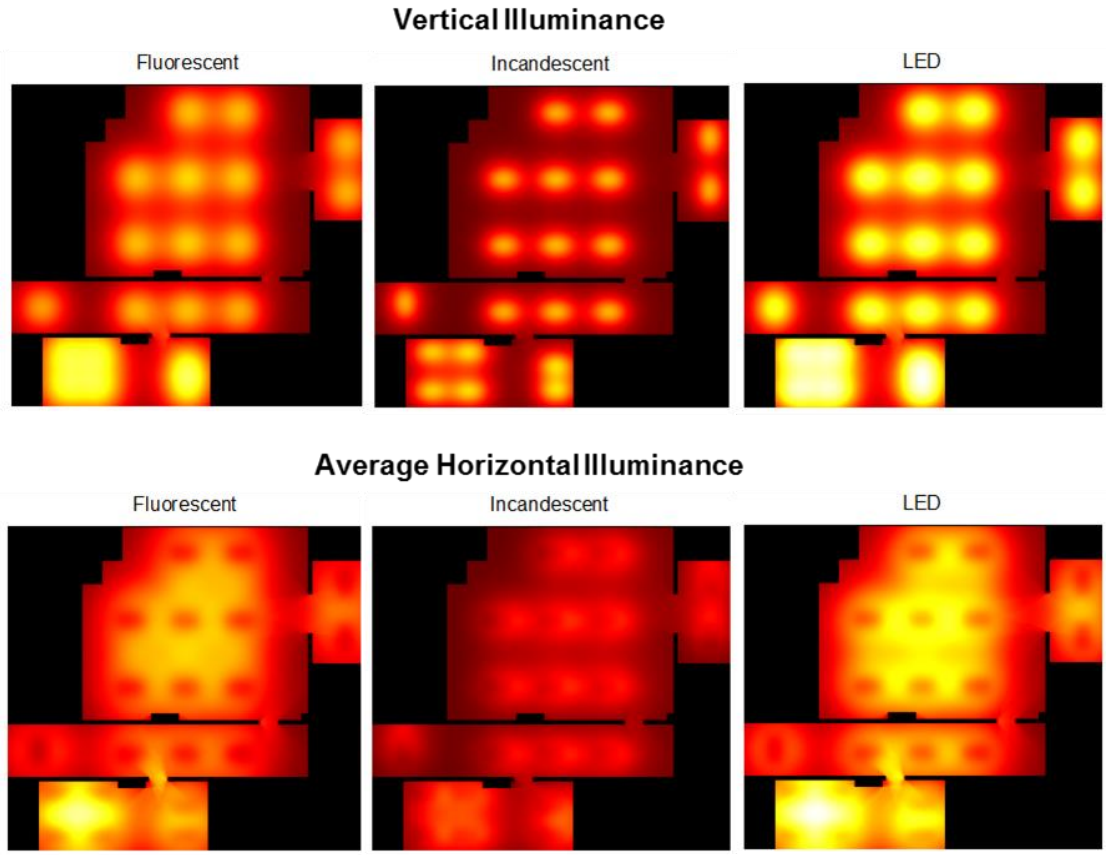

Figure 60. Simulation results for potential lighting types in the space. Both vertical and average horizontal illuminance is displayed.

As can be seen in the figure, the average light level for vertical illuminance for the fluorescent lighting is lower than LED and higher than incandescent. In average horizontal lighting, the light level is still higher than incandescent and comparable to LED. The light distribution is very similar for both fluorescent and LED lighting, and the average horizontal illuminance is fairly uniform over the space. Due to light reflection off walls, the predicted values from the simulation will tend to under predict light levels, particularly near the walls. The actual light level will be fairly uniform throughout the space as the edges will have higher light levels from reflection. There is

one instance in the bottom room where the space is potentially over lit, presenting an opportunity for energy savings for delamping while still maintaining adequate light levels.

4.1.6.3 Future Advances in Lighting Simulation

In future work, once the light topology map is generated, the map can be calibrated with several actual light readings taken during data collection on the vehicle to improve accuracy of the simulation. Additionally, the lighting plots can potentially be used in aiding the navigation to find additional lights with the addition of a light sensor on the vehicle. The computation burden of the simulation is low enough that the program can approach real-time in predicting light levels. As lights are identified and added to the simulation, measured lighting levels can be compared to predicted simulated values. Underpredicted values on the simulation program indicate the presence of undiscovered lights that will add to the lighting level at that point. Gradients on the difference between simulated and measured lighting level can predict the location of undiscovered lighting, providing feedback to the navigation program on where to search for additional lighting. This can be integrated into the program for more efficient searching.

4.1.7 Automated Report Generation

Once the type and quantity of lights are determined, the electrical energy and demand can be calculated to perform an energy efficiency analysis of lighting. The potential energy and cost savings are predicted for the different recommendations listed in Table 9. To estimate the energy savings from the different lighting recommendations, the energy use from different lighting types and operational conditions is quantified.

This reduction in usage is multiplied by the avoided cost of electrical energy to estimate the cost savings from the retrofit. If applicable, the reduction in costs for electrical demand is then calculated. Other potential savings are also calculated, including reduction in maintenance costs, personnel costs, and reduction in equipment replacement costs.

Also included in the recommendation report is a basic analysis on estimated implementation costs of the retrofit, which, coupled with the cost savings, provides a method to estimate simple payback periods. These implementation costs include the estimated costs for capital equipment and costs associated with installation. In the Texas A&M IAC, projects are typically recommended only if they yield a simple payback period under approximately two years.

The energy use is determined from Equation (57). Note that the variables used in the equation are summarized in Table 12. In this equation, the total lighting energy is the product of the lamp power, number of total lamps, the ballast factor and operational hours. The ballast factor accounts for the energy consumption of the ballast necessary to regulate the lamp, and this factor adjusts the lamp base output. For lighting without ballasts, the default value for the ballast factor is 1.

Table 12. Parameters to Calculate Energy and Demand from Lighting

Symbol	Parameter
x	Number of fixtures
L	Number of lamps per fixture
P	Lamp power (kW)
B	Ballast factor
T_O	Operational time
K	Annual billing cycles

Different lighting retrofit options will alter different parameters in the energy equation. For example, upgrading to a more efficient light type will alter the number of lamps per fixture and power per lamp. Likewise, installing occupancy sensors and photocells will reduce the operational hours of the lights. For each retrofit, the energy savings are calculated using the difference in energy prior to and following the retrofit as calculated in Equation (57).

(57) $$E_C = (x)(L)(P)(B)(T_O)$$

The calculation for annual electrical demand is similar to the calculation for consumption, but the operating hours are replaced by the number of billing cycles in one year as seen in Equation (58). The demand in a facility is a charge based on the peak rate of power in a facility over the course of each billing cycle. Only some of the possible lighting recommendations will consistently generate electrical demand savings. For example, occupancy sensors will only reduce lighting power when a space is unoccupied, which may not coincide with periods of peak demand. If the power

consumed by lighting is consistently reduced, the peak power draw over the month will decrease, generating demand savings. The cost associated with demand is the product of the annual demand and the avoided cost of demand.

(58) $$E_D = (x)(L)(P)(B)(K)$$

Once sufficient information has been gathered to determine current and potential electrical energy and demand use, the potential savings from different recommendations can be calculated and included in a report. The most common recommendation for lighting, as seen in Table 9, is lighting upgrades. More efficient lighting reduces both the electrical consumption and demand while maintaining similar illumination and quality of light. Additional savings are also possible through reduced maintenance if the recommended lamps have a longer lifetime than existing lamps. The following paragraphs demonstrate the general recommendation procedure that would be included in a report for a lighting upgrade with an example calculation. In this demonstration of the energy efficiency analysis, the actual lighting from the space (shown in Figure 44) is used.

In the calculations, recommendations are made for upgrading the lighting type from incandescent recessed lighting and T8 fluorescent lighting. Although the physical space is equipped with T8 fluorescent lighting, the calculations for a lighting upgrade from incandescent recessed lighting will be used to illustrate a recommendation within a simple two-year payback period. Table 13 shows the common operating parameters

found in an industrial or commercial setting that will be used in the calculations to determine the energy savings. These parameters are independent of lighting type and are consistent across all the types of recommendations found in Table 9. Usually the values are obtained from a building manager prior to collecting and analyzing recommendation data.

Table 13. Operating Parameters for Example Facility

Parameter	Units	Value
Annual Operating Hours	(hr/yr)	3,744
Avoided Cost of Consumption	(US $/kWh)	0.06
Avoided Cost of Demand	(US $/kW·mo)	5.2
Labor Cost	(US $/hr)	25
Number of Billing Cycles Annually	(mo/yr)	12

The parameters specific to each lighting type are listed in Table 14. The number of lamps and fixtures for the three different lighting types listed in the table were determined to ensure similar lumen output. The parameters for the T8 lighting are the actual parameters found in the space. The costs for lamps and fixtures were averaged from several manufacturers. The energy consumption before and after a lighting retrofit is used to determine the energy consumption savings. The initial energy consumption is calculated using Equation (57). The total cost of the current energy use is the product of the electrical energy consumption and the avoided cost of electrical consumption. The demand cost is also calculated in a similar manner, according to Equation (58).

Table 14. Lighting Type Specific Parameters

Parameter	Recessed Lighting (Incandescent)	T8	LED
Number of fixtures	140	20	20
Lamps per Fixture	1	3	3
Lamp Power (kW)	0.071	0.032	0.021
Ballast factor	1	0.87	1
Annual Energy Use (kWh)	37,215	6,254	4,717
Annual Consumption Cost (US $)	2,233	375	283
Annual Demand (kW*mo)	119	20	15
Annual Demand Cost (US $)	596	100	76
Cost per Lamp (US $)	N/A	3	24
Cost per Fixture (US $)	N/A	75	75

Once the base consumption and demand are calculated, energy and cost savings are determined and the simple payback period of the project is calculated using the updated number of lamps per fixture and power per lamp. The annual energy and demand savings from switching lighting types is the difference of energy consumption and demand, respectively. The capital cost is the cost associated with buying new lamps and fixtures. The installation cost is the product of the number of fixtures, the installation time per fixture, and the labor cost. The sum of the capital and installation costs results in the total cost of the project. The simple payback period of the project is then determined as the quotient of the total cost of the project and the annual cost savings.

The results for three potential retrofit options are summarized in Table 15. As seen in the table, there are significant cost savings from replacing recessed incandescent lighting with either T8 fluorescent lamps or LED tube lamps. Both produce sufficient savings to yield a simple payback of less than two years. Upgrading the currently

existing T8 lamps to LED lamps falls outside a typical payback period for industrial and commercial facilities. Although local rebates may be available to help defray initial costs, a large payback will typically result in low implementation rates.

In an energy assessment, similar calculations are performed for every potential lighting recommendation. As seen in this example calculation, the calculations are generally straightforward, and the necessary parameters to assess energy and cost savings are established beforehand and can be collected from the AuRAE.

Table 15. Calculated Savings for Three Upgrade Options

Current Lighting Type Recommended Lighting	Units	Incandescent T8 Fluor	Incandescent LED	T8 Fluor LED
Annual Energy Savings	(kWh/yr)	30,961	32,498	1,537
Annual Demand Savings	(kW·mo/yr)	99	104	5
Annual Cost Savings	(US $)	2,354	2,471	117
Capital Cost	(US $)	180	1,500	1,500
Installation Time	(hr/fixture)	1.75	1.75	0.25
Implementation Cost	(US $)	875	875	125
Total Cost	(US $)	1,055	2,375	1,625
Simple Payback	(yr)	0.45	0.96	14

4.1.8 System Integration

Each individual subsystem for both lighting analysis and navigation has been demonstrated in this chapter with experimental work on actual lighting conditions in buildings. Other ongoing work by undergraduate research teams has focused on the integration of individual components into a single lighting assessment package. The teams have integrated the sensors into an operator-driven package that autonomously identifies and analyzing lighting fixtures according to the algorithms

presented in this chapter. The robot requires manual navigation, but mapping of the facility and analysis of lighting are performed autonomously. Figure 61shows the manually-driven robot that contains the sensor package. In creating the integrated package, the researchers built additional visualization capability by marking the identified light locations on the map created by the SLAM mapping algorithms. Several experimental tests have been run, including on a large section on the campus library. Figure 62 shows the identified lights in the library positioned on the autonomously-created map. Current undergraduate research teams are creating a handheld lighting sensor package that can be used by human energy auditors as they walk through a facility. An initial prototype has been designed and built based on the lighting analysis tools presented in this chapter.

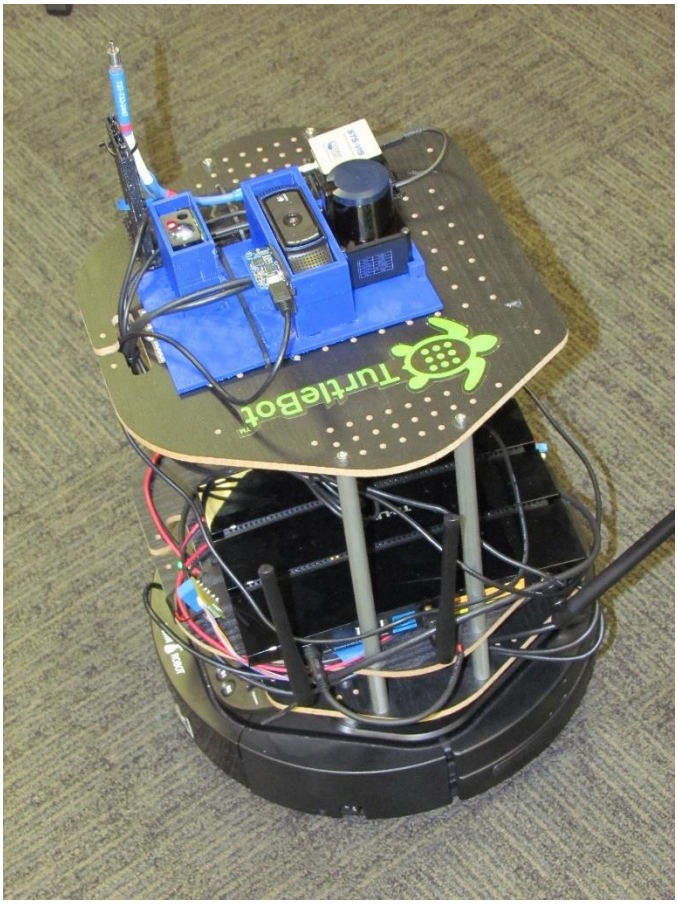

Figure 61. Prototype system that is manually driven but autonomously collects and analyzes lighting data.

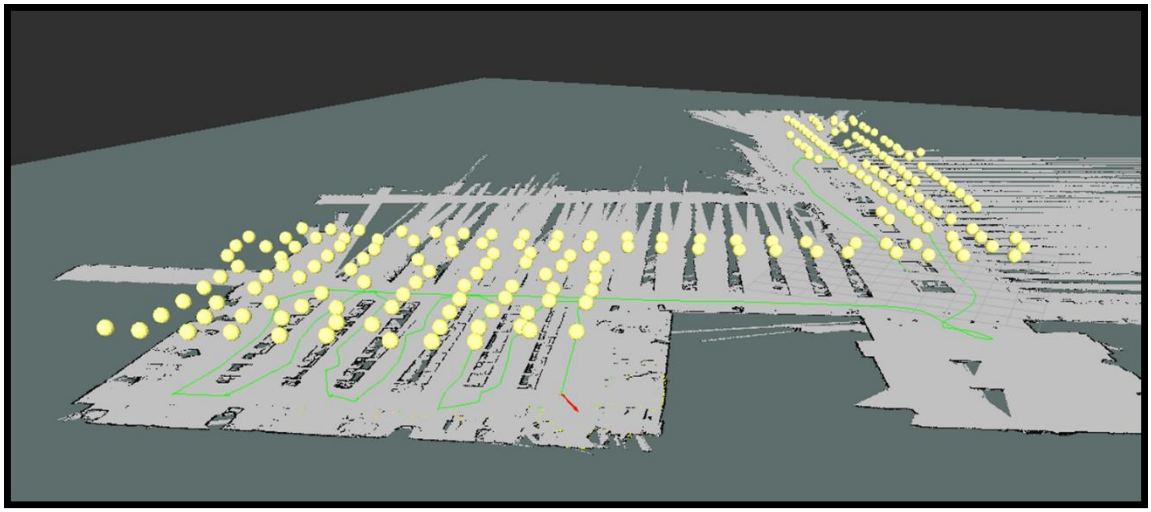

Figure 62. Identified lights in the campus library using the methods described in this chapter.

4.2 Summary

Energy assessments have demonstrated energy and cost savings by identifying opportunities for reduction in energy use with quick payback periods. However, the high cost and extensive time required for analysis inhibits the number of assessments performed. Automated data analysis for parts of or even entire assessments can decrease overall costs while improving assessment quality. Some of the most common recommendations in industrial energy assessments pertain to lighting. This chapter demonstrated a system that is capable of analyzing building lighting for identification and analysis of potential reductions in energy. The number of lights and their location were found in conjunction with exploration and navigation through image analysis and image stitching. The lighting type was determined through spectral power curves available through a spectrometer. Once the lighting type and location was determined, this information was used to predict lighting levels for both horizontal and vertical

illuminance in a space. Experimental evaluations have been performed for each component of the lighting analysis, and current work is focused on the integration of all these components.

CHAPTER 5

INTELLIGENT ACTUATION

Advanced sensors enable additional monitoring opportunities with the end goal of more intelligent control of equipment. The full potential of IIoT can be realized in achieving enhanced production, safety, or energy efficiency of the production facility. This chapter demonstrates how intelligent control can decrease energy usage with a case study on compressed air systems. Specifically, this chapter explores the use of a novel k-means algorithm to determine the production schedule of the compressed air system. Times with idle compressor load are identified, and the compressor is automatically turned off and then restarted prior to the next production period. This chapter explains the unique adaptations to the k-means algorithm and demonstrates its capabilities on both simulated and experimental data. A simulation of a typical industrial compressed air system is created for initial testing of the algorithm. The algorithm is then tested on twelve datasets of 15-minute demand data with profiles similar to expected profiles in compressed air systems. Finally, data are collected and tested from an actual industrial facility to compare the performance of the k-means algorithm to the performance of human actuation. The chapter concludes by exploring future avenues of research in this field.

5.1 Case Study: Improved Intelligent Compressor Operation using Machine Learning

5.1.1 Background and Introduction

While IAC assessments have demonstrated meaningful savings for participants, the overall percentage of total candidate facilities that receive a visit are minimal. Facilities that fall under the scope of the IAC program have to be small enough to not employ a devoted energy manager in the facility. Outside the IAC program, few small- to medium-sized manufacturing facilities receive energy assessments. With no allocated focus on energy efficiency, and with prohibitively high initial costs for energy assessments, common efficiency problems continue in these plants indefinitely.

One way to improve the realized energy savings in smaller manufacturers and achieve larger market penetration of energy efficiency best practices is through turn-key autonomous solutions. In the residential market, smart thermostats are one embodiment of this concept to deliver energy savings with minimal expertise and time required of the user [83]. There are many energy subsystems with potential autonomous solutions in manufacturing plants. The most visible and simple energy subsystem in plants is lighting, where occupancy-based actuation is a common energy efficiency upgrade. Similar opportunities exist in other energy subsystems such as boilers, motors, product heating, space conditioning, and compressed air.

5.1.1.1 IIoT and Compressed Air Systems

This chapter explores autonomous solutions in compressed air systems. Compressed air has widespread use throughout industry and is recognized for its

reliability. These systems can still meet necessary production requirements while being continuously undermaintained. This reliability ironically causes reduced system maintenance, resulting in energy efficiency penalties from compressed air leaks, high pressure, and unnecessary pressure drops in the system. Compressed air systems are widely considered a prime application for IIoT due to their high operating costs and maintenance issues [84-86]

The IIoT for compressed air systems is still in the early stages. As such, there is still uncertainty on the scope of possible compressed air solutions. Some have identified the IIoT for compressed air as an enhanced data stream for help operators make improved operation decisions [84, 85]. Other have envisioned entire automated control systems [87]. Since the field is still in its infancy, one current challenge is to identify actionably ways to improve system efficiency [88]. Researchers recognize that there is potential for improvement in overall compressed air system performance, but often they only provide a framework of how future solutions may work while providing no specific solutions [87]. Examples of potential actionable items to improve compressed air system efficiency include fixing air leaks, reducing system pressure, turning down or off the system during low production periods, fixing large pressure drops such as clogged filters, and performing compressor maintenance to ensure optimal compressor efficiency. The current scope of available market solution in the IIoT field fall primarily in two categories: intelligent monitoring of air systems and enhanced sequencing of compressors.

5.1.1.1.1 Intelligent Monitoring

Many market solutions have emerged that enhance monitoring capabilities of compressed air systems [85, 88-93]. Historically, individual sensors measured fragmented elements of the system, but emerging market solutions integrate multiple sensors to monitor overall system performance [84, 87]. These sensors can be integrated into plant-level monitoring and control systems, providing additional diagnostics for improved decision making by plant managers. Some of the sensors available for compressed air monitoring include pressure transducers, differential pressure transducers, power meters, dewpoint, thermocouples, and air flow meters [94, 95]. Newer compressors are being built with these sensors integrated into the compressor and require no additional third party add-on sensors [85]. Use of these built-in sensors can potentially reduce the cost of implementation of autonomous actuation algorithms.

In addition to data collection, many of the market solutions analyze and provide visualization of the data, enabling additional diagnostic capabilities. Some of the capabilities include efficiency monitoring of individual compressors in the context of the total system, identification of idle operating states in a system, and trend identification of overall air usage in the system [94]. Additional sensors also allow the detection of anomalous system operation or usage, such as unmet pressure setpoints or compressor short cycling [84, 87, 95]. A final common capability within current market solutions is in predictive maintenance for both compressors and other elements in the air system [85, 91]. Predictive maintenance prevents overall waste of materials and labor by only performing maintenance when necessary instead of on a predetermined schedule.

Current market solutions identify when maintenance is required and automatically notify plant personnel.

5.1.1.1.2 Intelligent Sequencing

The other current market solution in the IIoT field for air compressors regards intelligent compressor selection and loading [94, 96-98]. Often termed intelligent control, this strategy uses current plant conditions to select the optimal combination of compressors to both meet facility demand and optimize energy use. In systems with multiple compressors, individual compressors can operate inefficiently without a centralized control strategy for the compressors. As compressor efficiency is a function of compressor loading, correct sequencing of compressors can operate compressors near their optimal loading point to optimize overall system efficiency. This IIoT solution requires multiple compressors to implement, and multiple market solutions are available that calculate and implement this optimization.

5.1.1.2 Autonomous Solutions

The current market solutions enable improved decision making by leveraging the additional sensing capabilities and data analytics. However, these solutions require manual actuation to achieve the desired improved operation. These compressors are often equipped with remote start-stop and pressure setpoint capabilities. Plant managers can use the remote function to set timers on when the system should turn off and turn on to prefill the system [96]. Production schedule uncertainty limits the windows of time plant managers can confidently times to turn off the system. The methods proposed in

this chapter extend the timer operation by autonomously determining the appropriate times to schedule the system to turn off or reduce the pressure setpoints.

When there is no production on the plant floor, the compressed air systems can be turned off to conserve energy. Turning off equipment when not in use is a common recommendation on IAC assessments. Often compressors are left on during weekday nights and over weekends. The large leakage rates typical in plants result in significant cost penalties when not turning off compressors. Additionally, turning off compressors have maintenance benefits in reduced wear of components and decreased risk of catastrophic loss of the equipment if failure occurs when no one is present.

There are several primary barriers that block a universal adoption of turning off compressors when not in use. First, turning off compressors requires consistent effort from plant or other personnel. Although compressed air lines are located in convenient locations in a typical plant, often the compressors themselves are found in a separate building, outside, or on the far edges of the building. These locations are chosen specifically away from most plant personnel for noise reduction and compressor heat rejection. The last facility personnel to leave the building have little incentive to commit the time and effort to ensure all air usage has stopped and turn off the compressors, especially since there is no economic incentive to the particular employee. Second, the cost to run idle compressors is unexpectedly high, resulting in plant managers underestimating the operating cost. Furthermore, the cost of running individual equipment is hidden in the total facility electricity bills, which often is considered an unavoidable accounting expense to production. Therefore, the lost money due to

unnecessary operation is not fully realized. Finally, the variability of production and personnel schedules creates uncertainty on when turning off the compressor is appropriate. Without detailed monitoring of the system, the personnel in charge of compressor operating will unlikely know when extra production hours occur or even who will be the last employee to leave the facility. This uncertainty, together with the considerable consequences of turning off equipment during production, discourages manually turning off equipment at the end of production.

This case study specifically looks at autonomous actuation of compressed air systems when production is not running. The compressed air system loading is monitored throughout the day. The compressor automatically turns off after production has ceased and turns back on prior to the start of production. Most larger compressors are equipped with remote start-stop capabilities that can be used for actuation. These capabilities are unlocked and require no special equipment or propriety software to access. This case study specifically targets plants that operate on a 1 or 2 shift operation. A facility that operates on a 24-hour operation will have fewer time periods in which the compressors can be in the off state, lowering the cost effectiveness of the technology.

The cost to implement the technology is primarily in the cost of monitoring the compressor power and the cost of communication equipment. As these components are readily available and relatively low cost, the total implementation cost is expected to be modest. Additionally, newer compressors are being built with sensors integrated into the system, which could reduce the cost of implementation further. In general, the

anticipated energy savings, coupled with the modest implementation cost, would result in a manageable payback for implementation of this technology.

5.1.1.3 Machine Learning Algorithms for Load Forecasting

The proposed system in this case study will work by monitoring the load on the compressed air system. A novel k-means algorithm will be used to learned the operating schedules of the production plant. Once the algorithm has monitored the system for the prescribed training period, the compressor will be autonomously turned off during times when production is off and the compressor is idling. The algorithm will use the training data to determine the necessary start time of the compressor and restart the compressed air system prior to future production use.

5.1.1.3.1 K-means Clustering

K-means clustering is a data mining technique that identifies structure in unlabelled data. K-means is a pattern recognition algorithm in the general field of artificial intelligence. The basic algorithm finds a set of K clusters, where each data point belongs to the cluster with the shortest distance. It is an unsupervised learning technique, where no prior knowledge is available to specify how the data should be grouped. Initially, K data points are chosen at random to represent the K cluster centers, and each data point is assigned to its closest cluster center. The center of each cluster is updated to account for added data points to the cluster, and the algorithm continues to iterate and reassign profiles until a steady state has been reached.

There are many application fields for forecasting and clustering algorithms across many disciplines. Example application fields of k-means algorithms include

image classification, web text searching, and image denoising [99-104]. Within the energy field, the use of clustering algorithms varies from determining household electrical breakdown to comparison of electrical use between different cities [105-107]. Each field of clustering has unique clustering needs, and adaptations on the base k-means algorithm have often been proposed and implemented. Generally speaking, adjustment of the number of clusters and strategic selection of initial clusters provide a notable improvement over the base clustering algorithm, as these present the most common challenges in the k-means algorithm.

One common application for use for k-means that has many parallels to the work presented in this chapter is in load forecasting of electric demand by utilities. Load forecasting is used to by utilities to ensure sufficient electrical supply in the most cost-effective manner. Electrical load forecasts are broken into short-, medium-, and long-term time horizons [108] The work presented in this chapter aligns most closely with short-term load forecasting, so referenced papers to load forecasting focus primarily on this time horizon [108].

Another common use for clustering algorithms in the electrical utility industry is for grouping different customers [109-113]. This clustering is done to identify potential candidates for energy programs provided by the utilities. With large datasets of many customers, often a single representative load profile is used for each customer in the clustering algorithm. The algorithm divides the customers into different groupings based on these representative load profiles. This application domain is particularly well suited

for a k-means algorithm because of strengths is differentiating between distinctly different profiles.

5.1.1.3.2 Other Forecasting Algorithms

Other types of load forecasting algorithms have been explored in the literature. Most methods focus specifically on the time and value of the peak usage in order to guide utility providers in projected electrical demand requirements [108]. Lavin et. al. provide an overview of many of these methods in the literature [114]. The most common other methods are Artificial Neural Networks (ANN) and time-series methods. Regression models have also been proposed and can accurately characterize relationships within facility demand, but these methods require large computations and complicated modelling techniques [115]. Independent component analysis has also been explored for load forecasting [116].

Time series models are popular and still used today by many utilities, especially because they are capable of incorporating weather into the model. Time series approaches require a mathematical model to map the inputs to outputs, but they allow the ability to account for weather or other seasonal effects. One challenge in time series models is the large amount of required data to estimate the coefficients of the model. Sufficient training data are required to estimate dozens of coefficients without overfitting the data [117].

Neural networks, a common approach to load forecasting, are modelled after the working principle of the human brain with several layers of neurons [108]. There have been many papers on the topic published in the literature and some acceptance in

industry using this method [118]. These models can have multiple hidden layers and incorporate weather effects [119]. When setting up ANN models, the available data are divided into training and validation data. The hidden layers and network connections in the ANN is modelled after how the human brain learns new data, but the modelling of electrical load values provides no clear physical meaning for each of the nodes. As such, the available data is divided into the two datasets to ensure that the correct number of hidden nodes are used and that the data are not overfit [120].

ANN methods have shown success in load forecasting, but there are several limitations with this method that directly affect its applicability to the work presented in this chapter. This method, as opposed to other algorithms, the electrical demand data is not broken into profiles of individual days, but rather the method continuously forecasts the next look ahead period from the previous periods [118]. As one key element of the compressor actuation is in predicting when the compressor can be turned off and on, the knowledge of how production compares to previous days is desirable for this prediction. A second challenge in neural networks comes from load profiles that considerably differ from day to day [118]. Since the neural network continuously updates in the look ahead period, the algorithm has difficulty distinguishing different operating days. For example, in projecting the load forecast for a Friday or Saturday, the prior data used for the forecast will be similar. However, the anticipated forecast is expected to be different. To account for these differences, some research only forecasts based on weekdays [115, 120]. Other research clusters the daily profiles in a pre-processing step using a k-means algorithm. The ANN network is then applied only to the cluster with the most member

profiles [119]. Other challenges to ANN models include the large amount of required training data, the optimal amount of layers and inputs, and the possibility of over-training from the training dataset [118, 120].

5.1.1.3.3 Adaptations to K-means Clustering

The work presented in this chapter draws many parallels to the electric load forecasting and is well suited for a k-means algorithm. The k-means algorithm is simple in execution and allows for modification to incorporate physical constraints or phenomena [112]. The algorithm is capable of gracefully accounting for outlier profiles without affecting the overall estimation or forecasting. One day of compressor data is used for a single profile in the clustering algorithm, similar to what is done in electrical load forecasting.

However, the base k-means algorithm has several challenges that have been addressed in the literature with various approaches. Variations on the base k-means algorithm fall into three general categories. The first is in the autonomous selection of the number of final clusters. The second category is in the adaptive selection of the cluster centers used in the clustering. The final set of modifications come in extracting features to cluster instead of clustering the load profiles directly.

5.1.1.3.3.1 Autonomous Selection of the Final Number of Clusters

The selection on the number of clusters is a non-trivial problem. Use of the classical k-means with a pre-determined number of clusters and no change in the number of clusters gives non-optimal clustering results [113]. Some approaches use cluster validity indices to automatically check if the selected number of clusters is optimal [110,

121]. However, even using multiple cluster validation indices to evaluate the cluster, some work still relied on expert opinion to determine the final number of clusters [110].

Others have addressed this challenge by automatically adjusting the number of clusters during the algorithm. Sahu et. al. independently tested the adaptive k-means algorithm, and they found that the adaptive nature of finding the number of clusters improved the results over a standard k-means algorithm [122]. Kwac et. all use a threshold condition to determine when additional clusters are necessary, and then a trimming stage is done to merge clusters that are close to each other [112]. Both of these parameters for growth and trimming are user defined, so manual tuning is required to achieve an acceptable number of clusters.

Mutanen et. al. propose a clustering algorithm that automatically determines the number of clusters. The algorithm first removes outlying usage based on monthly average usage as well as intraday load variation. With the user-specified desired number of clusters and minimum samples in a cluster, the algorithm merges and divides the clusters iteratively under the number of final clusters is close to the desired number of clusters [109]. However, this merging and splitting process permanently removes profiles that, at any time, belong to clusters with fewer than the minimum samples in a cluster. To prevent the excessive removal of clusters as outliers, the addition and removal of clusters take place prior to removal of outliers. The authors note the challenges of clustering profiles that are distinctly different from each other [109].

5.1.1.3.3.2 Adaptive Selection of Cluster Centers

Another challenge in the classic k-means algorithm is in the initial selection clusters, as using random selection of initial clusters rarely provides the optimal solution. Tsekouras proposes a k-means algorithm for electric load profiling that examines many possible starting combinations for both number of clusters and initial cluster centers. Although computationally expensive, this exhaustive k-means algorithm demonstrated marked improvement over classical k-means algorithms [113].

Kwac proposes an adaptive k-means algorithm that adds clusters whenever a load shape violates a user-specified threshold condition. The user then specifies the number of target clusters, and then algorithm will hierarchically merge the closest clusters until the number of target clusters is reached [111]. Nikulen proposes a version of k-means clustering that involves both a threshold-based growth and merging stage [123]. Both stages require the user to pre-specify the thresholds for splitting and merging, values which are not known beforehand. This algorithm was tested against other network intrusion detection methods, and when threshold values were properly selected, this method outperformed other published methods.

Another approach to eliminate the dependency on the initial cluster centers is to use a hierarchal trimming procedure. Starting with the entire dataset, nearby datasets are grouped at different levels until the desired number of clusters is satisfied [124, 125]. One challenge to this method is the selection of the stopping criteria, similar to the challenge of selecting the number of clusters in the classical k-means algorithm.

5.1.1.3.3.3 Clustering on Features Extracted from the Load Profiles

Some approaches extract features from the dataset and cluster on these extracted features to improve the overall clustering or increase the algorithm speed. For example, Khan et. al. pre-process their 15-minute data in order to extract density distributions of the load profiles from the datasets [126]. Other pre-processing steps include using Markov models and wavelet transforms to identify load profile shapes that are then clustered [127, 128].

5.1.2 Intelligent Actuation Algorithm

The intelligent actuation algorithm builds on prior work in load forecasting for electrical demand. Key differences between peak load forecasting and this work require adaptations to algorithms found in the literature. The algorithm is broken into the following four segments:

1. Load Factor Estimation
2. Load State Information
3. K-means Estimation Algorithm
4. Compressor Actuation

The first segment of estimating the load factor is only necessary if pressure is the only available sensor and a load/unload cycle is utilized on the compressed air system. The timing of the load/unload cycles are determined based on the rise and fall of the pressure, and compressor loading is estimated from these cycles. Once the loading information is estimate or power information is recorded, the algorithm translates the load information into estimated load states. This second step is a key difference from

standard algorithms from the literature as the peak load during each profile is of minimal interest compared to knowledge of relative load states throughout each profile. After translating each profile to the estimated load states, a novel approach to a k-means algorithm is implemented for the training dataset. Standard approaches use a fixed number of clusters to classify each profile or adaptively determine this fixed number of clusters based on user-defined criteria. The approach described in this chapter allows an adaptive number of clusters that is autonomously determined through a growth and trimming stage. After identification of the training set, the algorithm steps through each day and determines periods of time the compressor can be shut down. The compressor is actuated for the identified idle time and autonomously turned back on prior to the next operating period. The following subsections detail how each of the algorithm segments are accomplished. In the algorithm, the user-defined parameters are the nominal buffer time (detailed in Compressor Actuation), the number of training and test days, and the cutoff threshold for outliers (detailed in K-means Estimation Algorithm).

5.1.2.1 Load Factor Estimation

The proposed algorithm in this chapter leverages the relative load on the compressor to estimate periods where the only energy usage comes from compressor idling power and air system leaks. To determine these times with low operating load, the relative compressor load must be known. If compressor power is measured, the average compressor loading is estimated by the mean power over a loading cycle or a defined period of time. Measuring the power of the compressor is the most direct way of monitoring the relative compressor loading, but pressure monitoring can also be utilized

if the compressor operates under a load/unload or on/off operating scheme. Use of pressure monitoring provides simultaneous monitoring of system pressure along with estimation of compressor loading.

The relative compressor loading can be estimated by observing each pressure cycle during operation. The estimated compressor load is the fraction of time the compressor is loaded over one compressor cycle. By autonomously determining the points at which the compressor cycles on and off, the loading can be determined as a pre-processing step to the intelligent actuation algorithm.

Figure 63 compares the compressor loading estimation for both pressure and power monitoring for one day. These data come from the industrial plant that is evaluated later in the chapter. In this figure, the compressor power is scaled to match the peak value in the day to provide clarity in the comparison. The compressor loading estimation from pressure compares favorably to the power monitoring of the system.

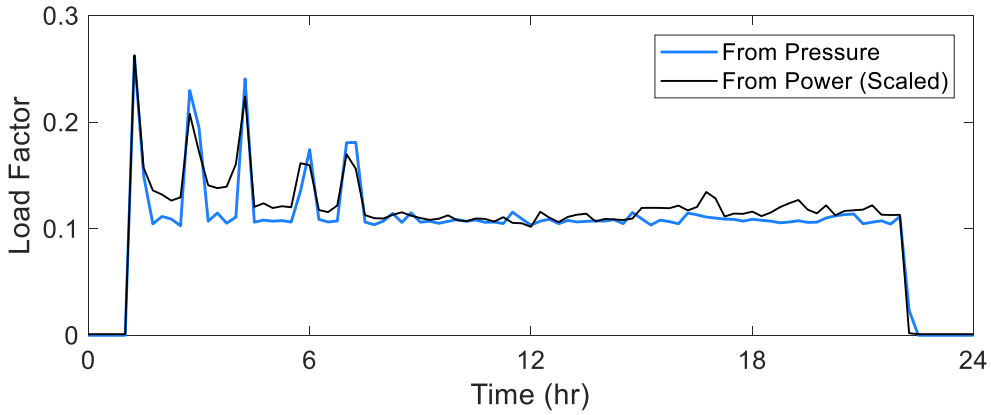

Figure 63. Evaluation of using pressure to estimate compressor loading for a load/unload operating scheme on a compressor in an industrial plant.

5.1.2.2 Load State Information

One of the primary purposes of the electrical load forecasting is to estimate the magnitude and time of the peak load to ensure adequate supply to meet electrical demand. The work presented in this chapter differs from electrical load forecasting as the peak load during a production day is of secondary interest. The primary objective is to determine periods of time in which production is idle and the compressor air system may be switched off. As such, the exact peak loading is not required, but rather the relative compressor loading throughout the day is of interest. One challenge in k-means load profiling of industrial data is that the relatively noisy profiles cause many profiles to be identified as outlier profiles. Many clustering algorithms give preference to these noisy data by assigning the outliers to unique clusters. The remaining production clusters are then grouped into a small number or even a single cluster. Once the load states are identified, the clustering algorithm is performed on profiles of these load states. The clustering algorithm becomes less sensitive to minor variations in daily usage and more dependent on general trends seen for each daily profile.

The number of necessary load states are determined from calculating the histogram of all points in the training dataset between zero and the maximum loading conditions. From this histogram, the local maxima indicate groups of similar loading conditions. These loading conditions regularly occur and indicate candidate load state locations. Peaks that include at least 5% of the total number of points are marked for potential candidate load state locations. Nearby peaks are combined into a single candidate load state location. If fewer than three load states are identified, additional

load state locations are added equidistant from the identified states so that at least three load states are identified.

Once the load state loading values are identified, each point in the training dataset is assigned to the nearest load state. After this initial assignment of load states, a low-pass filter is applied to the time-series dataset of load states to smooth the transitions between states.

5.1.2.3 K-means Estimation Algorithm

A novel k-means algorithm is used to classify the load profile and autonomously actuate the compressor. The training dataset is broken into days, where each profile is one day of operation. The k-means algorithm classifies each load profile in the training set into different clusters, which identifies different types of operating days. Once these clusters are identified, the test days are identified according to the clusters in order to identify times in which the compressor may be turned off.

As discussed in the literature review, many k-means clustering algorithms rely on a set number of identified clusters to assign profiles, and an incorrect number of clusters can result in large classification errors. If too many clusters are assigned, some clusters will contain nearly identical profiles. If too few cluster are used, either a significant number of clusters will be assigned as outliers, or a single cluster will contain profiles of distinctly difference shapes. In this work, the optimal number of clusters is not known *a priori*. Some work has been done on adaptively determining the number of clusters by allowing the number of clusters to grow or performing hierarchal trimming, but each of these algorithms require manual specification of parameters in each dataset.

The algorithm proposed in this work dynamically adds and removes clusters in two distinct phases during cluster assignment. The first phase allows the number of clusters to grow to accurately incorporate all profiles into unique clusters. A second trimming stage combines similar clusters and removes outlier profiles from the assigned clusters.

5.1.2.3.1 Growth stage:

After the load profiles are converted into profiles with defined load states, the algorithm starts with a user-defined number of clusters. This number is a starting point for the algorithm and does not constrain the number of clusters in the growth or trimming phase. In this work, four initial clusters are selected for all tests in simulation and real data sets. These initial clusters are selected by looking at the average load factor of each profile in the training set, similar to what is done in other work. The initial clusters are chosen such that the average loading is spaced equally among all the starting clusters. This does not guarantee that the starting profiles are unique or the final cluster centers, but rather serves as a simple metric to space out the starting clusters.

The growth stage is based on adaptive k-means clustering proposed by Bhatia [129]. In the referenced algorithm, the distance of the current profile to all the cluster centers is computer. Whenever this distance is further than the distance between the two closest cluster centers, the two closest clusters are merged and the profile becomes the center of a new cluster. However, this algorithm is sensitive to outlier profiles as an outlier profile will result in a new cluster formed specifically for the profile. Two existing cluster centers will be collapsed when creating this outlier cluster, thereby

creating non-optimal clusters. Ad-hoc measures are recommended in the paper to remove or account for these profiles.

In the growth stage, each profile is examined in the training set and the Euclidean distance to all cluster centers is computed. If the distance of the current profile to the nearest cluster center is closer than the minimum distance between cluster centers, then the current profile is assigned to that nearest cluster center. If the distance is further than the closest cluster centers, a new cluster is created with the current profile as the cluster center. Instead of collapsing two clusters to create this new cluster, an additional temporary cluster is created [129]. In this novel work, all clusters are retained and the number of clusters is allowed to grow. Note that the number of clusters after growth is largely dependent on the initial selection of cluster centers. If two of the initial clusters have a really small distance, then many new clusters will be created in this stage. Multiple iterations on the training set profiles are done until a steady state is reached and all of the profile have been assigned to their closest cluster center.

5.1.2.3.2 Trimming stage:

After all the cluster centers have been created, the clusters are trimmed. A merging factor, which is detailed in the following paragraph, is calculated for trimming. Once the merging factor is determined, the trimming stage operates as using a greedy algorithm that iterates through all the cluster centers in descending order according to current number of member profiles. The cluster with the largest number of member profiles first assigns all profiles within the merging factor to itself. The cluster with the largest number of remaining profiles after this first classification then assigns all

remaining profiles within the merging factor to itself. This continues until all the profiles have been assigned.

The merging factor can be user-specified or determined from characteristics of the training set profiles. In this work, the merging factor is determined by calculating the standard deviation of all training set profiles away from each cluster center. The minimum of these calculated standard deviation values determines the cluster center that is closest to all profiles if only a single cluster center was retained. This acts as a measure of the overall spread of the profiles from each other. The initial merging factor is calculated as half the minimum standard deviation of the profiles from the cluster centers. The final merging factor is determined as the minimum of the calculated merging factor and the minimum distance between cluster centers.

At the conclusion of the trimming phase, some clusters will have few or no remaining profiles remaining. Any cluster with less than a user-specified percentage of profiles is removed. In this work, the cutoff threshold is set at 4%. Any cluster with fewer than 4% of the total training set profiles is removed, as this corresponds to operating schedules which occur less frequently than approximately once a month. All of the member profiles in these removed clusters are labeled as outliers.

Once the clusters have been trimmed, the profiles are reassigned used a basic k-means cluster algorithm. This reassignment is done to ensure each profile is closest to its respective cluster center. The greedy nature of the trimming phase potentially assigns profiles to the largest cluster, even though a surviving cluster center may be closer to a particular profile.

5.1.2.4 Compressor Actuation

Once the training set is classified into respective clusters, each test day is evaluated to identify appropriate times when the compressor may be turned off. In turning off the compressor, the user specifies a buffer time window, which provides a buffer to ensure the compressor does not turn off immediately when reaching the idle state. Each test day is evaluated with the most recent 90 days as the specified training window. In other words, a rolling training window is used for the algorithm.

To determine periods of time to shut off the compressor, the load is observed until the idle load state is reached. Once the system has remained in the idle load state for the user-specified buffer time, the current profile up to that point in the day is compared to the cluster centers. Specifically, for each cluster, the distance of all member profiles at that instance from the cluster center is calculated. The maximum distance is recorded, as this represents the member profile that is furthest from the cluster center. The distance of the current day to each of the cluster centers is also calculated. For each cluster center in which the distance is less than the maximum observed distance for member profiles, the current profile is marked as a candidate fit. If all the candidate cluster centers are in the idle load state, the compressor is turned off. The compressor is turned back on at the earliest start time of the candidate cluster centers minus the user-specified buffer time.

If all the candidate clusters remain off for the remainder of the day, then the compressor is shut off for the rest of the day. The entire day profile is then classified according to the closest cluster center. If the error on the entire cluster is less than the

maximum error of all the respective member profiles, the current day is classified as a member of the cluster. If the current profile error exceeds the maximum error of member profiles, the current day is marked as an outlier, despite the cluster center being the closest match of all cluster centers. In this instance, the compressor is started at the end of the current day in anticipation of the next day.

If the profile is classified to one of the clusters, the next-day start time is estimated using both the day of week and the cluster type. The next-day start times of all member profiles are evaluated, and the initial minimum start time is set as the preliminary next -day start time based on member clusters. Additionally, as industrial schedules are typically heavily dependent on the day of week, the next-day start times of all profiles from the appropriate day of the week are determined. The start time is the minimum of the start time based on cluster assignment and the day of the week. The compressor is turned on at this evaluated start time minus the user-specified buffer time.

5.1.3 Compressed Air System Simulation

A physics-based dynamic simulation of a compressed air system was created in the Matlab/Simulink operating environment to test the intelligent actuation algorithm. This model was created using governing thermodynamic and fluid dynamic equations. The key elements created in the simulation include the compressor, receiver tanks, and pipelines. The goal of the simulation was to create realistic estimations of pressure drop, flow rates, and power usage for a specific plant layout.

5.1.3.1 Setup of Simulation

5.1.3.1.1 Plant Setup

Figure 64 shows the simulated plant that is modelled in Matlab/Simulink. This plant was based on a plant with a similar layout encountered during an IAC assessment. The compressed air system initially serviced two production lines, the main production line on Line 1 and the customs area. The main production line is serviced by a large, 3-inch compressed air line. The customs area is serviced by a smaller, 2-inch line due to the lower initial demand. As the plant expanded, an additional production line was added of the same line as the customs area. The smaller line creates an undersized line with a measurable pressure drop from the outlet of the receiver tank to the production lines.

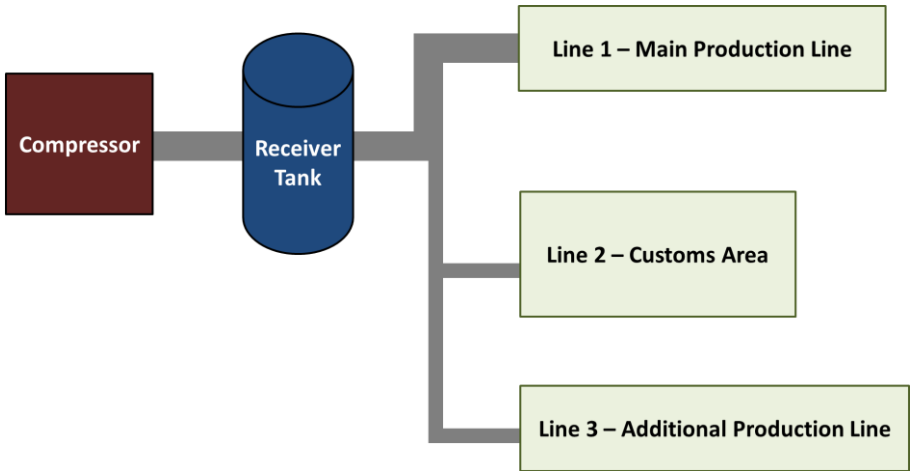

Figure 64. Plant layout used in the compressed air simulation.

5.1.3.1.2 Plant Schedule

The schedules for the different operating lines were created based on the maximum operating flow for each line. The approximate operating schedule of the actual

plant from Figure 64 was used as a basis for the operating schedules of the simulation. The two production lines exhibit relatively constant usage, and the customs line has more variation in usage.

Specifically, the plant nominally operates 5 days a week and on some Saturdays, with some reduced usage on Saturday. The second production line operates on a reduced schedule of approximately three or four days each week. There are eight holidays when production is shut down, and there is slow overall growth in energy over the course of the year. The different lines have different start and stop times, and there is random variation added to these times that account for variation in actual production schedules. In addition to the relatively constant usage in the production line, there is highly variable usage added to the production line schedules to account for air tool usage. An example of a one-week schedule for each of the operating lines is shown in Figure 65.

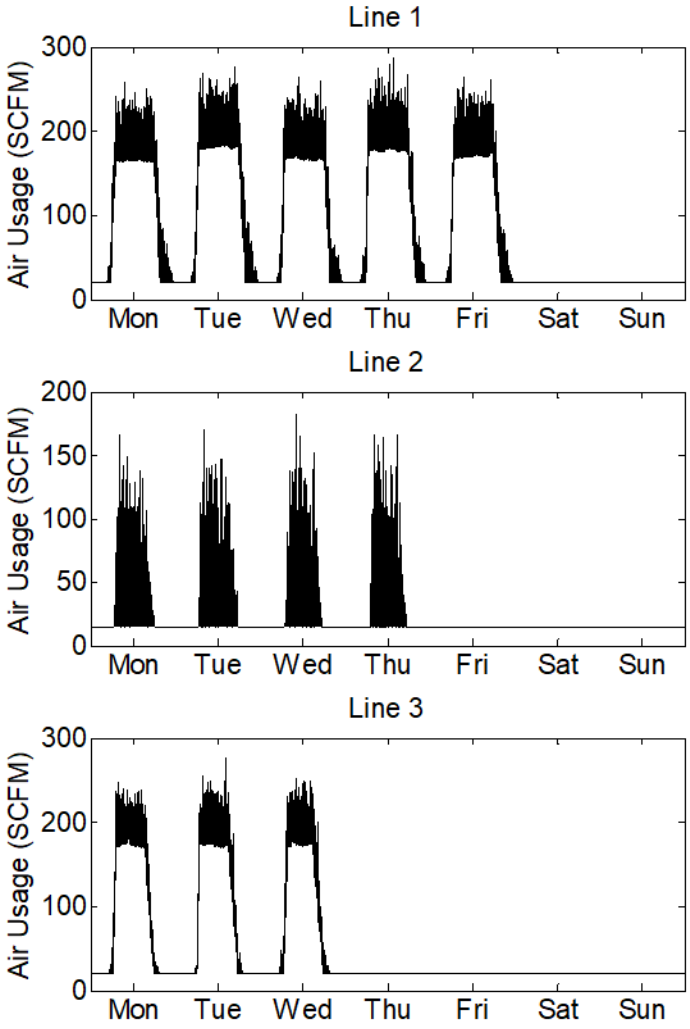

Figure 65. Typical weekly air usage for each of the three operating lines in the simulation.

5.1.3.1.3 Physics Models

5.1.3.1.3.1 Compressor Model

The modelled compressor was based heavily on a thesis from Chappell [130]. The purpose of the work of the thesis was to create a sound physics and thermodynamic model to accurately model a compressor for use in energy assessments. The modelling of the compressor is heavily influenced by the operating scheme and type of compressor

(reciprocator, scroll, etc.). A load-unload, oil-flooded screw compressor was chosen to model for the simulation. A load-unload operating scheme of the compressor was chosen in order to test the ability to convert pressure load-unload data to approximate load factors, and a rotary screw compressor was chosen because of their common use in manufacturing plants.

The basic governing equations for the simulation are derived based on theoretical equation with empirical correlations to account for observed phenomena. The theoretical rated volumetric flow through a compressor is a function of the compressor motor power, compressor motor efficiency, and the specific motor efficiency as seen in Equation (59).

(59) $$\dot{V}_{rated} = \frac{(W_{motor})\left(0.746\frac{kW}{hp}\right)}{(\eta_{motor})(\eta_{comp})}$$

This rated flow is given in standard cubic feet per minute, where the mass flow can be determined by multiplying the density of air at standard conditions. The actual volumetric flow is the product of the rated flow and an empirically-derived volumetric efficiency. This volumetric efficiency is defined in Equation (60). As the pressure increases, the mass output of the compressor will decrease slightly from the nominal rated value.

(60) $$n_{ov} = -0.0452P_{comp} + 91.932$$

The work required of the compressor to compress the air is a function of both the isentropic work with its associated isentropic efficiency and the additional work to overcome frictional forces on the compressor. The isentropic work scales with operating pressure as shown in Equation (61).

(61) $$\dot{W}_{compression} = \frac{\dot{V}_{actual}\rho_{atm}c_p T_{atm}}{\eta_{is}}\left[\left(\frac{P}{P_{atm}}\right)^{\frac{k_c-1}{k_c}} - 1\right]$$

The frictional work is assumed to be constant, and it is estimated in the thesis to consume approximately 10 percent of the rated power of the compressor as seen in Equation (62).

(62) $$\dot{W}_{friction} = (0.10)\frac{(W_{motor})\left(0.746\frac{kW}{hp}\right)}{(\eta_{motor})}$$

The values used in the compressed air simulation are summarized in Table 16. The values for air are all taken at standard conditions. The pressure load and unload setpoints were set at 105 and 115 psi, respectively. These pressure setpoints were defined to meet minimal operating pressures at each of the lines with an additional small pressure buffer.

5.1.3.1.3.2 Receiver Tank

The receiver tank is assumed to be at uniform pressure by assuming that sudden pressure changes in the upstream and downstream systems quickly equilibrate within the tank. Additionally, the temperature of the air inside the tank is assumed to be constant with time. The pressure in the tank is a function of the mass flows into and out of the tank as seen in Equation (63). The mass flow rates into and out of the tank are determined based on compressor output and downstream uses of the compressed air. The size of receiver tank used in the simulation is shown in Table 16.

(63) $$\frac{dP}{dt} = \frac{RT}{V}\left(\sum \dot{m}_{inlet} - \sum \dot{m}_{outlet}\right)$$

5.1.3.1.3.3 Pipelines

The pipelines are assumed to contain negligible volume, so the outlet pressure of the pipe is determined from the inlet pressure and the pressure drop in the pipe. The pressure loss of the pipelines is calculated based on the Darcy-Weisbach equation [131]. Additionally, the pressure loss through pipe bends and other obstructions is also calculated for a user-defined equivalent number of 90° bends. The pressure loss is given as a function of the pipe dimensions, friction factor, speed of the air in the pipe, and the density of the air. As mass flow is calculated for the compressor and used in the simulation, the velocity term in the equation is substituted with the mass flow rate. Additionally, the density of air is substituted using the ideal gas equation, which consists

of the temperature, pressure, and specific gas constant. When these terms are introduced and simplified, the pressure loss in the pipeline is given in Equation (64).

$$P_{downstream} = \left(P_{upstream} - \frac{\dot{m}^2}{P_{upstream} D_i^4} \left(7.340 * 10^{-8} \frac{Lf_D}{D_i} + 3.670 * 10^{-8} \right) \right) \quad (64)$$

The friction factor for rough cast iron pipes was used in the simulation, and the length and diameter of the pipes were sized based on actual industrial plants. The mass flow rate is a dependent parameter that depends on conditions upstream and downstream of the pipe.

5.1.3.1.3.4 End Uses

To determine the mass flow rate of end uses, the end uses are approximated as orifices of equivalent diameter. This captures the reduced flow rate at lower pressure as seen in typical end uses of compressed air. The input parameter is the flow rate at 100 psig, and the orifice is sized to meet the flow rate. The orifice has a default coefficient of flow of 1.0. Using these assumptions, the mass flow for the end uses are calculated using Equation (65) with the equivalent orifice size [131].

$$\dot{m} = \frac{12.22 D_{orifice} P}{\sqrt{T}} \quad (65)$$

Table 16. Summary of Parameter Values for the Compressed Air Simulation

Parameter	Units	Value
W_{motor}	hp	125
η_{motor}	$-$	0.95
η_{comp}	$\frac{kW}{SCFM}$	0.171
η_{is}	$-$	0.90
$\eta_{reciever}$	m^3	11.36
T_{atm}	K	311

5.1.3.1.4 Connection of Subsystems

Figure 66 shows the Simulink diagram for the system depicted in Figure 64. Each of the major subsystems on the diagram has a colored background and is labelled on the figure. One challenge with creating Simulink diagrams is the limitation that each subsystem may contain a single output. As such, the subsystems must be connected in a specific way to ensure each subsystem has exactly one output. In order to accomplish this, the defined outputs and inputs for each system is shown in Table 17. Future expanded simulations using these connection types will ensure that the simulation is linked properly.

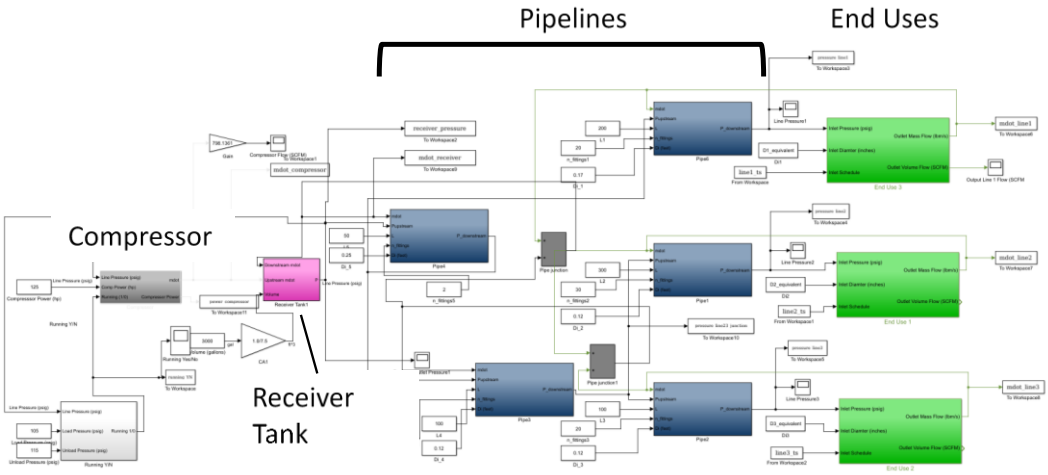

Figure 66. Simulink diagram to create compressed air simulation.

Table 17. Defined Inputs and Outputs for Simulation Subsystems

Subsystem	Defined Output	Defined Input
Compressor	Mass Flow	System Pressure
Receiver Tank	Tank Pressure	Input and Output Mass Flow
Pipelines	Downstream Pressure	Upstream pressure, Mass Flow
End Uses	Mass Flow	Inlet Pressure

5.1.3.2 Simulation Outputs

The simulation was run with time steps of one second for a period of one year. The simulation provides information on the pressure at every point in the system, the compressor power, and the mass flow rates for all subsystems. As full state information is not available in actual industrial plants, only the pressure at the outlet of the receiver tank and the power of the compressor were measured for use in the algorithm to learn the schedules. Figure 67 shows the total system mass flow for one operational day.

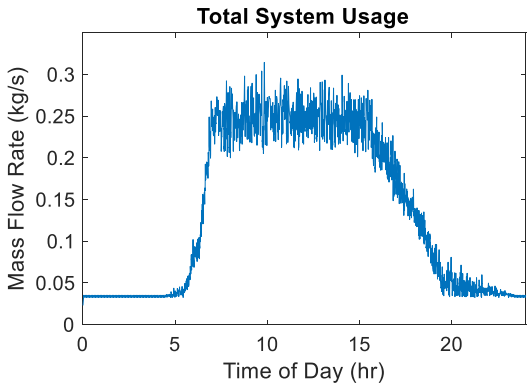

Figure 67. Compressed air usage in the simulated facility over the course of one day.

Figure 68 shows selected times of pressure and flow data. The figure on the left highlights how the undersized line creates a measurable pressure drop downstream of the compressor. The compressor setpoints must be raised to compensate and ensure the production line receives the correct system pressure. The figure on the right shows the facility compressed air usage as it ramps down during a production cycle. The pressure at the receiver tank is shown as well, with noticeable changes in the load-unload cycles of the compressor. At low loading conditions, the pressure ramps up quickly and decays slowly, while the reverse is true during periods of heavy loading. Both of these conditions result in extended cycle times, while periods with moderate loading have shorter compressor cycles.

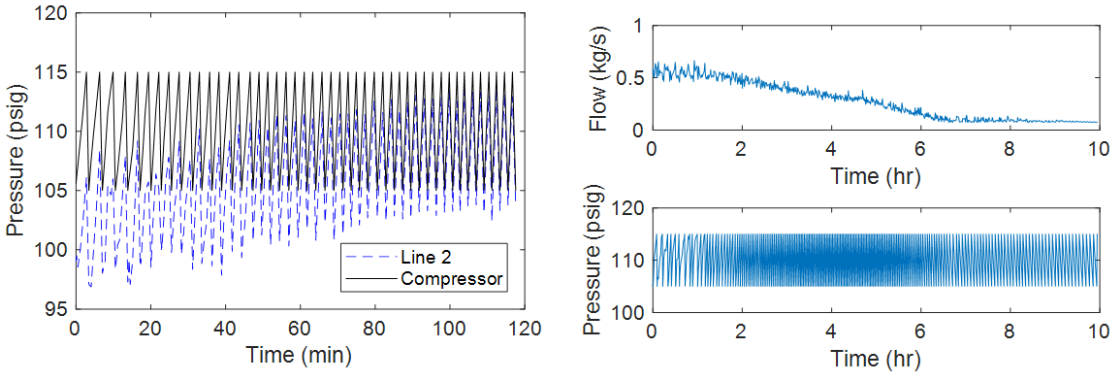

Figure 68. Pressure data at selected points from the simulation. The left figure shows the noticeable pressure drop from the undersized line, and the right figure shows the differences in load-unload cycles as the flow rate changes in the facility.

5.1.3.3 Simulation Results for the Intelligent Actuation Algorithm

The clustering algorithm was tested on the simulation dataset prior to testing on experimental data. The simulation described above was tested using the algorithm presented earlier in the chapter. A total of 360 days of operation were simulated. An initial training set of 90 days was used, and the remaining 270 days were used as the test days. A rolling training window of the most recent 90 days was used for the training set for each individual test day.

5.1.3.3.1 Load State Information

The histogram of the loading states is shown in Figure 69 for the first 90 training days. There are three distinct peaks in the different loading conditions. The highest peak with a load factor around 0.2 is the idle load state when no production occurs. The other two peaks, while lower in maximum intensity, have a wider tail that reflects the greater variability of usage during production. The main production line has a wider range of hours, and times when the main production line is the dominant usage is seen in the

second peak. The third peak comes from when both main production lines are operating near full capacity. The loading states between the peaks do not go completely to zero as there are transitional states or periods of time when the sporadic usage in the Customs Area is larger. The three identified load states are marked on the figure for reference.

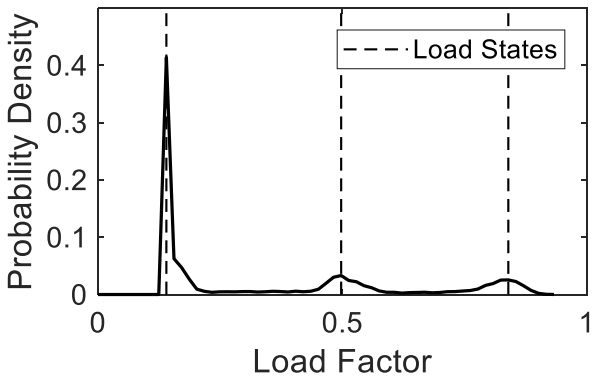

Figure 69. Histogram of loading conditions for the simulated compressed air system.

Figure 70 shows the daily load profiles for 30 days of the training dataset. The top plot shows the original compressor loading values. The bottom plot shows the loading values translated into load states with the lowpass filter applied. Translating the loading conditions into load states maintains the trends and general structure of the daily profiles, but it reduces the noise in the signal used for clustering.

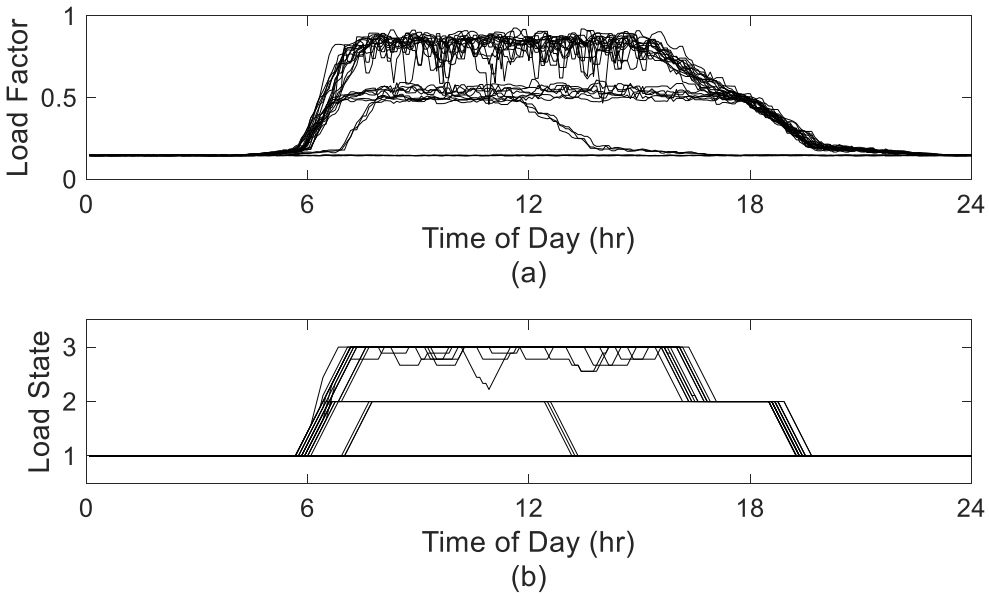

Figure 70. Original simulation compressor loading is seen in (a) for 30 training days. The load states of the same profiles are shown in (b).

5.1.3.3.2 K-means Estimation Algorithm

Figure 71 displays the initial clusters used on one day of the clustering algorithm. These clusters were chosen based on equidistant average loading conditions for all the training days. They are intended only as starting points in the algorithm. This is evident in the third and fourth clusters, which are similar in shape and total loading.

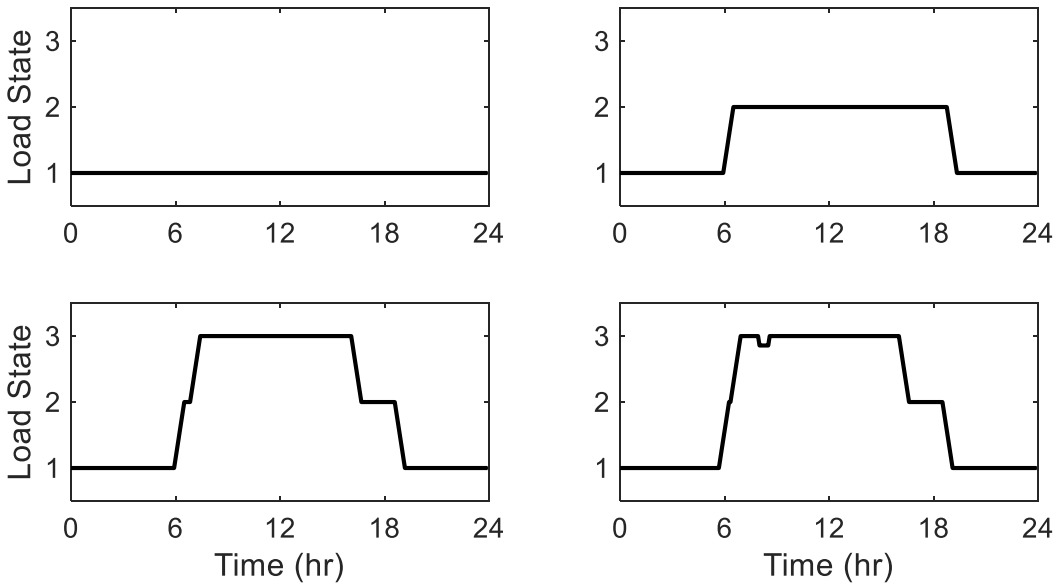

Figure 71. Initial clusters in the k-means clustering algorithm for the simulation data.

All of the clusters after the growth stage are shown in Figure 72 for this particular test day. In this instance, there were nine clusters after growth. Throughout the 270 test days in the simulation, the number of clusters after growth varied from 5 to 54. As mentioned previously, this number is highly dependent on the initial choice of clusters. Several of these clusters are very similar in shape and overall loading.

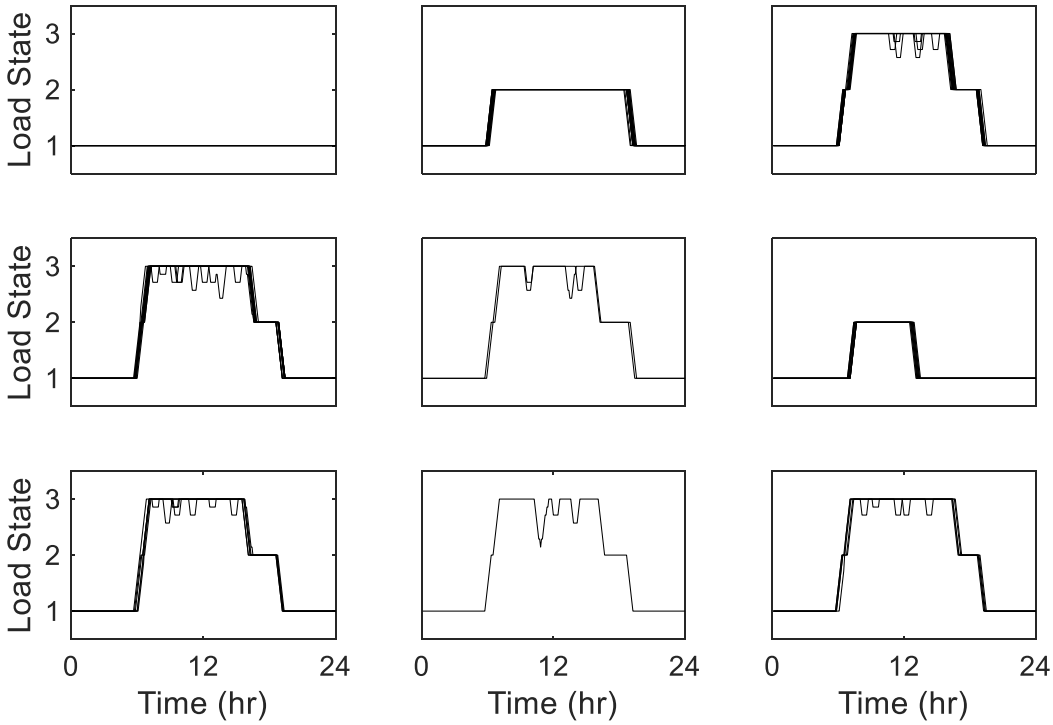

Figure 72. All of the clusters and the associated profiles after the growth stage for the simulation data.

The final identified clusters after the trimming stage are displayed in Figure 73. Four distinctly different profile types are identified, and all of the profiles in the training set are associated with one of the clusters. When contrasting these final clusters to the initial clusters, the third and fourth initial clusters were collapsed into the same clusters and Cluster 4 in the final clusters was a newly identified cluster type. For each of the 270 test days, there were either 3 or 4 identified clusters after the trimming stage. While the number of clusters after the growth stage is largely dependent on the initial selection of clusters, the final number of clusters is much more consistent between test days. These

clusters and associated member profiles are notably clean and well defined in the simulation, in large part due to the limited noise and uncertainty of simulation data.

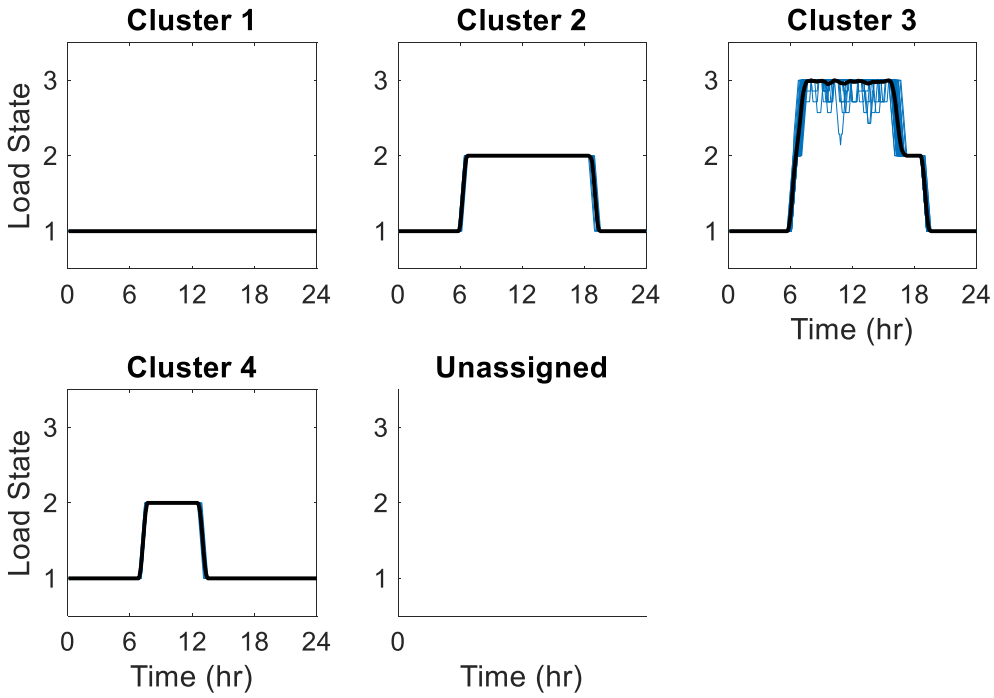

Figure 73. Final clusters after trimming stage for the simulation data.

5.1.3.3.3 Compressor Actuation

The original loading profile of the compressor and actuated load profile are compared in Figure 74. In the simulation, the actuation of the compressor always occurred within times when production was off, in part because of the conservative nature of the algorithm. The algorithm achieved an 8.3% reduction in energy usage of the compressor. If the compressor was turned off always when in the idle state, a 27.3%

reduction in energy would be possible. The algorithm achieved approximately one third of potential energy savings.

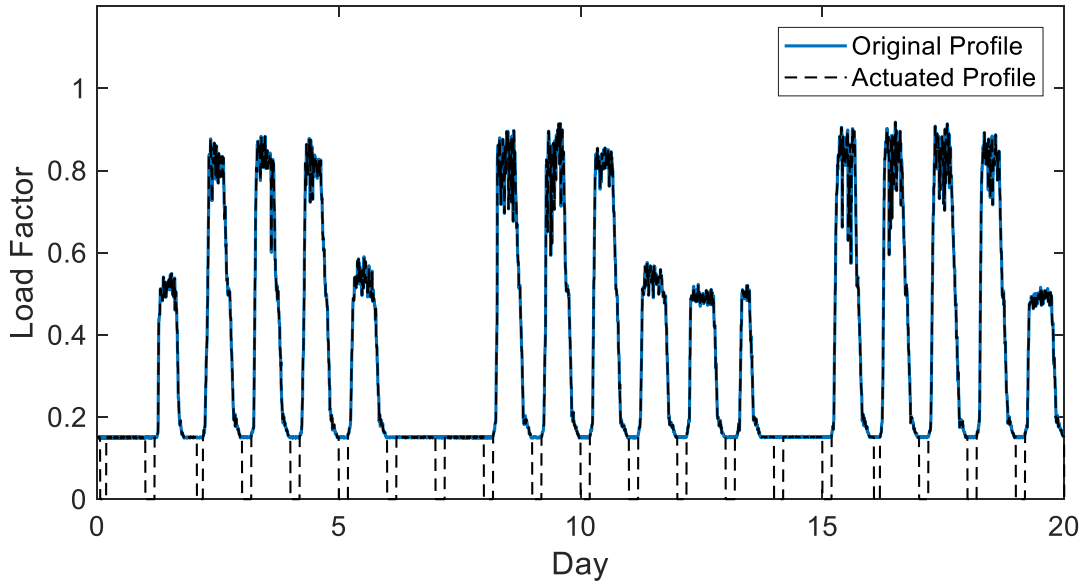

Figure 74. Comparison of original and actuated compressor load profile for approximately 3 weeks in the simulation data.

5.1.4 Experimental Verification using 15-Minute Demand Data Sets

As data collection for the evaluation of research algorithms is expensive, the k-means algorithm presented in this chapter was tested initially on several datasets of 15-minute demand data. These datasets record the average demand of a facility over each 15-minute period over the course of data collection. A total of 12 datasets were tested, each with operating characteristics similar to a manufacturing plant with either a one- or two-shift operation. While the data for the entire facility from these datasets will be

notably different that then compressor loading, they provide similar overall usage patterns along with realistic noise and process variation from a physical facility.

One year of data were available for 10 of the sets, and the other two had approximately 120 days of data. The time series data of each of the complete datasets are shown in Figure 75. While all show periods of idle and production usage, there is variability between overall patterns of usage in the datasets to test the robustness of the k-means algorithm. Figure 76 shows the same 12 datasets, but this figure illustrates the data over a 30-day window. This figure highlights the differences in daily use and operating pattern between datasets. The algorithm was tested on all of the datasets, but the figures shown in the following sections are representative from a single test day from one of the datasets.

For each dataset, a 90-day rolling window was used for the training set, and the remaining days in each dataset were used as test days with the actuated compressor. The results are summarized in Table 18 below. Note that since there is a 90-day training window with a dataset of an entire year, later test days will have completely different training sets than earlier test days. As can be seen in Figure 75, there is considerable variation in usage throughout the year which can result in considerable variation in the identified clusters and times for actuation.

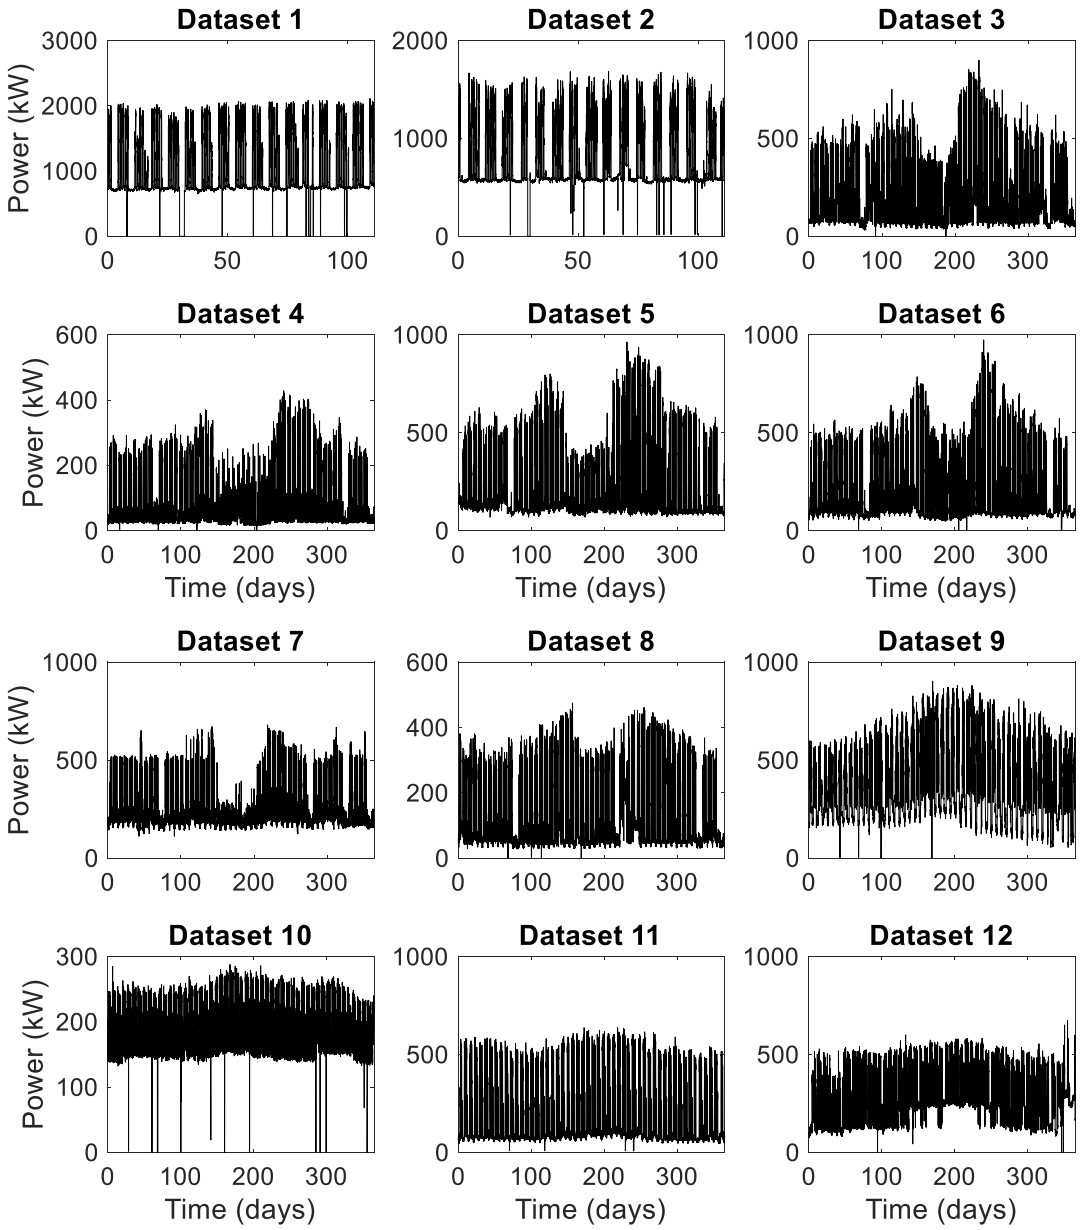

Figure 75. Complete time-series data for the 12 datasets of 15-minute demand data used to test the k-means algorithm.

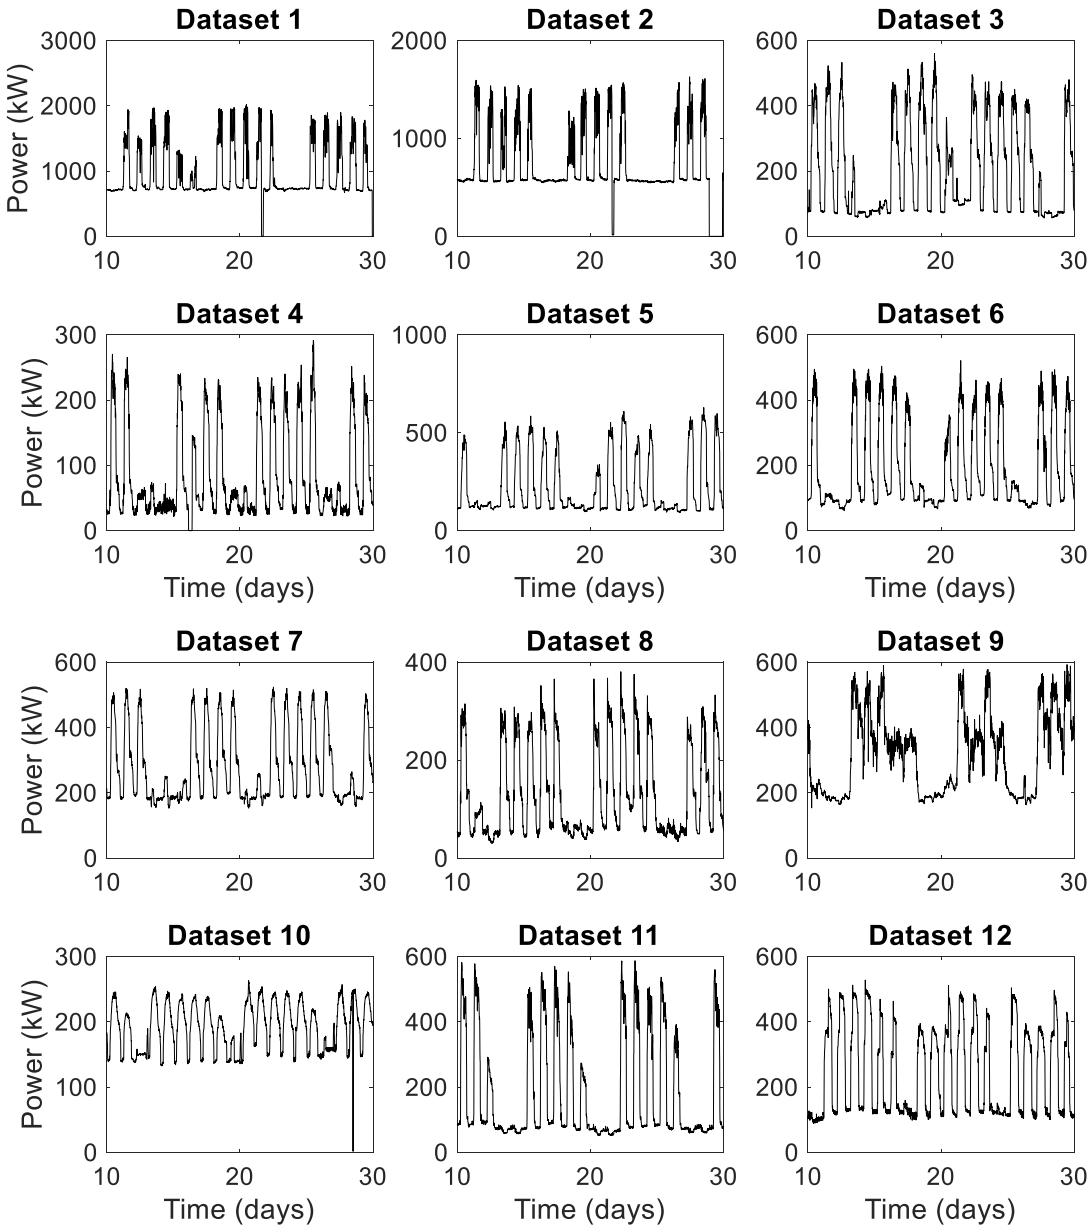

Figure 76. A selected 30-day window for each of the 12 datasets to show the variation in production usage for each of the 15-minute demand datasets.

5.1.4.1 Load State Information

As the available data provides facility demand instead of relative compressor loading, the load states were determined based on the facility demand. An example of the identified load states for one test day in one of the datasets is shown in Figure 77. As expected, the identified load states are much less distinct than in simulation. The most distinct peak occurs in the idle load state, and the remaining usage is much more dispersed. Although nuanced, there are three small peaks in the three highest load states. Table 18 shows the range of identified load states for each of the tested datasets.

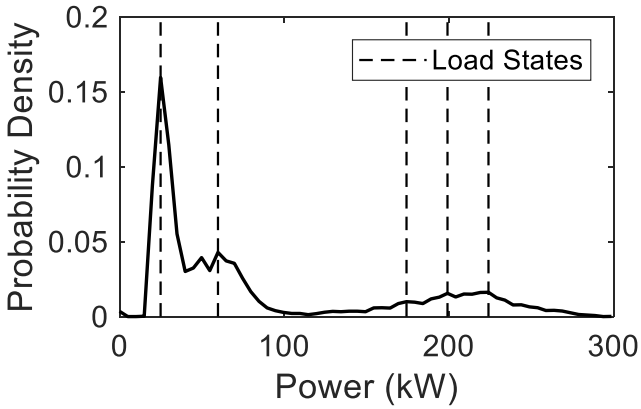

Figure 77. Example load state identified for one of the 15-minute demand datasets.

Figure 78 shows the translated load states from Figure 77 for 30 daily profiles. Similar to simulation, translating the initial data to load states maintains overall usage patterns while reducing noise in the signal.

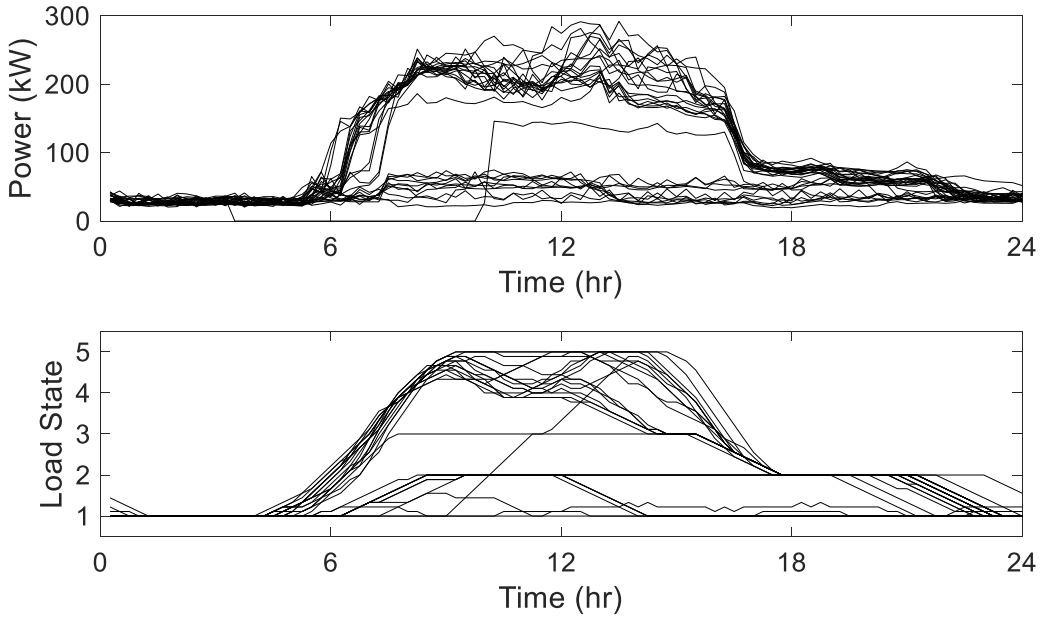

Figure 78. Power data translated to load states for 30 days in one of the 15-minute demand datasets.

5.1.4.2 K-means Estimation Algorithm

The initial clusters are shown in Figure 79 based on equidistant spacing of average daily loading. All five load states are found in the initial cluster profiles, but the third and fourth profiles are similar in overall shape.

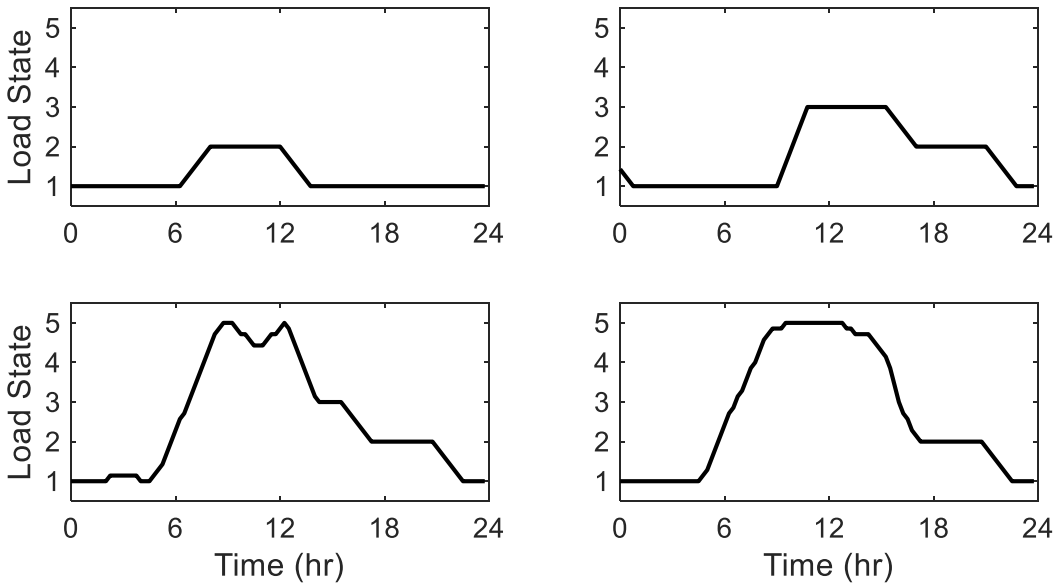

Figure 79. Initial clusters for one of the 15-minute demand datasets.

After growth for this particular test day, a total of 12 clusters were created and are shown in Figure 80. The clusters after the trimming phase are shown in Figure 81. With the increased variability and noise in the dataset, some profiles fall outside the limits for every cluster and are classified as outliers. On this particular test day, four final clusters were identified with distinct operating profiles. Of note is the difference in final cluster centers from the initially identified cluster centers. Table 18 shows the minimum and maximum number of clusters after both the growth and trimming stage. As noted earlier, initial selection of clusters that are close together causes an explosion in the number of total clusters after growth, but the number of clusters is greatly reduced through the trimming phase.

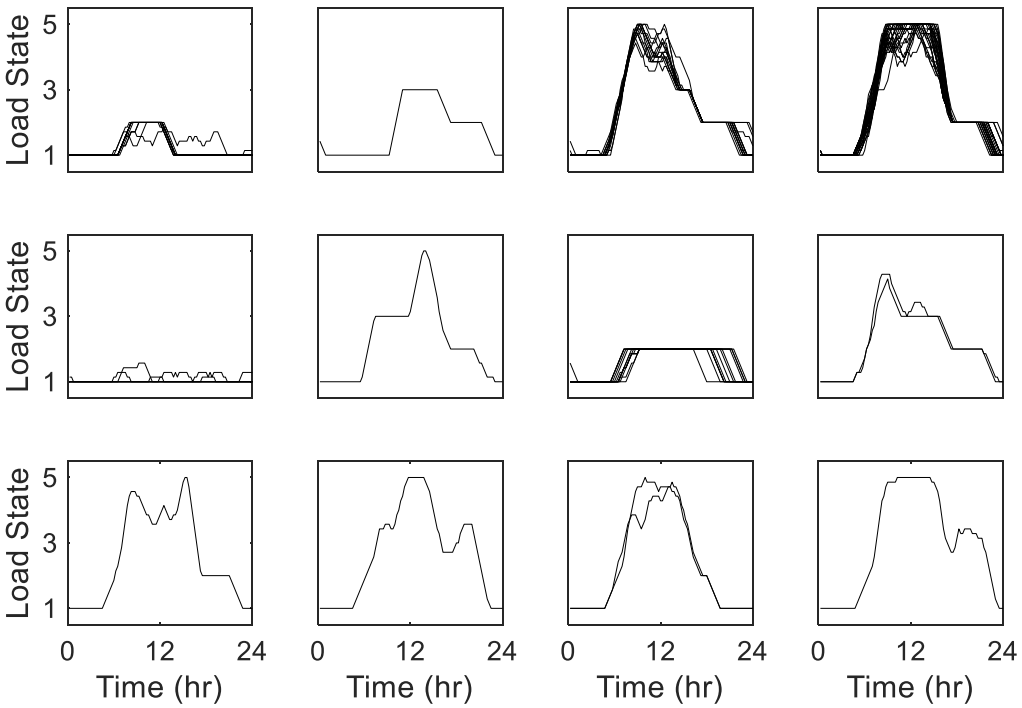

Figure 80. Clusters after growth for one of the 15-minute demand datasets.

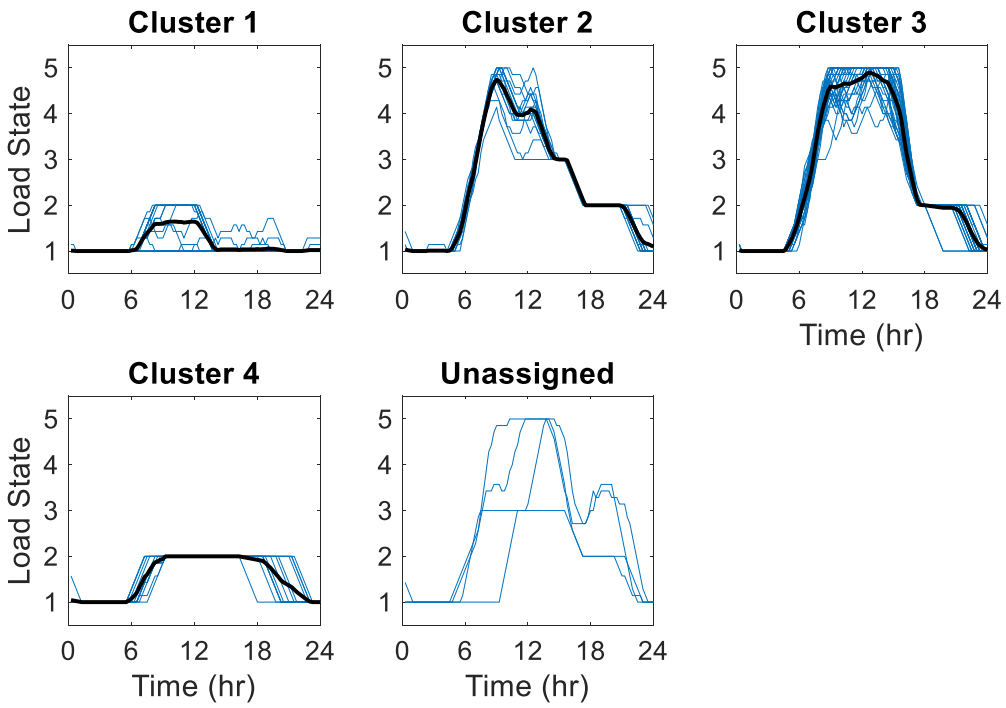

Figure 81. Final clusters after growth for one of the 15-minute demand datasets.

5.1.4.3 Compressor Actuation

The actuated profile for one of the datasets is displayed in Figure 82, and a summary of all the actuated datasets is shown in Table 18. This table shows the actual percent savings that were achieved in each of the datasets, which ranges from 1% to 14% of the total energy use of the facility. This metric partially depends on the facility operating characteristics because not all energy may be reduced in the facility. The distinctly different overall usage patterns between datasets causes differences in the total possible reduction in usage based on both the load during idle periods and the time spent in idle periods. As such, the maximum possible reduction in energy usage for each dataset is also shown in the table. The maximum reduction was calculated by reducing

the demand to zero at all times in the idle load state. Also shown in the table is the ratio of the actual achieved savings to the maximum possible savings. This ratio ranges from 11% to 36% and indicate potential improvements to the k-means algorithm.

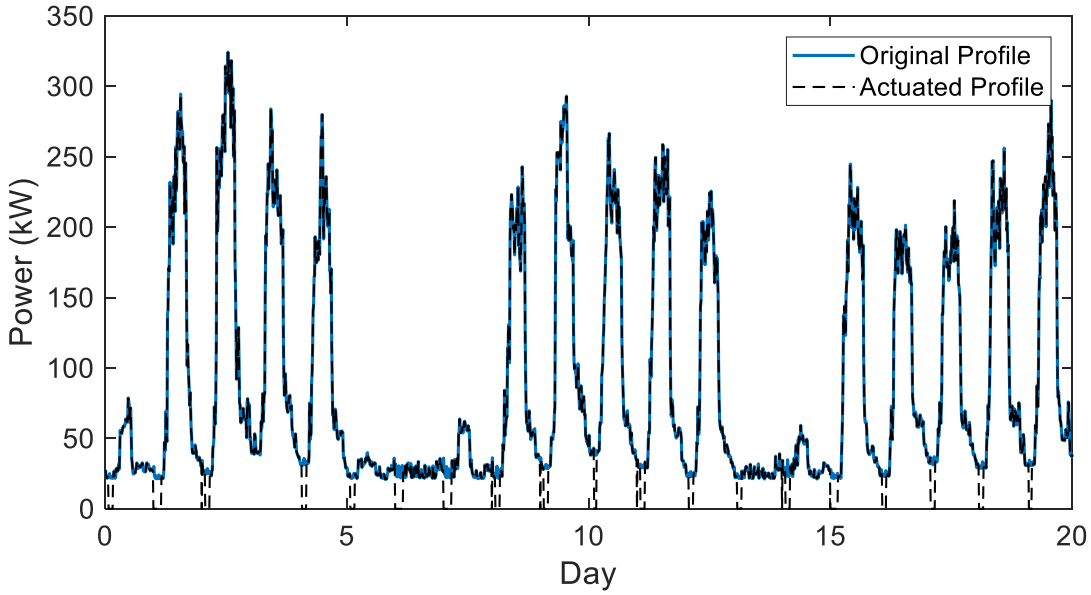

Figure 82. Comparison of the original and actuated profile for one of the 15-minute demand datasets.

Table 18. Summary of Results for the 15-minute Demand Datasets

Dataset Number	Number of Test Days	Annual Savings	Maximum Possible Savings	Second Law Savings	Min Identified Load States	Max Identified Load States	Min Clusters After Growth	Max Clusters After Growth	Min Clusters After Trimming	Max Clusters After Trimming	Number of Type 1 Error Instances	Number of Type 2 Error Instances
1	20	13%	53%	25%	3	3	8	59	3	5	0	0
2	20	14%	58%	23%	3	3	30	55	3	4	0	0
3	270	2%	11%	17%	3	5	4	64	3	7	9	0
4	270	5%	17%	29%	3	5	4	78	2	8	7	0
5	270	6%	22%	25%	3	5	4	48	2	9	2	0
6	270	4%	19%	23%	3	5	5	60	2	6	1	0
7	270	13%	36%	36%	3	5	4	75	2	8	4	0
8	270	3%	15%	21%	3	5	4	65	2	6	2	0
9	270	1%	12%	11%	3	5	4	74	3	8	3	0
10	270	4%	20%	21%	3	4	9	73	3	7	7	0
11	270	6%	20%	30%	3	5	4	58	2	6	1	0
12	270	3%	12%	23%	3	5	4	63	3	7	3	0

5.1.4.4 Errors in Actuation

In evaluating the effectiveness of the k-means algorithm, errors in actuation were evaluated for each dataset. The errors are broken into Type 1 and Type 2 errors. Type 2 errors are errors in which the system is turned off in the non-idle operating state. Due to the conservative nature of the algorithm and the ability to use measured system data to determine when to turn off the system, no Type 2 errors were observed in any dataset. Type 1 errors are errors in which the system is not turned prior to the start of production. This is the more likely error as it requires prediction using past performance data, which is subject to additional uncertainty. This error typically is not catastrophic to production, but instead is inconvenient in requiring plant personnel to manually turn on equipment. The maximum number of observed Type 1 errors in a dataset was 9 errors, which corresponds to approximately 3% of the test days. Overall, there are minimal occurrences of this type of error.

Several examples of this error are illustrated in Figure 83. The two plots on top show errors from potential misclassification of load states. This nominally higher increased use could potentially still be in a non-production state and have minimal negligible impact from the error. In general, this misclassification of idle load states as operating load states adds an additional conservative layer to the algorithm. The bottom two plots show errors when the start of production was incorrectly estimated. As mentioned previously, typically the consequence of this incorrect estimation is inconvenience to facility workers in manually actuating the system.

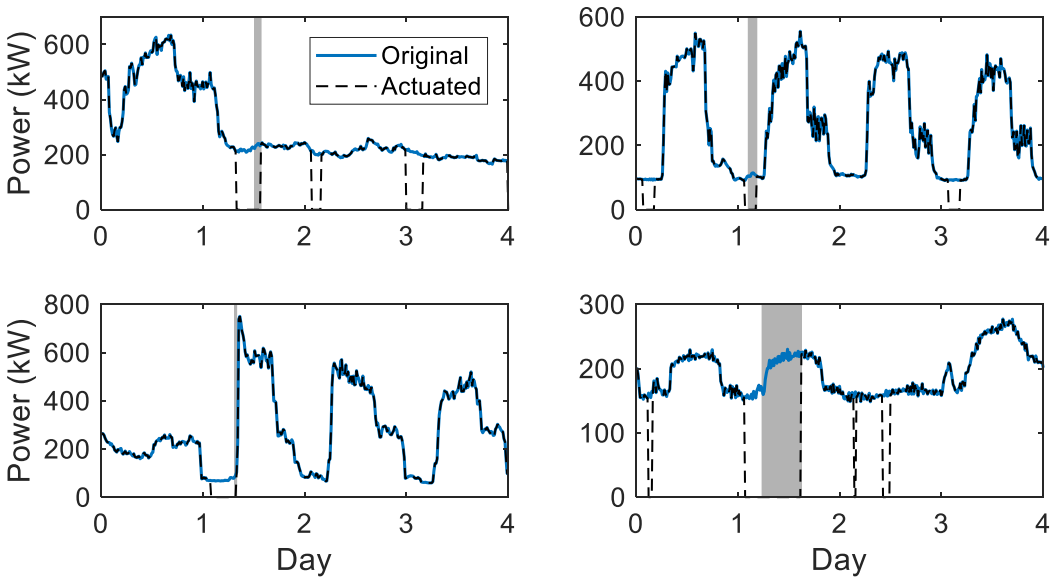

Figure 83. Illustrative examples of Type 1 actuation errors. The periods of time when the system was actuated off the but actual load state was in a production load state are highlighted in gray.

5.1.4.5 Variation in User-Specified Buffer Time

The number of Type 1 errors in the dataset is correlated to the user-specified buffer time in the k-means algorithm. The system will turn on earlier with increased buffer time. While this will reduce the number of Type 1 errors, the total energy savings will be decreased. Figure 84 shows the effect of changing the buffer time on both the number of Type 1 error instances and the system energy savings. Each dataset is one of the lines on the graph. For reference, all of the figures for the 15-minute demand datasets and results in Table 18 were done with a buffer time of 1.5 hours. The number of errors steeply drops off with increased buffer time, but the realized energy savings drops off more gradually.

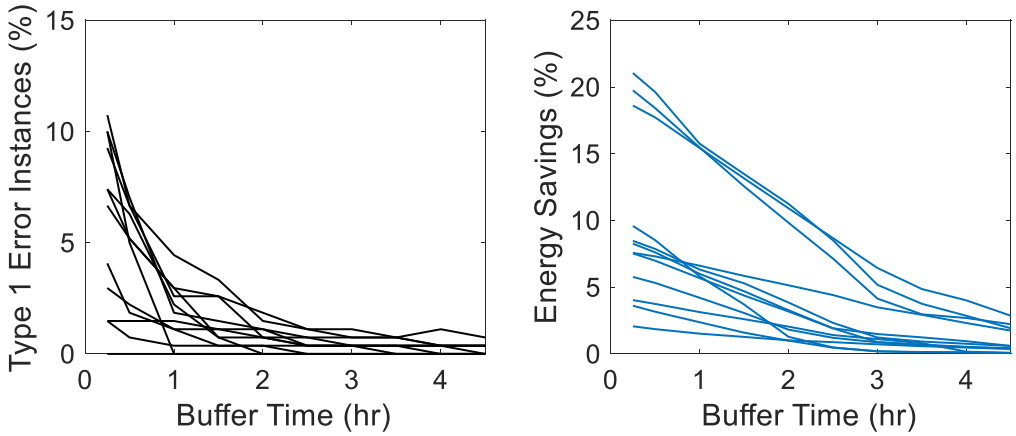

Figure 84. Effect of the use-specified buffer time on Type 1 error occurrences and energy savings for the 15-minute demand datasets.

5.1.5 Experimental Verification using Sensed Data

Data were collected on an industrial compressed air system in a facility that fabricates steel. The facility manufactures customs assemblies and relies on compressed air for several steps in the production process. The facility operates between one and two shifts for five days each week. The only constant usage during production comes from air leaks in the lines, and all other usage in the production process is sporadic due to the custom nature of the manufactured product.

The compressed air system was manually turned off at the end of each production day by production personnel. The energy savings achieved from the k-means algorithm is compared to the energy savings from manual actuation. As mentioned earlier in the chapter, there are several reasons that inhibit manually turning off the compressor in most plants, but this comparison highlights the current performance of the algorithm and future possible improvements.

Figure 85 shows the compress air usage over the entire time window of data collection on the left and a shortened window on the right. As expected, the usage goes to zero at nights and non-production day. The usage is relatively constant from leaks throughout each production day with visible spikes during periods of air usage.

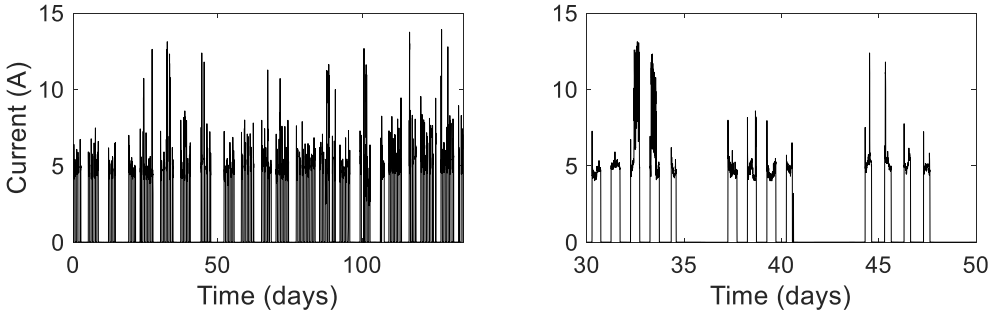

Figure 85. Current from the industrial compressed air system over the course of data collection.

When evaluating the autonomous actuation algorithm, there will be periods of time in which the compressor should still be on according to the algorithm but the actual compressor was off. In these instances, the idle compressor power is estimated to account for the compressor power. The average power of the compressor is computed over all times in which the compressor is in the idle load state, and this average value is used in all instances with zero measured power usage.

Data were collected for a total of 134 days at the facility. Similar to simulation and the 15-minute demand datasets, 90 training days were used for the initial training set. The remaining 44 days were used for test days. The results of the k-means algorithm

are presented in Table 19. These results are discussed in the following sections for the different stages of the k-means algorithm.

Table 19. Test Results of the k-mean Algorithm on the Measured Data from an Industrial Manufacturer.

Parameter	Value
Number of Test Days	44
Savings: Autonomous Control	19%
Savings: Human Control	70%
Maximum Savings	94%
Ratio of Savings: Autonomous Control	20%
Ratio of Savings: Human Control	74%
Min Identified Load States	3
Max Identified Load States	3
Min Clusters After Growth	17
Max Clusters After Growth	36
Min Clusters After Trimming	1
Max Clusters After Trimming	1
Number of Type 1 Error Instances	1
Number of Type 2 Error Instances	0

5.1.5.1 Load Factor Estimation

Both power and pressure data were monitored for the dataset. The facility changed the operating conditions of the compressor from load/unload operating scheme to a VFD configuration midway through the data collection period. The period of time with the load/unload operating scheme was used to evaluate the use of pressure data to estimated compressor loading. The methodology of this estimation was discussed previously, and this dataset is used in the creation in Figure 63.

5.1.5.2 Load State Information

The load states are calculated in the same manner as previous datasets for all periods of time when the compressor was on. At times when the compressor was manually turned off, the load state of those times is set to the first load state, which corresponds to the idle operating state. In every test day, as seen in Table 19, three load states were identified and used. Figure 86 shows the comparison of the compressor current and translated load states for 30 of the training days. When production is running, the compressor remains in the idle operating state for large periods of time. The relatively sporadic usage of the compressed air system is captured in the load states, and the pattern of usage is much more irregular that was seen in both simulation and the 15-minute demand datasets.

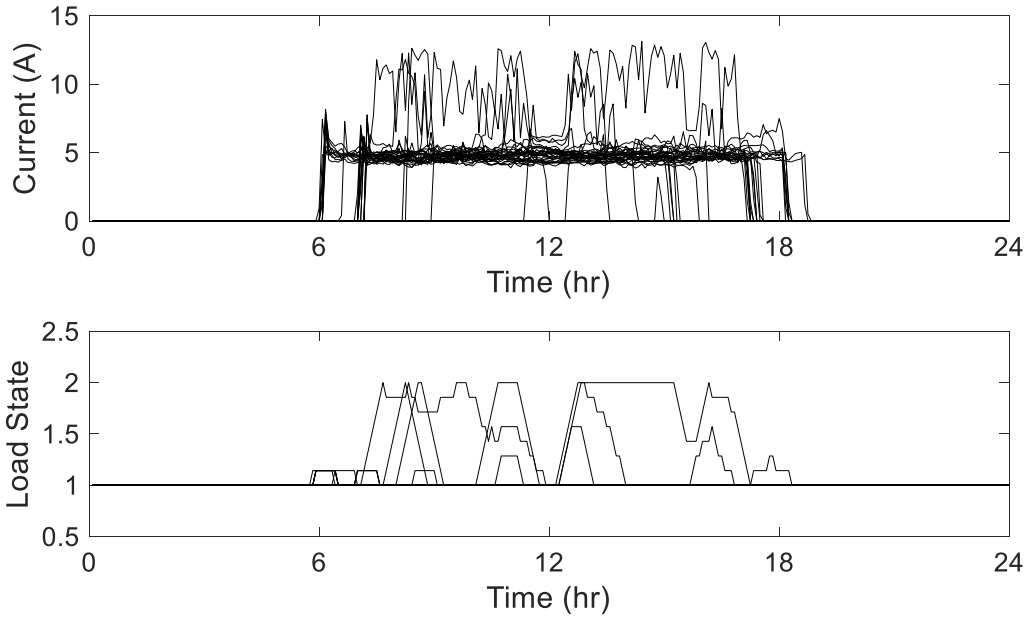

Figure 86. Comparison of daily load profiles based on compressor current and translated compressor load states.

5.1.5.3 K-means Estimation Algorithm

The results of the k-means algorithm are summarized in Table 19. Due to the sporadic nature of the load profile, there are many days with unique time periods of heavy air use. As such, there are many identified clusters after the growth stage for each of the test days. However, the algorithm always consolidated to a single cluster after the trimming phase due to the lack of uniformity in periods of high use. A single cluster after trimming results in more conservative actuation, as the uncertainty from all the profiles is captured and accounted for in the actuation. Each cluster defines the start and stop times of the cluster center as the earliest and latest times of the member profiles.

5.1.5.4 Compressor Actuation

Figure 87 shows the human actuation and the autonomous actuation compared to the estimated unactuated profile. In the estimated unactuated profile, periods with zero measured usage were replaced with the average idle usage as discussed previously. As seen in the figure, the human actuation performs better than the autonomous actuation due to the conservative approach of the autonomous actuation. This is verified in Table 19, where the savings from human control is markedly improved over the autonomous control. This difference in savings highlights the potential improvements to the autonomous actuation algorithm. The ratio of the achieved savings to the maximum possible savings is also displayed in the table for reference. Only one instance of error in actuation occurred, during which the compressor failed to autonomously turn on prior to the start of production usage.

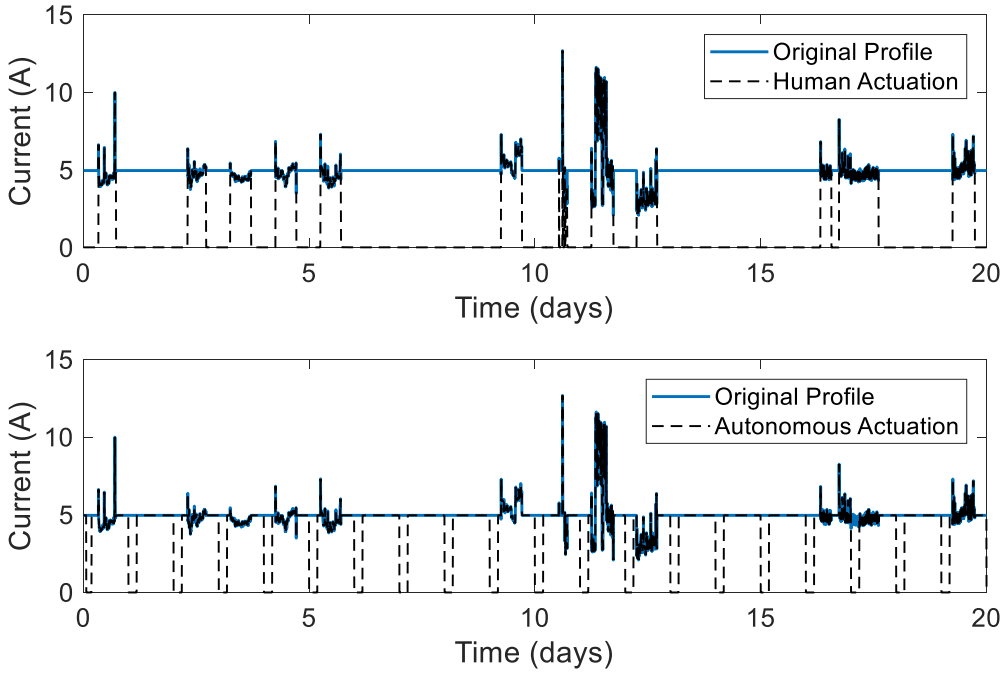

Figure 87. Comparison of human and autonomous actuation to the unactuated compressor loading for 20 days in the industrial manufacturing dataset.

5.1.6 Future Avenues of Research

While several future directions of research are envisioned, including implementation of the presented k-means algorithm on actual industrial compressed air systems, the following section presents future improvements to the k-means algorithm along with two specific directions of future research. The first direction is in using intelligent storage to reduce compressor energy while production is off. The second is an extension of the presented work in this chapter to autonomously adjust the pressure setpoint of the compressor during production periods to optimize energy usage while maintaining necessary pressure setpoints.

5.1.6.1 Improvements to the K-means Algorithm

There are several possible improvements that could be explored to improve the k-means algorithm. First, the standard deviation away from the cluster centers was used as the metric in this work as the factor to merge clusters in the trimming phase. Other potential metrics that compare the uniformity of clusters [109] can be evaluated in future work in merging the clusters. Additionally, future work can pair compressor data with occupancy data, other equipment data, or even entire facility demand data. These additional data could provide better estimates of plant production in order to improve actuation. Finally, reinforcement learning concepts can be used to improve the estimation. When the compressors fail to turn on at the start of production, this information can be incorporated into future predictions.

5.1.6.2 Intelligent Compressed Air Storage

Another avenue of potential research is in intelligent storage of compressed air. The algorithm presented in this chapter actuate the compressor based on production. Turning off the compressor is based on the analysis of observed data to determine when the compressor may be turned off. Because turning off the compressor is reactive instead of predictive, the compressor is consistently turned off when no production is running. In all of the tests with the algorithm, not a single instance was observed when the compressor was turned off prematurely. Turning the compressor back on, however, involves prediction on when production will resume. The inherent uncertainty of prediction results in either an overly conservative algorithm or instances when the compressor fails to turn on in time for production.

To avoid the uncertainty of prediction, the potential future research would use intelligent air storage. The system would be monitored for production, similar to what was done in this chapter. When production is stopped, the compressor setpoint is reduced to a lower pressure instead of completely shut down. A second, smaller air tank remains at high pressure. When air use is detected, the second air tank supplies the necessary high pressure until the main air system can reach the operating pressure setpoint.

A diagram of the proposed system is shown in Figure 88. The installation of the system is done at the location of the main receiver tank. The system is installed by tapping into the header lines on both the input and output of the receiver tank. The proposed system consists of two solenoid valves, a one-way value, a three-way valve, and a secondary air tank. Both solenoid valves and the three-way valve are connected and actuated by the central controller. During production, air flows through the receiver tank with the nominally open solenoid valve and three-way valve allowing air through the main header. The one-way check valve ensures the secondary air tank remains pressurized to the system pressure. Once production is off, the system will remain in the configuration, and the secondary tank will remain at operating pressure.

Once air use is detected while in the reduced pressure state, the three-way valve switches and releases the air from the secondary tank for production. Simultaneously, the nominally-open solenoid valve closes, redirecting all air from the compressor through the smaller, pressurized air tank.

While the compressor is running, the nominally-closed solenoid valves oscillates to slowly fill the receiver tank while maintaining pressure in the system through the smaller air tank.

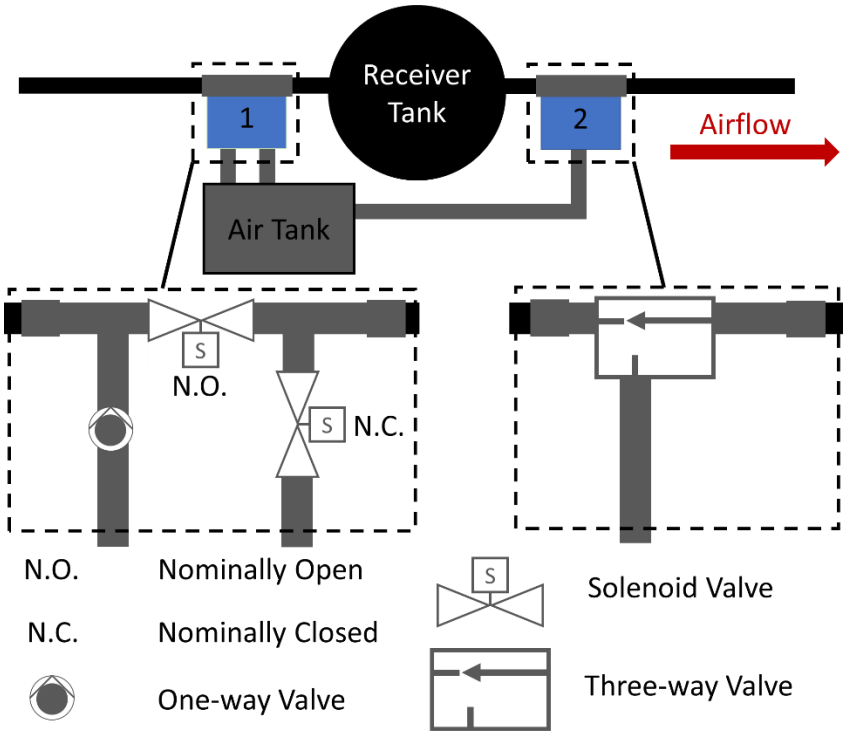

Figure 88. Schematic of potential intelligent air storage concept.

The future research would focus on detection methods to determine the initial air use when production starts. One potential challenge is in achieving sufficient measurement sensitivity to correctly identify when the compressed air system should be turned on. Other work has used the rate of pressure change to identify when compressor sequencing, providing initial evidence on the feasibility of using pressure changes to detect air usage [94].

Another potential research effort in this project is in correct sizing of the air tank. The tank should be large enough to meet the initial inrush of demand but small enough for economic purposes. Once the system is installed, the algorithm can experimentally determine the appropriate reduced pressure setpoint. This setpoint would be based on the compressor power, production air usage, and air storage capacity and can be estimated based on the system response.

Finally, the feasibility of slowly filling the main tank while maintaining pressure in the secondary tank will be explored. Solenoid valves can nominally switch very quickly, allowing the compressor to remain in the loaded pressure band while redirecting some of the airflow into the main tank. Potential implementation challenges will be explored, including potential control issues, air flow capacity limitations, pressure waves, etc.

5.1.6.3 Pressure Setpoint Adjustment using Flow on Multiple Lines

The algorithm presented in this chapter provides system level operation based on the pressure and power of the total compressed air system. Many industrial plants have several compressed air lines with major equipment and significant use. Often, the various uses of compressed air have different operating pressure requirements. Additionally, incorrectly sized headers create measurable pressure drop between the compressor and end use of equipment. This occurrence typically results from plant expansion of equipment without commensurate upgrades to compressed air headers. This pressure drop creates a virtual increase of the pressure setpoint at the compressor in order to maintain the necessary setpoint after the pressure drop. This is done due to the

lower initial cost of running the compressor at a higher pressure rather than upgrading the entire network of compressed air lines. Figure 89 shows an example of the pressure drop in third operating line from the simulation in Figure 64. The top plot shows the pressure at the outlet of the compressor to verify that the pressure drop results from pressure loss in the line. This pressure drop resulted from an undersized header due to the addition of the third production line using the existing air system.

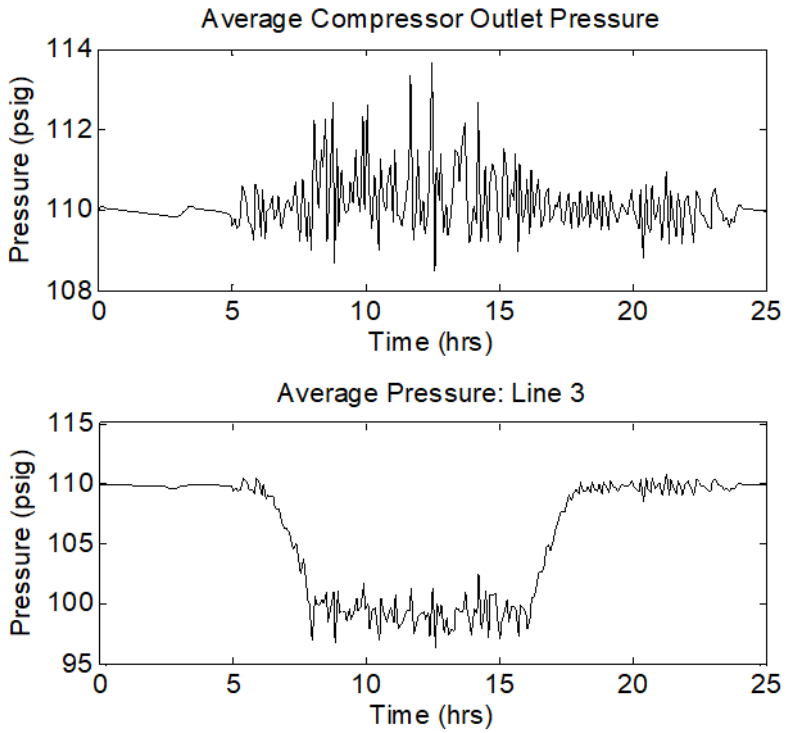

Figure 89. Example of a significant pressure drop due to an insufficiently-sized compressed air line.

The plant compressed air pressure setpoint is determined by the end use with the highest pressure requirement. When equipment with this high pressure requirement is not running, the compressor setpoint may be reduced to meet the highest setpoint among

equipment that is in production. This strategy requires additional knowledge on how the system is operating. With knowledge on compressed air use on major lines along with necessary setpoints, the compressor may be turned down during production hours.

As mass flow meters are expensive and at times have reliability issues, differential pressure between the compressor and the end use of air can be used to approximate the flow. Figure 90 shows the difference in pressure along with flow information for the second operating line in the simulation. While exact flow measurements would be difficult using pressure, the difference in pressure can identify periods of air use on the production line. This approach would require pressure gauges on major operating lines in order to determine when the production line is running.

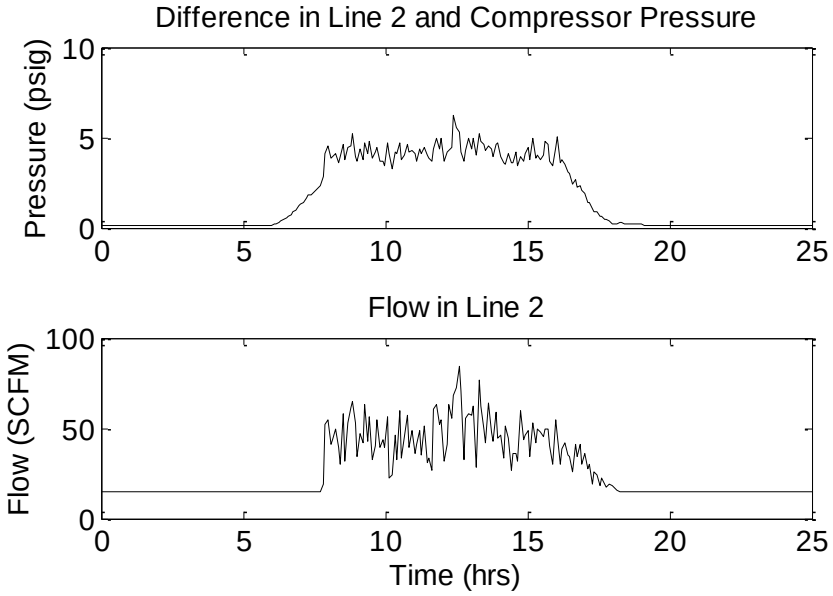

Figure 90. Use of pressure differences to estimate the overall production use of compressed air.

With a user-defined setpoint for each line, the k-means algorithm presented in this chapter can be used to identify times when each line is not being used. The current pressure setpoint is determined by the maximum pressure requirement of all lines that are currently in production. Figure 91 shows an example where the compressor can be turned down during production prior to automatically shutting off at the end of production. The compressor is still turned off at the end of all production, but the pressure is reduced during periods of time with some production still in the plant.

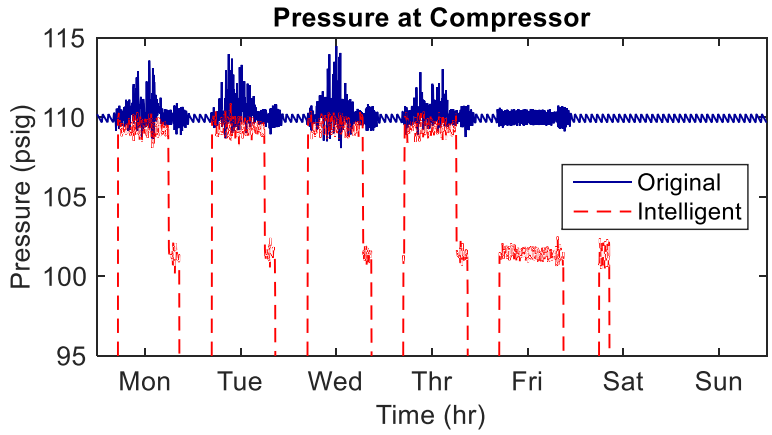

Figure 91. Example of how intelligent compressor actuation can decrease compressor energy usage further by monitoring individual lines.

5.2 Summary

This chapter demonstrated the potential improvement in energy efficiency through autonomous actuation with a test case on improved actuation of compressed air systems. Compressed air systems use relatively large amounts of electricity when the system is idle, and autonomously turning off equipment can generate energy savings with no effect on production.

A novel adaptation of a k-means algorithm was presented in this chapter to identify periods of time to turn off the compressor. The algorithm converted the measured power into discrete load states. These load states were used in the clustering algorithm, which included distinct growth and trimming phases to adaptively determine the number of necessary clusters. The effectiveness of the algorithm was evaluated with a simulation of a compressed air system. The algorithm was then tested on a set of 15-miniute demand datasets. Finally, the algorithm was tested on measured data from an industrial manufacturer. The effectiveness of the autonomous actuation was compared to human actuation, where production personnel manually turned off the compressor at the end of production periods. Overall, the algorithm effectively actuated the compressor to reduce energy usage. The only types of observed error occurred when the compressor did not turn on prior to the start of the following production period, an error that is primarily an inconvenience instead of production critical.

CHAPTER 6

CONCLUSIONS

This dissertation discusses how greater energy efficiency can be achieved within the Industrial Internet of Things. While multiple frameworks have been proposed in the literature, very few applications have emerged in the marketplace specific to the Industrial Internet of Things. This slow development in IIoT solutions is directly related to the slow adoption of new technologies in industry. This hesitation in technology adoption will likely slow the development of IIoT solutions in the field. While complete networking of all plant equipment with automated plant improvements is the end vision of IIoT, the realization of this vision will require extended effort and time. Focusing on targeted solutions for energy support systems, as proposed in the dissertation, will expediate IIoT solutions in plants. Specifically, by combining energy, maintenance, and safety solutions in common energy systems, faster technology development and implementation can be achieved.

This dissertation provided examples of how IIoT solutions can improve the energy efficiency of industrial energy systems either through improved diagnostics or intelligent actuation. These demonstrated examples target common energy systems that are typically found in industrial plants. By identifying improvements in common energy systems, the implementation of identified solutions can be greatly enhanced.

The first demonstrated example in the dissertation is a non-intrusive flow meter for compressed air systems. In both simulation and experimental validation, the

proposed method of using thermal responses to measured compressed air flow shows promise. The experimental results closely follow the expected trend seen in simulation with limited scatter in the data away from the correlation. These experimental results justify additional exploration and experimentation to improve the method. Since the proposed method is inexpensive, requires minimal setup, and is insensitive to changes of pressure in the pipe, the method can be economically developed for flow measurement in energy assessments. Additionally, the method of using the thermal response to measure flow can be applied to other fields of fault detection and diagnosis. Nonintrusive, inexpensive flow measurement with minimal calibration and setup is difficult in many fields. Some of these fields include BTU meters in industrial support systems, HVAC refrigerant flow in fault detection and diagnosis, and state monitoring of water and gas pipelines. This method can potentially provide flow solutions to these other application areas.

This method was advanced further by dynamically determining multiple parameters when measuring the temperature profile. While this advanced method shows promise in the ability to measure flow, accounting for heater dynamics is the most prominent research task going forward. The effects of the heater dynamics are seen both in the experimental tests on the black iron pipe and in the test case with the copper pipe. Accounting for these dynamics will ensure that differences in the temperature response are primarily caused by flow differences in the pipe, thereby improving the robustness of the method. In addition to accounting for heater dynamics, extensive testing needs to be performed on multiple materials to identify any other potential limitations that were not

manifest in the original experiments. If no further limitations are identified, the low cost of materials and minimal assembly will enable a market solution that is economical and can meet all the requirements for a non-intrusive flow meter for compressed air and other fluid systems.

The second presented example shows the automated analysis of building lighting systems to identify appropriate energy efficiency measures. This work lays the groundwork for more extensive capabilities in the automated analysis of building retrofit opportunities. Current building energy assessments are time-consuming and expensive, and automation of many repetitive tasks can reduce cost while improving overall assessment results. Additionally, this automation allows professional building auditors to focus their effort on unique challenges. Improving overall assessment efficacy will lead to greater participation rates among facilities, accelerating realized energy savings.

In IIoT, improved decision making and actuation is a necessary step to achieve improved energy or production savings. Specifically in IIoT for energy efficiency, autonomous actuation is important due to the relatively modest achievable savings. The final case study evaluated using intelligent actuation to achieve energy savings in compressed air systems. The presented method demonstrated the ability to provide energy savings for compressed air systems. However, the algorithm was shown to be conservative, and there exists potential improvements to increase the realized energy savings from the actuation. Approaches to increase the effectiveness of the algorithm decrease the amount of time when the system is running with no production air use. Potential future research avenues to achieve this decrease in idle time include

improvement of the k-means algorithm, adaptively adjusting pressure setpoints to optimize compressor use during production, and implementing intelligent storage to autonomously start the compressed air system when use is detected. Each of these methods can increase the realized savings with minimal additional cost. Additionally, other approaches can integrate other sensors in IIoT to provide contextual information on production. Total system power, electrical signals from equipment, and occupancy monitoring of the facility can all potentially improve the identification of idle periods and thereby increase the potential energy savings.

NOMENCLATURE

Parameter	Description	Units/Definition
A_c	Cross sectional area	m^2
$\alpha = \frac{k}{\rho c}$	Thermal diffusivity	$\frac{m^2}{s}$
k	Thermal conductivity	$\frac{W}{m \cdot K}$
ρ	Density	$\frac{kg}{m^3}$
c	Specific heat	$\frac{kJ}{kg \cdot K}$
E	Energy in material	J
T	Temperature	K
h	Convection coefficient	$\frac{W}{m^2 \cdot K}$
T_∞	Ambient temperature	K
P_i	Interior perimeter of pipe	m
q	Heat flux	$\frac{W}{m^2}$
erf(x)	Error Function	
$u = T - T_\infty$	Adjusted Temperature	K
$w = \frac{Q_c x}{4}$	Estimated Parameter	K
$f=\sqrt{x^2/4\alpha t}$	Non-dimensional Estimated Parameter	
$g= \sqrt{bt}$	Non-dimensional Estimated Parameter	
x	Axial location of pipe	m
t	Time	s
x0	x offset	m
Q	Thermal energy transfer	J

Re	Reynolds number	
Pr	Prandtl number	
D	Diameter	m
v	Velocity	$\frac{m}{s}$
μ	Dynamic viscosity	$\frac{kg \cdot m}{s}$
$\dot{V}$	Volumetric flow rate	$\frac{m^3}{s}$
r	Radius	m
V	Volume	m^3
W_{motor}	Motor power rating	hp
η_{motor}	Motor Efficiency	–
η_{comp}	Compressor Specific Efficiency	$\frac{\text{kW}}{SCFM}$
η_{is}	Isentropic Efficiency	–
P	Pressure	psi
k_c	Heat Capacity Ratio	–
R	Specific Gas Constant	$\frac{\text{kJ}}{\text{kg K}}$
$\dot{m}$	Mass Flow	$\frac{\text{kg}}{\text{s}}$
V	Volume	m^3
L	Length of Pipeline	m
f_D	Friction Factor	–
n	Number of 90° Turns	–
atm	Atmospheric	–
comp	Compressor	–

REFERENCES

[1] Gubbi, J., et al., *Internet of Things (IoT): A vision, architectural elements, and future directions.* Future generation computer systems, 2013. **29**(7): p. 1645-1660.

[2] Sundmaeker, H., et al., *Vision and challenges for realising the Internet of Things.* Cluster of European Research Projects on the Internet of Things, European Commision, 2010. **3**(3): p. 34-36.

[3] Tsai, C.-W., et al., *Data mining for Internet of Things: A survey.* IEEE Communications Surveys and Tutorials, 2014. **16**(1): p. 77-97.

[4] Da Xu, L., W. He, and S. Li, *Internet of things in industries: A survey.* IEEE Transactions on industrial informatics, 2014. **10**(4): p. 2233-2243.

[5] Perera, C., C.H. Liu, and S. Jayawardena, *The emerging internet of things marketplace from an industrial perspective: A survey.* IEEE Transactions on Emerging Topics in Computing, 2015. **3**(4): p. 585-598.

[6] Miorandi, D., et al., *Internet of things: Vision, applications and research challenges.* Ad Hoc Networks, 2012. **10**(7): p. 1497-1516.

[7] Perera, C., et al., *A survey on internet of things from industrial market perspective.* IEEE Access, 2014. **2**: p. 1660-1679.

[8] Chen, M., S. Mao, and Y. Liu, *Big data: A survey.* Mobile Networks and Applications, 2014. **19**(2): p. 171-209.

[9] Gartner, I., *Gartner's 2015 hype cycle for emerging technologies identifies the computing innovations that organizations should monitor*. Gartner Web site.. Published http://www. gartner. com/newsroom/id/3114217. Updated, 2015.

[10] Gordon, R.J., *The rise and fall of American growth: The US standard of living since the civil war*. 2017: Princeton University Press.

[11] Jeschke, S., et al., *Industrial Internet of Things*. Cham, Switzerland: Springer, 2017.

[12] Stankovic, J.A., *Research directions for the internet of things.* IEEE Internet of Things Journal, 2014. **1**(1): p. 3-9.

[13] Wang, H., et al. *Big data and industrial internet of things for the maritime industry in northwestern norway*. in *TENCON 2015-2015 IEEE Region 10 Conference*. 2015. IEEE.

[14] Teich, P., *Connecting with the Industrial Internet of Things (IIoT).* Moor Insights & Strategy, Tech. Rep, 2013.

[15] Lee, J., B. Bagheri, and H.-A. Kao. *Recent advances and trends of cyber-physical systems and big data analytics in industrial informatics*. in *International Proceeding of Int Conference on Industrial Informatics (INDIN)*. 2014.

[16] Chiang, L.H., E.L. Russell, and R.D. Braatz, *Fault detection and diagnosis in industrial systems*. 2000: Springer Science & Business Media.

[17] Jiang, H., et al., *Energy big data: A survey.* IEEE Access, 2016. **4**: p. 3844-3861.

[18] Zhou, K., C. Fu, and S. Yang, *Big data driven smart energy management: From big data to big insights.* Renewable and Sustainable Energy Reviews, 2016. **56**: p. 215-225.

[19] *Industrial Assessment Center Database*. 2014, U.S. Department of Energy.

[20] Sandberg, P. and M. Söderström, *Industrial energy efficiency: the need for investment decision support from a manager perspective.* Energy policy, 2003. **31**(15): p. 1623-1634.

[21] Moreno, M.V., et al., *Big data: the key to energy efficiency in smart buildings.* Soft Computing, 2016. **20**(5): p. 1749-1762.

[22] Bandyopadhyay, S., et al., *Role of middleware for internet of things: A study.* International Journal of Computer Science and Engineering Survey, 2011. **2**(3): p. 94-105.

[23] Atzori, L., A. Iera, and G. Morabito, *The internet of things: A survey.* Computer networks, 2010. **54**(15): p. 2787-2805.

[24] Liu, A., et al., *A green and reliable communication modeling for industrial internet of things.* Computers & Electrical Engineering, 2016.

[25] Katipamula, S. and M.R. Brambley, *Review article: Methods for fault detection, diagnostics, and prognostics for building systems—a review, part II.* HVAC&R Research, 2005. **11**(2): p. 169-187.

[26] Katipamula, S. and M.R. Brambley, *Review article: methods for fault detection, diagnostics, and prognostics for building systems—a review, Part I.* HVAC&R Research, 2005. **11**(1): p. 3-25.

[27] Li, H. and J. Braun, *An Improved Method for Fault Detection and Diagnosis Applied to Packaged Air Conditioners.* ASHRAE Transactions, 2003. **109**(2): p. 683-692.

[28] Li, H. and J. Braun, *Decoupling Features and Virtual Sensors for Diagnosis of Faults in Vapor Compression Air Conditioners.* International Journal of Refrigeration, 2007. **30**(3): p. 546-564.

[29] Li, H. and J. Braun, *A Methodology for Diagnosis Multiple Simultaneous Faults in Vapor-Compression Air Conditioners.* HVAC&R Research, 2007. **13**(2): p. 369-395.

[30] Janecke, A., T.J. Terrill, and B.P. Rasmussen, *A comparison of static and dynamic fault detection techniques for transcritical refrigeration.* International Journal of Refrigeration, 2017.

[31] Hsu, J.S., et al., *Comparison of induction motor field efficiency evaluation methods.* Industry Applications, IEEE Transactions on, 1998. **34**(1): p. 117-125.

[32] Challenge, M., *Determining Electric Motor Load and Efficiency.* Program of the US Department of Energy, 1997.

[33] Perera, C., et al., *Context aware computing for the internet of things: A survey.* IEEE Communications Surveys & Tutorials, 2014. **16**(1): p. 414-454.

[34] Jin, X., et al., *Significance and challenges of big data research.* Big Data Research, 2015. **2**(2): p. 59-64.

[35] Penrose, H., *Large electric motor reliability: what did the studies really say.* Dreisilker Electric Motors http://www. mt-online. com/feb2012/large-electric-motor-reliability-what-did-the-studies-really-say, 2012.

[36] Radu, C., *The Most Common Causes of Bearing Failure and the Importance of Bearing Lubrication.* RKB Technical Review-February, 2010.

[37] Singh, G. and S.a.A.S. Al Kazzaz, *Isolation and identification of dry bearing faults in induction machine using wavelet transform.* Tribology international, 2009. **42**(6): p. 849-861.

[38] Mijares, J.A., B.P. Rasmussen, and A.G. Parlos, *Detection of Lubrication Starvation in Ball Bearings and Preload Effects.*

[39] Yin, S., et al., *A comparison study of basic data-driven fault diagnosis and process monitoring methods on the benchmark Tennessee Eastman process.* Journal of Process Control, 2012. **22**(9): p. 1567-1581.

[40] Saidur, R., N. Rahim, and M. Hasanuzzaman, *A review on compressed-air energy use and energy savings.* Renewable and Sustainable Energy Reviews, 2010. **14**(4): p. 1135-1153.

[41] Beals, C., et al., *Improving compressed air system performance.* U.S. Department of Energy, 2003.

[42] Benton, N., *Compressed-Air Evaluation Protocol*, N.R.E. Laboratory, Editor. 2003, U.S. Department of Energy.

[43] Xenergy, *Assessment of the Market for Compressed Air Efficiency Services*, O.o.E.E.a.R.E. Office of Industrial Technologies, Editor. 2001, U.S. Department of Energy.

[44] Cai, M. and T. Kagawa, *Energy consumption assessment and energy loss analysis in pneumatic system.* Chinese Journal of Mechanical Engineering, 2007. **43**(9): p. 69.

[45] Dudić, S.P., et al., *Leakage quantification of compressed air on pipes using thermovision.* Thermal science, 2012. **16**(suppl. 2): p. 555-565.

[46] Pittard, M.T., et al., *Experimental and numerical investigation of turbulent flow induced pipe vibration in fully developed flow.* Review of scientific instruments, 2004. **75**(7): p. 2393-2401.

[47] Zhao, X., M. Yang, and H. Li, *Development, evaluation, and validation of a robust virtual sensing method for determining water flow rate in chillers.* HVAC&R Research, 2012. **18**(5): p. 874-889.

[48] Song, L., G. Wang, and M.R. Brambley, *Uncertainty analysis for a virtual flow meter using an air-handling unit chilled water valve.* HVAC&R Research, 2013. **19**(3): p. 335-345.

[49] Swamy, A., P. Li Song PhD, and P. Gang Wang PhD, *A virtual chilled-water flow meter development at air handling unit level.* ASHRAE Transactions, 2012. **118**: p. 1013.

[50] Arvoh, B.K., et al., *Online estimation of reject gas flow rates in compact flotation units for produced water treatment: A feasibility study.* Chemometrics and Intelligent Laboratory Systems, 2012. **114**: p. 87-98.

[51] Ying, S., et al., *Research of Optical Fiber Fluid Flow Monitoring System Based on Flow-Induced Pipe Vibration*, in *Photonics and Optoelectronics (SOPO), 2012 Symposium on*. 2012, IEEE. p. 1-3.

[52] Medeiros, K., C. Barbosa, and E. de Oliveira, *Flow Measurement by Piezoelectric Accelerometers: Application in the Oil Industry.* Petroleum Science and Technology, 2015. **33**(13-14): p. 1402-1409.

[53] Olama, M.M., et al., *An energy signature scheme for steam trap assessment and flow rate estimation using pipe-induced acoustic measurements*, in *SPIE Defense, Security, and Sensing*. 2012, International Society for Optics and Photonics. p. 83660F-83660F-8.

[54] Bergman, T.L., F.P. Incropera, and A.S. Lavine, *Fundamentals of heat and mass transfer*. 2011: John Wiley & Sons.

[55] Skogestad, S. and I. Postlethwaite, *Multivariable feedback control: analysis and design*. Vol. 2. 2007: Wiley New York.

[56] Thompson, A.S., D. Maynes, and J.D. Blotter, *Internal Turbulent Flow Induced Pipe Vibrations With and Without Baffle Plates*, in *ASME 2010 3rd Joint US-European Fluids Engineering Summer Meeting*. 2010, American Society of Mechanical Engineers. p. 649-659.

[57] Thompson, A.S., et al., *Experimental characterization of flow induced vibration in fully developed turbulent pipe flow*, in *ASME 2009 Fluids Engineering Division Summer Meeting*. 2009, American Society of Mechanical Engineers. p. 1771-1778.

[58] Evans, R.P., J.D. Blotter, and A.G. Stephens, *Flow rate measurements using flow-induced pipe vibration.* Journal of fluids engineering, 2004. **126**(2): p. 280-285.

[59] Fan, Z., M. Cai, and W. Xu, *Non-invasive and non-intrusive gas flow measurement based on the dynamic thermal characteristics of a pipeline.* Measurement Science and Technology, 2012. **23**(10): p. 105303.

[60] Xu, W., et al., *Soft Sensing Method Of LS-SVM Using Temperature Time Series For Gas Flow Measurements.* Metrology and Measurement Systems, 2015. **22**(3): p. 383-392.

[61] Parker, W., et al., *Flash method of determining thermal diffusivity, heat capacity, and thermal conductivity.* Journal of applied physics, 1961. **32**(9): p. 1679-1684.

[62] Polyanin, A.D., *Handbook of linear partial differential equations for engineers and scientists*. 2001: CRC press.

[63] Ng, E.W. and M. Geller, *A table of integrals of the error functions.* Journal of Research of the National Bureau of Standards B, 1969. **73**: p. 1-20.

[64] Ioannou, P.A. and J. Sun, *Robust adaptive control*. 2012: Courier Corporation.

[65] WolframAlpha. *Online Integral Calculator*. 2017 [cited 2016 July 13]; Available from: http://www.wolframalpha.com/calculators/integral-calculator/.

[66] Strauss, W.A., *Partial differential equations*. Vol. 92. 1992: Wiley New York.

[67] Bay, C.J., T.J. Terrill, and B.P. Rasmussen. *Autonomous Lighting Audits: Part 1—Building Navigation and Mapping*. in *ASME 2014 Dynamic Systems and Control Conference*. 2014. American Society of Mechanical Engineers.

[68] Terrill, T.J., C.J. Bay, and B.P. Rasmussen. *Autonomous Lighting Audits: Part 2—Light Identification and Analysis*. in *ASME 2014 Dynamic Systems and Control Conference*. 2014. American Society of Mechanical Engineers.

[69] Bay, C.J. and B.P. Rasmussen. *Simulation and validation of interior and exterior navigational strategies for autonomous robotic assessments of energy*. in *American Control Conference (ACC), 2015*. 2015. IEEE.

[70] DOE, *Buildings Share of U.S. Primary Energy Consumption*. 2012. p. Table 1.1.3.

[71] Standard, A., *Energy standard for buildings except low-rise residential buildings.* ASHRAE/IESNA Standard, 2013. **90**(1).

[72] *Commissioning for Federal Facilities*. 2014.

[73] Engineers, A.o.E. *Certified Energy Manager*.

[74] Halonen, L., E. Tetri, and P. Bhusal, *Guidebook on energy efficient electric lighting for buildings.* Espoo, Finland: Department of Electrical Engineering, Aalto University, 2010.

[75] Otsu, N., *A threshold selection method from gray-level histograms.* Automatica, 1975. **11**(285-296): p. 23-27.

[76] Chen, C.-Y. and R. Klette. *Image stitching—Comparisons and new techniques*. in *Computer Analysis of Images and Patterns*. 1999. Springer.

[77] Brown, M. and D.G. Lowe, *Automatic panoramic image stitching using invariant features.* International journal of computer vision, 2007. **74**(1): p. 59-73.

[78] Lowe, D.G. *Object recognition from local scale-invariant features*. in *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*. 1999. Ieee.

[79] Fischler, M.A. and R.C. Bolles, *Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography.* Communications of the ACM, 1981. **24**(6): p. 381-395.

[80] Elvidge, C.D., et al., *Spectral identification of lighting type and character.* Sensors, 2010. **10**(4): p. 3961-3988.

[81] Simons, R.H. and A.R. Bean, *Lighting engineering: applied calculations*. 2008: Routledge.

[82] DiLaura, D.L., et al., *The lighting handbook: Reference and application*. 2011: Illuminating Engineering Society of North America New York (NY).

[83] Yang, R. and M.W. Newman. *Living with an intelligent thermostat: advanced control for heating and cooling systems*. in *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. 2012. ACM.

[84] Perry, W. *The IIoT Comes to Comrpessed Air*. 2016 [cited 2017 12/29]; Available from: https://www.plantservices.com/articles/2016/ca-iiot-comes-to-compressed-air/.

[85] Camber, M. *Air Compressor Control for Industry 4.0*. 2017; Available from: https://www.airbestpractices.com/technology/compressor-controls/air-compressor-control-industry-40.

[86] Perry, W. *What the Industrial Interent of Things Means for Compressed Air*. 2016 [cited 2017 12/29]; Available from: https://www.chemicalprocessing.com/industrynews/2016/what-does-the-industrial-internet-of-things-mean-for-compressed-air/.

[87] Manzke, T.D.J. *Webinar: IoT and Compressed Air Management Systems*. 2017 [cited 2017 12/29]; Available from: https://www.airbestpractices.com/system-assessments/compressor-controls/iot-and-compressed-air-management-systems.

[88] Bay-Controls. *Three Things You Need to Know About Actionable Data for Compressed Air Systems*. 2017 [cited 2017 12/29]; Available from: http://www.baycontrols.com/blog/actionable-data-for-compressed-air/.

[89] Aamir, A. *Industrial Internet of Things for the Compressed Air Industry*. 2017 [cited 2017 12/29]; Available from: https://www.linkedin.com/pulse/industrial-internet-things-compressed-air-industry-athar-aamir/.

[90] Kaeser. *Kaeser Industrial Screw Compressors Now Come with Sigma Control 2*. 2017 [cited 2017 12/29]; Available from: https://www.airbestpractices.com/technology/industry-news/kaeser-industrial-screw-compressors-now-come-sigma-control-2.

[91] Intel. *Improving Downtime and Energy Efficiency with IoT-Connected Air Comrpessors*. 2017 [cited 2017 12/29]; Available from:

https://www.intel.com/content/www/us/en/internet-of-things/solution-briefs/improving-downtime-energy-efficiency-iot-connected-brief.html.

[92] Clark, P. *How Can You Sell Air? As a Service, Of Course*. 2017; Available from: http://www.digitalistmag.com/iot/2016/05/09/how-can-you-sell-air-as-a-service-04192752.

[93] Honeywell. *Honeywell and Spark Dynamics Collaborate on Industrial Internet of Things (IIoT) Solutions*. 2016 [cited 2017 12/22]; Available from: https://www.honeywellprocess.com/en-US/news-and-events/Pages/pr11162016-honeywell-and-sparks-dynamics-collaborate-on-industrial-internet-of-things-solutions.aspx.

[94] Hoetzel, J. *Easy to Implement Master Control and Monitoring Systems*. 2017 [cited 2017 12/29]; Available from: https://www.airbestpractices.com/technology/compressor-controls/easy-implement-master-control-and-monitoring-systems?page=6.

[95] Spark-Dynamics. *Remote Monitoring of Industrial Equipment*. 2017 [cited 2017 12/29]; Available from: https://www.sparksdynamics.com/.

[96] CompAir. *CompAir Compressed Air Intelligent System Controller*. 2017 [cited 2017 12/22]; Available from: http://www.engnetglobal.com/c/c.aspx/COM110/productdetail/compair-smartair-compressed-air-intelligent-system.

[97] Bonanomi, A. *Intelligent compressed air management with M.I.E.M.* 2017 [cited 2017 12/29]; Available from:

http://www.powertransmissionworld.com/intelligent-compressed-air-management-with-m-i-e-m/.

[98] van Tonder, A.M., J. Marais, and G. Bolt. *Practical considerations in energy efficient compressor set-point control strategies*. in *Industrial and Commercial Use of Energy (ICUE), 2014 International Conference on the*. 2014. IEEE.

[99] Bhatia, S.K. *Hierarchical clustering for image databases*. in *Electro Information Technology, 2005 IEEE International Conference on*. 2005. IEEE.

[100] Al-Jubouri, H., H. Du, and H. Sellahewa. *Adaptive clustering based segmentation for image classification*. in *Computer Science and Electronic Engineering Conference (CEEC), 2013 5th*. 2013. IEEE.

[101] Gupta, P. and A. Sharma. *A framework for hierarchical clustering based indexing in search engines*. in *Parallel Distributed and Grid Computing (PDGC), 2010 1st International Conference on*. 2010. IEEE.

[102] Agrawal, U., et al. *K-means clustering for adaptive wavelet based image denoising*. in *Computer Engineering and Applications (ICACEA), 2015 International Conference on Advances in*. 2015. IEEE.

[103] Alain, M., et al. *Clustering-based methods for fast epitome generation*. in *Signal Processing Conference (EUSIPCO), 2014 Proceedings of the 22nd European*. 2014. IEEE.

[104] Chauhan, E. and A. Asthana, *A Review of Hierarchal Clustering Algorithm and Evaluation.* International Journal of Engineering Science, 2017. **13442**.

[105] Costacurta, L.G. and M.A. Sanz-Bobi. *Application of clustering methods for discovering patterns of energy use in regional areas for the residential sector*. in *PowerTech, 2017 IEEE Manchester*. 2017. IEEE.

[106] Yao, R. and K. Steemers, *A method of formulating energy load profile for domestic buildings in the UK.* Energy and Buildings, 2005. **37**(6): p. 663-671.

[107] Widén, J., et al., *Constructing load profiles for household electricity and hot water from time-use data—Modelling approach and validation.* Energy and Buildings, 2009. **41**(7): p. 753-768.

[108] Hahn, H., S. Meyer-Nieberg, and S. Pickl, *Electric load forecasting methods: Tools for decision making.* European Journal of Operational Research, 2009. **199**(3): p. 902-907.

[109] Mutanen, A., et al., *Customer classification and load profiling method for distribution systems.* Power Delivery, IEEE Transactions on, 2011. **26**(3): p. 1755-1763.

[110] Ramos, S., et al. *A data mining framework for electric load profiling*. in *Innovative Smart Grid Technologies Latin America (ISGT LA), 2013 IEEE PES Conference On*. 2013. IEEE.

[111] Kwac, J., J. Flora, and R. Rajagopal, *Household energy consumption segmentation using hourly data.* Smart Grid, IEEE Transactions on, 2014. **5**(1): p. 420-430.

[112] Kwac, J., et al. *Utility customer segmentation based on smart meter data: Empirical study*. in *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*. 2013. IEEE.

[113] Tsekouras, G., et al., *A pattern recognition methodology for evaluation of load profiles and typical days of large electricity customers.* Electric Power Systems Research, 2008. **78**(9): p. 1494-1510.

[114] Lavin, A. and D. Klabjan, *Clustering time-series energy data from smart meters.* Energy efficiency, 2015. **8**(4): p. 681-689.

[115] Chow, T. and C. Leung, *Neural network based short-term load forecasting using weather compensation.* Power Systems, IEEE Transactions on, 1996. **11**(4): p. 1736-1742.

[116] Liao, H. and D. Niebur, *Load profile estimation in electric transmission networks using independent component analysis.* Power Systems, IEEE Transactions on, 2003. **18**(2): p. 707-715.

[117] Al-Hamadi, H. and S. Soliman, *Short-term electric load forecasting based on Kalman filtering algorithm with moving window weather and load model.* Electric power systems research, 2004. **68**(1): p. 47-59.

[118] Khotanzad, A., et al., *ANNSTLF-a neural-network-based electric load forecasting system.* Neural Networks, IEEE Transactions on, 1997. **8**(4): p. 835-846.

[119] Beccali, M., et al., *Forecasting daily urban electric load profiles using artificial neural networks.* Energy conversion and management, 2004. **45**(18): p. 2879-2900.

[120] Sousa, J., L. Neves, and H. Jorge, *Assessing the relevance of load profiling information in electrical load forecasting based on neural network models.* International Journal of Electrical Power & Energy Systems, 2012. **40**(1): p. 85-93.

[121] Halkidi, M., Y. Batistakis, and M. Vazirgiannis. *Clustering algorithms and validity measures*. in *Scientific and Statistical Database Management, 2001. SSDBM 2001. Proceedings. Thirteenth International Conference on*. 2001. IEEE.

[122] Sahu, M., K. Parvathi, and M.V. Krishna, *Parametric Comparison of K-means and Adaptive K-means Clustering Performance on Different Images.* International Journal of Electrical and Computer Engineering (IJECE), 2017. **7**(2): p. 810-817.

[123] Nikulin, V., *Threshold-based clustering with merging and regularization in application to network intrusion detection.* Computational statistics & data analysis, 2006. **51**(2): p. 1184-1196.

[124] Xu, T.-S., et al., *Hierarchical K-means method for clustering large-scale advanced metering infrastructure data.* IEEE Transactions on Power Delivery, 2017. **32**(2): p. 609-616.

[125] Pal, R., et al., *Towards Dynamic Demand Response On Efficient Consumer Grouping Algorithmics.* IEEE Transactions on Sustainable Computing, 2016. **1**(1): p. 20-34.

[126] Khan, I., et al., *Segmentation of Factories on Electricity Consumption Behaviors Using Load Profile Data.* IEEE Access, 2016. **4**: p. 8394-8406.

[127] Wang, Y., et al., *Clustering of electricity consumption behavior dynamics toward big data applications.* IEEE transactions on smart grid, 2016. **7**(5): p. 2437-2447.

[128] Mets, K., F. Depuydt, and C. Develder, *Two-stage load pattern clustering using fast wavelet transformation.* IEEE Transactions on Smart Grid, 2016. **7**(5): p. 2250-2259.

[129] Bhatia, S.K. *Adaptive K-Means Clustering*. in *FLAIRS Conference*. 2004.

[130] Chappell, J.S., *A transient fluid and thermodynamic model of a compressed air system*. 2011: The University of Alabama.

[131] Batchelor, G.K., *An introduction to fluid dynamics*. 2000: Cambridge university press.

[132] Carslaw, H.S. and J.C. Jaeger, *Conduction of heat in solids.* Oxford: Clarendon Press, 1959, 2nd ed., 1959. **1**.

APPENDIX A

DERIVATION OF INDIVIDUAL NODE EQUATIONS FOR NUMERIC MODEL

The individual state equations can be quite involved when considering the heat transfer effects from each neighboring element. As such, the main dissertation presents the heat transfer equation for a state on the inside surface of the pipe. Equations for internal states, states on the exterior of the pipe, and states for the corners of the pipes are all found in this appendix.

The model also captures the dynamics of air heating up in the pipe and the resulting asymmetric temperature profile around the heater. As the equations for the states that include air are significantly different from the states of the pipe, these equations are also derived in this section

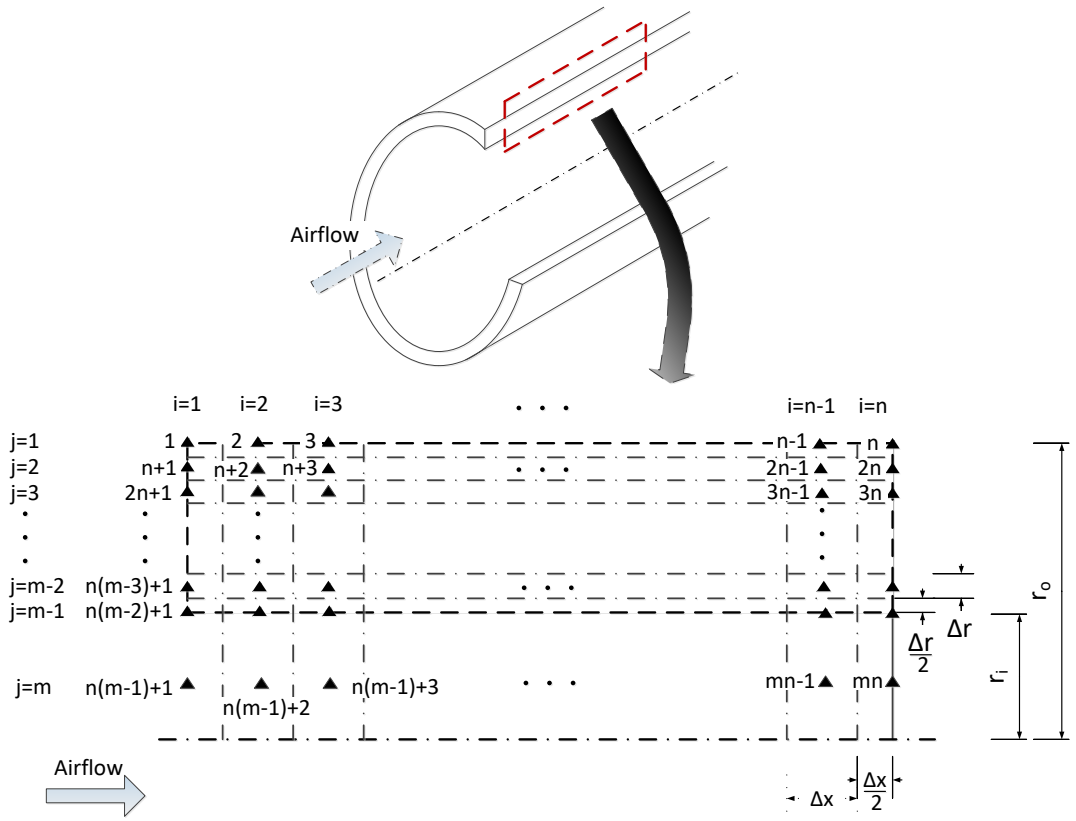

Figure 92. Outline of the model used for analyzing the dynamic response of the pipe.

A.1 Derivation of Control Volume Equations

A.1.1 Interior Nodes

The interior nodes are interior nodes in the pipe material in which no boundary conditions exist. Thermal conduction is the only mode of heat transfer in these nodes, and so the energy balance equation for these nodes is given by the following equation.

For i = 2:n-1, j = 2:m-2

$$\rho c\left(2\pi r_j \Delta r \Delta x\right)\frac{\partial T_{(j-1)n+i}}{\partial t}$$

$$= k2\pi r_j \Delta r \frac{\left(T_{(j-1)n+i-1} - T_{(j-1)n+i}\right)}{\Delta x} + k2\pi(r_j$$

$$+ \frac{\Delta r}{2})\Delta x \frac{\left(T_{(j-2)n+i} - T_{(j-1)n+i}\right)}{\Delta r} + k2\pi r_j \Delta r \frac{\left(T_{(j-1)n+i+1} - T_{(j-1)n+i}\right)}{\Delta x}$$

$$+ k2\pi(r_j - \frac{\Delta r}{2})\Delta x \frac{\left(T_{jn+i} - T_{(j-1)n+i}\right)}{\Delta r}$$

This equation can be simplified to the following form.

$$\frac{\partial T_{(j-1)n+i}}{\partial t} = \alpha\left[\frac{1}{(\Delta x)^2}\left(T_{(j-1)n+i-1} - T_{(j-1)n+i}\right) + \frac{(r_j + \frac{\Delta r}{2})}{r_j(\Delta r)^2}\left(T_{(j-2)n+i} - T_{(j-1)n+i}\right)\right.$$

$$\left. + \frac{1}{(\Delta x)^2}\left(T_{(j-1)n+i+1} - T_{(j-1)n+i}\right) + \frac{(r_j - \frac{\Delta r}{2})}{r_j(\Delta r)^2}\left(T_{jn+i} - T_{(j-1)n+i}\right)\right]$$

Finally, the temperature terms at the node of interest can be grouped together.

$$\frac{\partial T_{(j-1)n+i}}{\partial t} = -2\alpha\left[\frac{1}{(\Delta x)^2} + \frac{1}{(\Delta r)^2}\right]T_{(j-1)n+i}$$

$$+ \left[\frac{\alpha}{(\Delta x)^2}\left(T_{(j-1)n+i-1} + T_{(j-1)n+i+1}\right) + \frac{\alpha\left(r_j + \frac{\Delta r}{2}\right)}{r_j(\Delta r)^2}\left(T_{(j-2)n+i}\right)\right.$$

$$\left. + \frac{\alpha\left(r_j - \frac{\Delta r}{2}\right)}{r_j(\Delta r)^2}\left(T_{jn+i}\right)\right]$$

A.1.2 Outer Pipe Surface

The outer pipe surface has similar equations to the interior nodes. Since the outer surface is insulated, no heat transfer occurs off the outer edge.

For i =2:n-1, j = 1

$$\rho c\left(\pi r_o \Delta r \Delta x\left(1-\frac{\Delta r}{4r_o}\right)\right)\frac{\partial T_i}{\partial t}$$

$$= k\pi r_o \Delta r\left(1-\frac{1}{4r_o}\right)\frac{(T_{i-1}-T_i)}{\Delta x} + k\pi r_o \Delta r\left(1-\frac{1}{4r_o}\right)\frac{(T_{i+1}-T_i)}{\Delta x}$$

$$+ k2\pi\left(r_o-\frac{\Delta r}{2}\right)\Delta x\ \frac{(T_{n+i}-T_i)}{\Delta r}$$

The simplified equation is the following.

$$\frac{\partial T_i}{\partial t} = \frac{\alpha}{1-\frac{\Delta r}{4r_o}}\left[\left(1-\frac{\Delta r}{4r_o}\right)\frac{(T_{i-1}-T_i)}{(\Delta x)^2} + \left(1-\frac{\Delta r}{4r_o}\right)\frac{(T_{i+1}-T_i)}{(\Delta x)^2}\right.$$

$$\left. + \frac{2(r_o-\frac{\Delta r}{2})}{r_o(\Delta r)^2}\ (T_i - T_{n+i})\right]$$

Grouping the temperature terms at the node results in the final equation as shown below.

$$\frac{\partial T_i}{\partial t} = -\alpha\left[\frac{2}{(\Delta x)^2} + \frac{2\left(r_o-\frac{\Delta r}{2}\right)}{\left(r_o-\frac{\Delta r}{4}\right)(\Delta r)^2}\right]T_i$$

$$+\left[\frac{\alpha}{(\Delta x)^2}(T_{i+1}+T_{i-1}) + \frac{2\alpha\left(r_o-\frac{\Delta r}{2}\right)}{\left(r_o-\frac{\Delta r}{4}\right)(\Delta r)^2}\ (T_{n+i})\right]$$

A.1.3 Inner Pipe Surface

The inner pipe surface nodes have similar equations as before, except the surface that interfaces with the fluid has a convection term. The following equation describes these nodes.

For i =2:n-1, j = m-1

$$\rho c\left(\pi r_I \Delta r \Delta x\left(1+\frac{\Delta r}{4r_I}\right)\right)\frac{\partial T_{(m-2)n+i}}{\partial t}$$

$$= k\pi r_I \Delta r\left(1+\frac{1}{4r_I}\right)\frac{\left(T_{(m-2)n+i-1}-T_{(m-2)n+i}\right)}{\Delta x}$$

$$+ k\pi r_I \Delta r\left(1+\frac{1}{4r_I}\right)\frac{\left(T_{(m-2)n+i+1}-T_{(m-2)n+i}\right)}{\Delta x}$$

$$+ k2\pi\left(r_I+\frac{\Delta r}{2}\right)\Delta x\,\frac{\left(T_{(m-3)n+i}-T_{(m-2)n+i}\right)}{\Delta r}$$

$$+ h2\pi r_I \Delta x\left(T_{(m-1)n+i}-T_{(m-2)n+i}\right)$$

The simplified equation is the following.

$$\frac{\partial T_{(m-2)n+i}}{\partial t} = \left[\frac{\alpha}{(\Delta x)^2}\left(T_{(m-2)n+i-1}-T_{(m-2)n+i}\right)+\frac{\alpha}{(\Delta x)^2}\left(T_{(m-2)n+i+1}-T_{(m-2)n+i}\right)+\right.$$

$$\left.\frac{2\alpha}{(\Delta r)^2}\frac{\left(r_I+\frac{\Delta r}{2}\right)}{\left(r_I+\frac{\Delta r}{4}\right)}\left(T_{(m-3)n+i}-T_{(m-2)n+i}\right)+\frac{2h}{\left(1+\frac{\Delta r}{4r_I}\right)\rho c\Delta r}\left(T_{(m-1)n+i}-T_{(m-2)n+i}\right)\right]$$

The temperature terms at the node of interest can be grouped together.

$$\frac{\partial T_{(m-2)n+i}}{\partial t} = -\left[\frac{2\alpha}{(\Delta x)^2}+\frac{2\alpha}{(\Delta r)^2}\frac{\left(r_I+\frac{\Delta r}{2}\right)}{\left(r_I+\frac{\Delta r}{4}\right)}+\frac{2h}{\left(1+\frac{\Delta r}{4r_I}\right)\rho c\Delta r}\right]T_{(m-2)n+i}+\left[\frac{\alpha}{(\Delta x)^2}\left(T_{(m-2)n+i+1}+\right.\right.$$

$$\left.\left.T_{(m-2)n+i-1}\right)+\frac{2\alpha}{(\Delta r)^2}\frac{\left(r_I+\frac{\Delta r}{2}\right)}{\left(r_I+\frac{\Delta r}{4}\right)}T_{(m-3)n+i}+\frac{2h}{\left(1+\frac{\Delta r}{4r_I}\right)\rho c\Delta r}T_{(m-1)n+i}\right]$$

A.1.4 Pipe Edges

The modeled pipe is assumed to be infinitely long in considering thermal effects. Therefore, the location of the pipe edges along the x-coordinate is chosen such that the boundaries do not affect the temperatures at points of interest. Although any boundary condition could theoretically be chosen, as there will be no effects from the boundary, a boundary condition is chosen that will approximate the temperature conduction through the boundary. The true response will deviate most from this approximated response when a frequency input is given to the system. The approximate response is developed by assuming a temperature in the boundary node. The temperature is approximated by assuming a proportionality term with the difference between the node temperature and the ambient temperature. Essentially, this acts as if convection takes place on the edges of the pipe.

For i=1, j = 2:m-2

$$2\rho c\left(\pi r_j \Delta r \frac{\Delta x}{2}\right)\frac{\partial T_{(j-1)n+1}}{\partial t}$$

$$= h_r 2\pi r_j \Delta r\left(T_\infty - T_{(j-1)n+1}\right) + k2\pi\left(r_j + \frac{\Delta r}{2}\right)\frac{\Delta x}{2}\,\frac{\left(T_{(j-2)n+1} - T_{(j-1)n+1}\right)}{\Delta r}$$

$$+ k2\pi r_j \Delta r\,\frac{\left(T_{(j-1)n+2} - T_{(j-1)n+1}\right)}{\Delta x}$$

$$+ k2\pi\left(r_j - \frac{\Delta r}{2}\right)\frac{\Delta x}{2}\,\frac{\left(T_{jn+1} - T_{(j-1)n+1}\right)}{\Delta r}$$

The equation is simplified to the following form.

$$\frac{\partial T_{(j-1)n+1}}{\partial t} = \left[\frac{2h_r}{\rho c \Delta x}\left(T_\infty - T_{(j-1)n+1}\right) + \alpha \frac{\left(r_j + \frac{\Delta r}{2}\right)}{r_j (\Delta r)^2}\left(T_{(j-2)n+1} - T_{(j-1)n+1}\right) \right.$$

$$\left. + \frac{2\alpha}{(\Delta x)^2}\left(T_{(j-1)n+2} - T_{(j-1)n+1}\right) + \alpha \frac{\left(r_j - \frac{\Delta r}{2}\right)}{r_j (\Delta r)^2}\left(T_{jn+1} - T_{(j-1)n+1}\right) \right]$$

The temperature terms at the node of interest can be grouped together.

$$\frac{\partial T_{(j-1)n+1}}{\partial t} = -\left[\frac{2\alpha}{(\Delta x)^2} + \frac{2\alpha}{(\Delta r)^2} + \frac{2h_r}{\rho c \Delta x} \right]\left(T_{(j-1)n+1}\right)$$

$$+ \left[\frac{2\alpha}{(\Delta x)^2}\left(T_{(j-1)n+2}\right) + \frac{\alpha\left(r_j + \frac{\Delta r}{2}\right)}{r_j (\Delta r)^2}\left(T_{(j-2)n+1}\right) + \frac{\alpha\left(r_j - \frac{\Delta r}{2}\right)}{r_j (\Delta r)^2}\left(T_{jn+1}\right) \right.$$

$$\left. + \frac{2h_r}{\rho c \Delta x} T_\infty \right]$$

For i=n, j = 2:m-2

$$2\rho c\left(\pi r_j \Delta r \frac{\Delta x}{2}\right)\frac{\partial T_{jn}}{\partial t}$$

$$= h_r 2\pi r_j \Delta r\left(T_\infty - T_{jn}\right) + k2\pi\left(r_j + \frac{\Delta r}{2}\right)\frac{\Delta x}{2}\frac{\left(T_{(j-1)n} - T_{jn}\right)}{\Delta r}$$

$$+ k2\pi r_j \Delta r \frac{\left(T_{jn-1} - T_{jn}\right)}{\Delta x} + k2\pi\left(r_j - \frac{\Delta r}{2}\right)\frac{\Delta x}{2}\frac{\left(T_{(j+1)n} - T_{jn}\right)}{\Delta r}$$

The equation is simplified to the following form.

$$\frac{\partial T_{jn}}{\partial t} = \left[\frac{2h_r}{\rho c \Delta x}\left(T_\infty - T_{jn}\right) + \alpha \frac{\left(r_j + \frac{\Delta r}{2}\right)}{r_j (\Delta r)^2}\left(T_{(j-1)n} - T_{jn}\right) + \frac{2\alpha}{(\Delta x)^2}\left(T_{jn-1} - T_{jn}\right) \right.$$

$$\left. + \alpha \frac{\left(r_j - \frac{\Delta r}{2}\right)}{r_j (\Delta r)^2}\left(T_{(j+1)n} - T_{jn}\right) \right]$$

The temperature terms at the node of interest can be grouped together.

$$\frac{\partial T_{jn}}{\partial t} = -\left[\frac{2\alpha}{(\Delta x)^2} + \frac{2\alpha}{(\Delta r)^2} + \frac{2h_r}{\rho c \Delta x} \right]\left(T_{jn}\right)$$

$$+ \left[\frac{2\alpha}{(\Delta x)^2}\left(T_{jn-1}\right) + \frac{\alpha\left(r_j + \frac{\Delta r}{2}\right)}{r_j (\Delta r)^2}\left(T_{(j-1)n}\right) + \frac{\alpha\left(r_j - \frac{\Delta r}{2}\right)}{r_j (\Delta r)^2}\left(T_{(j+1)n}\right) \right.$$

$$\left. + \frac{2h_r}{\rho c \Delta x} T_\infty \right]$$

A.1.5 Corner Nodes

The equations for the corner nodes are given below. These are derived in the same method as below, but they are included here for completeness.

For i=1, j = 1;

$$2\rho c\left(\pi r_o\left(1 - \frac{\Delta r}{4r_o}\right)\frac{\Delta r}{2}\frac{\Delta x}{2}\right)\frac{\partial T_1}{\partial t}$$

$$= h_r 2\pi r_o\left(1 - \frac{\Delta r}{4r_o}\right)\frac{\Delta r}{2}\left(T_\infty - T_1\right) + k2\pi r_o\left(1 - \frac{\Delta r}{4r_o}\right)\frac{\Delta r}{2}\frac{\left(T_2 - T_1\right)}{\Delta x}$$

$$+ k2\pi\left(r_o - \frac{\Delta r}{2}\right)\frac{\Delta x}{2}\frac{\left(T_{n+1} - T_1\right)}{\Delta r}$$

The equation is simplified to the following form.

$$\frac{\partial T_1}{\partial t} = \frac{2h_r}{\rho c \Delta x}(T_\infty - T_1) + \frac{2\alpha}{(\Delta x)^2}(T_2 - T_1) + \frac{2\alpha}{(\Delta r)^2}\frac{\left(1 - \frac{\Delta r}{2r_o}\right)}{\left(1 - \frac{\Delta r}{4r_o}\right)}(T_{n+1} - T_1)$$

The temperature terms at the node of interest can be grouped together.

$$\frac{\partial T_1}{\partial t} = -\left[\frac{2h_r}{\rho c \Delta x} + \frac{2\alpha}{(\Delta x)^2} + \frac{2\alpha}{(\Delta r)^2}\frac{\left(1 - \frac{\Delta r}{2r_o}\right)}{\left(1 - \frac{\Delta r}{4r_o}\right)}\right](T_1) + \frac{2h_r}{\rho c \Delta x}T_\infty + \frac{2\alpha}{(\Delta x)^2}(T_2)$$

$$+ \frac{2\alpha}{(\Delta r)^2}\frac{\left(1 - \frac{\Delta r}{2r_o}\right)}{\left(1 - \frac{\Delta r}{4r_o}\right)}(T_{n+1})$$

For i = n, j = 1;

$$2\rho c\left(\pi r_o\left(1 - \frac{\Delta r}{4r_o}\right)\frac{\Delta r}{2}\frac{\Delta x}{2}\right)\frac{\partial T_n}{\partial t}$$

$$= h_r 2\pi r_o\left(1 - \frac{\Delta r}{4r_o}\right)\frac{\Delta r}{2}(T_\infty - T_n) + k2\pi r_o\left(1 - \frac{\Delta r}{4r_o}\right)\frac{\Delta r}{2}\frac{(T_{n-1} - T_n)}{\Delta x}$$

$$+ k2\pi\left(r_o - \frac{\Delta r}{2}\right)\frac{\Delta x}{2}\frac{(T_{2n} - T_n)}{\Delta r}$$

The equation is simplified to the following form.

$$\frac{\partial T_n}{\partial t} = \frac{2h_r}{\rho c \Delta x}(T_\infty - T_n) + \frac{2\alpha}{(\Delta x)^2}(T_{n-1} - T_n) + \frac{2\alpha}{(\Delta r)^2}\frac{\left(1 - \frac{\Delta r}{2r_o}\right)}{\left(1 - \frac{\Delta r}{4r_o}\right)}(T_{2n} - T_n)$$

The temperature terms at the node of interest can be grouped together.

$$\frac{\partial T_n}{\partial t} = -\left[\frac{2h_r}{\rho c \Delta x} + \frac{2\alpha}{(\Delta x)^2} + \frac{2\alpha}{(\Delta r)^2}\frac{\left(1-\frac{\Delta r}{2r_o}\right)}{\left(1-\frac{\Delta r}{4r_o}\right)}\right](T_n) + \frac{2h_r}{\rho c \Delta x}T_\infty + \frac{2\alpha}{(\Delta x)^2}(T_{n-1})$$

$$+ \frac{2\alpha}{(\Delta r)^2}\frac{\left(1-\frac{\Delta r}{2r_o}\right)}{\left(1-\frac{\Delta r}{4r_o}\right)}(T_{2n})$$

For i =1, j = m-1

$$\rho c\left(2\pi r_I\left(1+\frac{\Delta r}{4r_I}\right)\frac{\Delta r}{2}\frac{\Delta x}{2}\right)\frac{\partial T_{(m-2)n+1}}{\partial t}$$

$$= h_r 2\pi r_I\left(1+\frac{\Delta r}{4r_I}\right)\frac{\Delta r}{2}\left(T_\infty - T_{(m-2)n+1}\right)$$

$$+ k2\pi r_I\left(1+\frac{\Delta r}{4r_I}\right)\frac{\Delta r}{2}\frac{\left(T_{(m-2)n+2} - T_{(m-2)n+1}\right)}{\Delta x}$$

$$+ k2\pi\left(r_I+\frac{\Delta r}{2}\right)\frac{\Delta x}{2}\frac{\left(T_{(m-3)n+1} - T_{(m-2)n+1}\right)}{\Delta r}$$

$$+ h2\pi r_I\frac{\Delta x}{2}\left(T_{(m-1)n+1} - T_{(m-2)n+1}\right)$$

The equation is simplified to the following form.

$$\frac{\partial T_{(m-2)n+1}}{\partial t} = \frac{2h_r}{\rho c \Delta x}\left(T_\infty - T_{(m-2)n+1}\right) + \frac{2\alpha}{(\Delta x)^2}\left(T_{(m-2)n+2} - T_{(m-2)n+1}\right)$$

$$+ \frac{2\alpha}{(\Delta r)^2}\frac{\left(1+\frac{\Delta r}{2r_I}\right)}{\left(1+\frac{\Delta r}{4r_I}\right)}\left(T_{(m-3)n+1} - T_{(m-2)n+1}\right)$$

$$+ \frac{2h}{\rho c \Delta r\left(1+\frac{\Delta r}{4r_I}\right)}\left(T_{(m-1)n+1} - T_{(m-2)n+1}\right)$$

The temperature terms at the node of interest can be grouped together.

$$\frac{\partial T_{(m-2)n+1}}{\partial t} = -\left[\frac{2h_r}{\rho c \Delta x} + \frac{2\alpha}{(\Delta x)^2} + \frac{2\alpha}{(\Delta r)^2}\frac{\left(1+\frac{\Delta r}{2r_I}\right)}{\left(1+\frac{\Delta r}{4r_I}\right)} + \frac{2h}{\rho c \Delta r\left(1+\frac{\Delta r}{4r_I}\right)}\right]\left(T_{(m-2)n+1}\right)$$

$$+ \frac{2h_r}{\rho c \Delta x}T_\infty + \frac{2\alpha}{(\Delta x)^2}\left(T_{(m-2)n+2}\right) + \frac{2\alpha}{(\Delta r)^2}\frac{\left(1+\frac{\Delta r}{2r_I}\right)}{\left(1+\frac{\Delta r}{4r_I}\right)}\left(T_{(m-3)n+1}\right)$$

$$+ \frac{2h}{\rho c \Delta r\left(1+\frac{\Delta r}{4r_I}\right)}\left(T_{(m-1)n+1}\right)$$

For i = n, j = m-1

$$\rho c\left(2\pi r_I\left(1+\frac{\Delta r}{4r_I}\right)\frac{\Delta r}{2}\frac{\Delta x}{2}\right)\frac{\partial T_{(m-1)n}}{\partial t}$$

$$= h_r 2\pi r_I\left(1+\frac{\Delta r}{4r_I}\right)\frac{\Delta r}{2}\left(T_\infty - T_{(m-1)n}\right)$$

$$+ k2\pi r_I\left(1+\frac{\Delta r}{4r_I}\right)\frac{\Delta r}{2}\frac{\left(T_{(m-1)n-1} - T_{(m-1)n}\right)}{\Delta x}$$

$$+ k2\pi\left(r_I + \frac{\Delta r}{2}\right)\frac{\Delta x}{2}\frac{\left(T_{(m-2)n} - T_{(m-1)n}\right)}{\Delta r} + h2\pi r_I\frac{\Delta x}{2}\left(T_{mn} - T_{(m-1)n}\right)$$

The equation is simplified to the following form.

$$\frac{\partial T_{(m-1)n}}{\partial t} = \frac{2h_r}{\rho c \Delta x}\left(T_\infty - T_{(m-1)n}\right) + \frac{2\alpha}{(\Delta x)^2}\left(T_{(m-1)n+1} - T_{(m-1)n}\right)$$

$$+ \frac{2\alpha}{(\Delta r)^2}\frac{\left(1+\frac{\Delta r}{2r_I}\right)}{\left(1+\frac{\Delta r}{4r_I}\right)}\left(T_{(m-2)n} - T_{(m-1)n}\right)$$

$$+ \frac{2h}{\rho c \Delta r\left(1+\frac{\Delta r}{4r_I}\right)}\left(T_{mn} - T_{(m-1)n}\right)$$

The temperature terms at the node of interest can be grouped together.

$$\frac{\partial T_{(m-1)n}}{\partial t} = -\left[\frac{2h_r}{\rho c \Delta x} + \frac{2\alpha}{(\Delta x)^2} + \frac{2\alpha}{(\Delta r)^2}\frac{\left(1+\frac{\Delta r}{2r_I}\right)}{\left(1+\frac{\Delta r}{4r_I}\right)} + \frac{2h}{\rho c \Delta r\left(1+\frac{\Delta r}{4r_I}\right)}\right]\left(T_{(m-1)n}\right)$$

$$+ \frac{2h_r}{\rho c \Delta x}T_\infty + \frac{2\alpha}{(\Delta x)^2}\left(T_{(m-1)n-1}\right) + \frac{2\alpha}{(\Delta r)^2}\frac{\left(1+\frac{\Delta r}{2r_I}\right)}{\left(1+\frac{\Delta r}{4r_I}\right)}\left(T_{(m-2)n}\right)$$

$$+ \frac{2h}{\rho c \Delta r\left(1+\frac{\Delta r}{4r_I}\right)}\left(T_{mn}\right)$$

A.2 Node Equations for a Temperature Boundary Condition

For setting a temperature boundary condition at one end of the pipe, the equations are the same except for the pipe edge when $i = n$. At these edges, the only parts of the equations that change are when the h_r and T_∞ terms are present. In these instances, the following substitutions are made:

$$h_r = \frac{k}{\left(\frac{\Delta x}{2}\right)} = \frac{2k}{\Delta x}$$

$$T_\infty = T_B$$

where T_B is the boundary temperature. In this relation, the nodes at $i = n$ are located at $= \frac{\Delta x}{2}$. In this instance, the ***B*** term in the state space model will be populated by the relations that include a T_B term. The input vector, u, will contain the values of the temperature. For example,

For i=n, j = 2:m-2

$$B_{jn} = \frac{4\alpha}{(\Delta x)^2}$$

$$u_{jn}(t) = T_B(t)$$

A.3 Incorporating Air into the Model

In adding the effects of air heating up inside the test section of the pipe, the following assumptions were made:

- Between time steps, the temperature of the air at each node remains constant.
- The effects of conduction of the air between nodes upstream and downstream are negligible.
- At each node, air is well mixed.

For each time step, the temperature of control volume k at time $t + \Delta t$ is equal to the temperature of the control volume that is n_s control volumes earlier plus the

temperature increases as the volume of air passes through the pipe. Visually, this is displayed in Figure 93. Between time steps, the control volume travels n_s time steps and heats up during that section. For the final ending node, the volume of air may start between two nodes, and this is accounted for in the derivation of the equations.

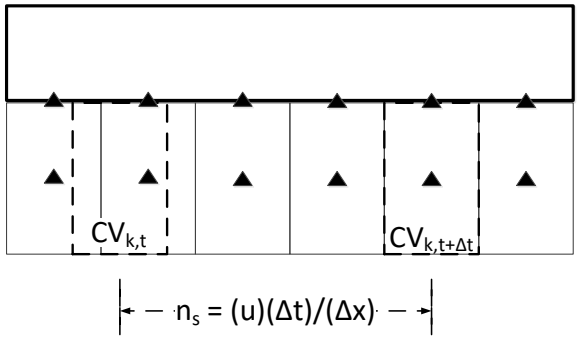

Figure 93. Illustration of how the control volumes are selected for the air nodes in the pipe model.

Note that although the previous equations are derived in continuous state space, the actual implementation and simulation occurs in discrete time, so the method here can be applied. For the control volume (CV) k of air, the number of nodes the CV travels in one time step is equal to:

$$n_s = \frac{u\,\Delta t}{\Delta x}$$

The amount of time that the CV k spends at each node is equal to:

$$t_s = \frac{n_s}{\Delta t} = \frac{u}{\Delta x}$$

Let f_s be the fractional part of n_s, $fl(n_s)$ be the integer part of n_s, and $cl(n_s)$ be the ceiling of n_s. Also, let $c = t_s \frac{hA_i}{p_a c_{p,a} V}$. Then the temperature of the CV k which is located at node $n(m-1)+i$ at the time $t + \Delta t$ is given by:

$$\begin{aligned} T_{n(m-1)+i,t+\Delta t} = {} & (1-c)^{fl(n_s)}\left(1 - c\, fr(n_s)\right) T_{n(m-1)+i-n_s,t} \\ & + (1-c)^{fl(n_s)} fr(n_s) c\, T_{n(m-2)+i-cl(n_s),t} \\ & + \sum_{q=1}^{fl(n_s)} (1-c)^{fl(n_s)-q} c\, T_{n(m-2)+i-fl(n_s)+q-1,t} \end{aligned}$$

Note that the summation is only used when $fl(n_s) \geq 1$. In this equation, the temperature at the starting location of the control volume is given by the following.

$$T_{n(m-1)+i-n_s,t} = (f_s)\left(T_{n(m-1)+i-cl(n_s),t}\right) + (1-f_s)\left(T_{n(m-1)+i-fl(n_s),t}\right)$$

Since the state space form is set up in terms of the temperature derivative at the specific location, the temperature relation is rearranged.

$$\frac{T_{n(m-1)+i,t+\Delta t} - T_{n(m-1)+i,t}}{\Delta t}$$

$$= -\frac{1}{\Delta t} T_{n(m-1)+i,t} + \frac{1}{\Delta t}(1-c)^{fl(n_s)}\left(1 - c\, fr(n_s)\right)(f_s)\left(T_{n(m-1)+i-cl(n_s),t}\right)$$

$$+ \frac{1}{\Delta t}(1-c)^{fl(n_s)}\left(1 - c\, fr(n_s)\right)(1-f_s)\left(T_{n(m-1)+i-fl(n_s),t}\right)$$

$$+ \frac{1}{\Delta t}(1-c)^{fl(n_s)} fr(n_s) c\, T_{n(m-2)+i-cl(n_s),t}$$

$$+ \frac{1}{\Delta t}\sum_{q=1}^{fl(n_s)}(1-c)^{fl(n_s)-q} c\, T_{n(m-2)+i-fl(n_s)+q-1,t}$$

APPENDIX B

MODEL DERIVATION OF TEMPERATURE PROFILE WITH NO CONVECTION

In the equation, the time and space terms both appear in the error function terms. As the convection term approaches zero, the term on the left side of the equation approaches infinity while the error function initially starts at zero. But as time increases, the absolute value term of the function increases but remains finite. Therefore, this model is only valid for b>0. A separate model must be used when comparing the temperature profile for a no-convection case. Appendix B shows the model derivation and solution for the no-convection case.

This model determines the temperature profile for a constant heat flux if the inside surface of the pipe was insulated. This model is still influenced by heat capacity and thermal conduction, but the difference in the profile can be used to determine the effect of convection on the temperature profile.

The setup and evaluation of the analytical model for this case is similar to the analytical model developed in Chapter 3. In the development of this model, all of the assumptions are the same as the previous model. In this instance, the setup of the model matches Equation 6 without the convection term. The boundary and initial conditions are the same as the previous model.

$$\frac{\partial u(x,t)}{\partial t} = \alpha \cdot \frac{\partial^2 u(x,t)}{\partial x^2}$$

The boundary and initial conditions, listed below, are the same as was seen in the previous model.

Boundary Conditions:

$$1.\ \frac{\partial u(0,t)}{\partial t} = g(t) = -Q_c$$

$$2.\ u(x \to \infty, t) < \infty$$

Initial Condition

$$u(x,0) = f(x) = 0$$

B.1 Analytical Derivation with No Convection

When solving the differential equation, the resulting temperature profile is defined by the following [62].

$$u(x,t) = \int_0^{\infty} \frac{1}{2\sqrt{\pi\alpha t}} \left\{ e^{\frac{-(x-\sigma)^2}{4\alpha t}} + e^{\frac{-(x+\sigma)^2}{4\alpha t}} \right\} \cdot f(\sigma)\, d\sigma + Q_c \sqrt{\frac{\alpha}{\pi}} \int_0^{t} \frac{g(\tau)}{\sqrt{t-\tau}} \cdot e^{\frac{-x^2}{4\alpha(t-\tau)}}\, d\tau$$

With the initial temperature distribution at the ambient temperature and using a substitution variable, the resulting temperature distribution is the following (found reference in the partial diff for engineers book) [132].

$$u(x,t) = Q_c\left[-\sqrt{\frac{4\alpha t}{\pi}}e^{\frac{-x^2}{4\alpha t}} + x\left(1 - \operatorname{erf}\left(\sqrt{\frac{x^2}{4\alpha t}}\right)\right)\right]$$

This simplified model still has similar terms to those found in the model with conduction, but this analytical modeler is overall simpler with the fewer terms.