2016-12-01

Extending Time Until Failure During Leaking in Inflatable, Pneumatically Actuated Soft Robots

Joshua Parker Wilson
Brigham Young University

Extending Time Until Failure During Leaking in Inflatable, Pneumatically Actuated Soft Robots

Joshua Parker Wilson

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Master of Science

Marc D. Killpack, Chair
Mark B. Colton
Timothy W. McLain

Department of Mechanical Engineering

Brigham Young University

August 2016

ABSTRACT

Extending Time Until Failure During Leaking in Inflatable,
Pneumatically Actuated Soft Robots

Joshua Parker Wilson
Department of Mechanical Engineering, BYU
Master of Science

Soft robots and particularly inflatable robots are of interest because they are lightweight, compact, robust to impact, and can interact with humans and their environment relatively safely compared to rigid and heavy traditional robots. Improved safety is due to their low mass that results in low-energy collisions and their compliant, soft construction. Inflatable robots (which are a type of soft robot) are also robust to impact and have a high torque to weight ratio. As a result inflatable robots may be used for many applications such as space exploration, search and rescue, and human-robot interaction. One of the potential problems with inflatable or pneumatically actuated robots is air leaking from the structural or actuation chambers.

In this thesis methods are demonstrated to detect leaks in the structural and actuation chambers of inflatable and pneumatically actuated robots. It is then demonstrated that leaks can be slowed by lowering a target pressure which affects joint stiffness to prolong the life of the system. To demonstrate the effects of lowering the target pressure it is first shown that there exists a trade-off between the commanded target pressures at steady-state and the steady-state error at the robot end effector under normal operation. It is then shown that lowering the target pressure (which is related to stiffness) can extend the operational life of the system when compressed air is a limited resource. For actuator leaks a lower target pressure for the leaking joint is used to demonstrate the trade-off between slowing the leak rate and system performance. For structural leaks a novel control algorithm is demonstrated to lower target pressure as much as possible to slow the leak while maintaining a user specified level of accuracy.

The method developed for structural leaks extends the operational life of the robot. Long-term error during operation is decreased by as much as 50% of the steady-state error at the end effector when compared to performance during a leak without the control algorithm. For actuation leaks in a joint with a high-torque load the possibility of a 30% increase in operation time while only increasing steady-state error by 2 cm on average is demonstrated. For a joint with a low-torque load it is shown that up to a 300% increase in operation time with less than 1 cm increased steady-state error is possible.

The work presented in this thesis demonstrates that varying stiffness may be used to extend the operational life of a robot when a leak has occurred. The work discussed here could be used to extend the available operation time of pneumatic robots. The methods and principles presented here could also be adapted for use on other types of robots to preserve limited system resources (e.g., electrical power) and extend their operation time.

Keywords: failure mitigation, leak detection, mpc, soft robotics, inflatable robotics

ACKNOWLEDGMENTS

I would first like to thank Dr. Marc Killpack for all of his hard work and assistance to me in my work as a graduate student. I would like to thank all the other faculty and staff at BYU who have been so helpful. Appreciation also goes to the people at NASA who funded and supported this research through NASA Space Technologies Research Fellowship NNX15AP40H and NASA Early Career Faculty Grant NNX14A051G. I would also like to thank the people at Pneubotics and Otherlab who supplied the robotic hardware that made all this research possible.

I would also like to thank both of my children for helping me to laugh and smile after the many stressful days of graduate school. Finally and most importantly I would like to thank my amazing and supportive wife Stephanie for her consideration and support through this whole process. I couldn't have done it without her and without her it wouldn't have been worth it anyway.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1. INTRODUCTION

Soft and inflatable robotics is a new and emerging area in the field of robotics. Inflatable robots offer several benefits over traditional robotics including being lightweight, compact (when deflated), and robust to impact. These benefits make inflatable robots interesting for purposes such as space exploration where weight and volume come at a high price. Inflatable robots also typically are made of elastic or fabric based materials which means that the robot structure has a very high surface and structural compliance which allows for safe interaction with humans and delicate environments because the risk of high energy impacts is greatly reduced. When an inflatable robot utilizes pneumatic actuation then the joints are also compliant (due to the inherent compliance of compressed air) which further increases the level of safety for robot-environment or robot-human interactions.

In addition to the increased level of safety for the environment around the robot the compliance in the structure and joints of the robot also results in increased safety for the robot itself. A typical motor joint or rigid link may be damaged during an unplanned collision. Motors can be stressed unduly by being driven backwards or forced to suddenly stop, gearing can be stripped, and rigid links can crack or buckle. Compliant mechanisms (like the inflatable links of our robot), in contrast can absorb impacts and recover with minimal damage. The robustness of inflatable robots to blunt impacts makes them perfectly suited to environments and tasks that would normally be ill suited for a more traditional robot. Tasks such as search and rescue, delicate excavation, exploration, and most tasks in unmodeled environments will certainly benefit from the robustness to impact and contact that is possible with inflatable robots.

Robots in many applications such as space exploration or search and rescue are resource-limited. Often, there is a relatively limited opportunity for renewing these finite resources and as a result resource conservation is important and methods that can improve efficiency and resource conservation are desirable. For an inflatable soft robot, one key resource is compressed air. An

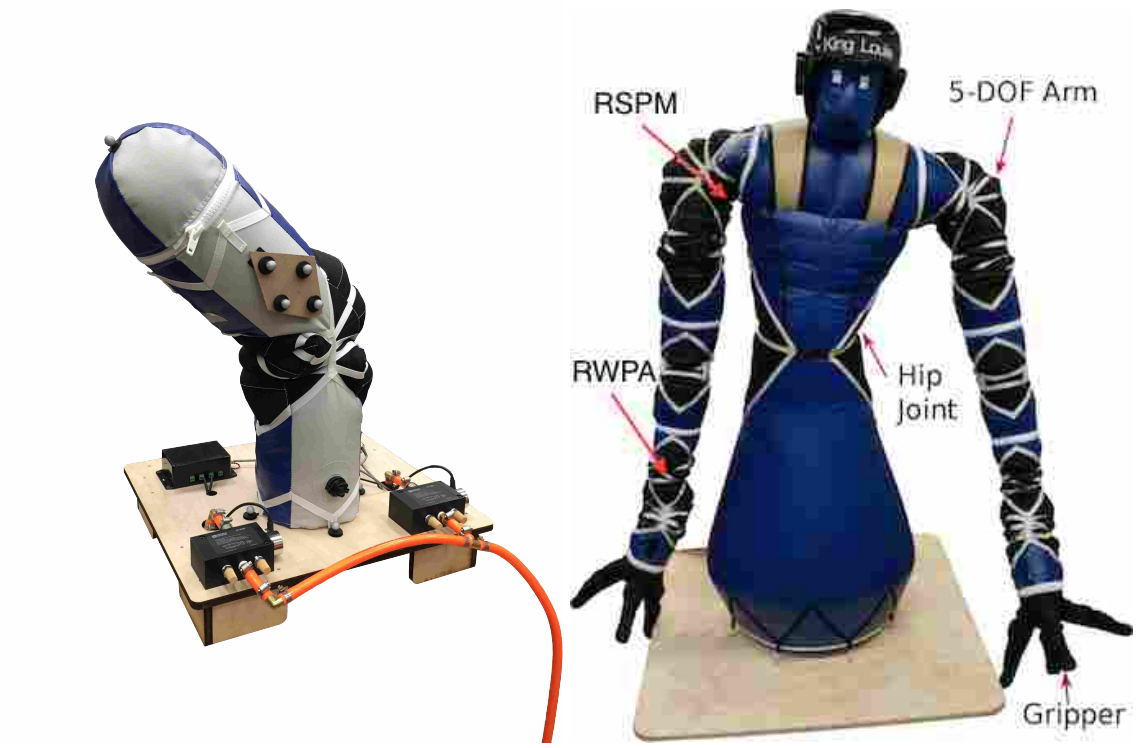

Figure 1.1: Image showing our inflatable robots. On the left is a single degree of freedom (DoF) robot called a ”Grub”. On the right is a humanoid robot with two 5-DoF arms and an additional joint at the hip. The humanoid robot is named ”King Louie”.

inflatable robot used for space exploration would use a closed pneumatic system to recover spent air during actuation but a leak would vent air to the environment that would be difficult or impossible to recover in a thin or no-atmosphere environment. Repair of a leak is time consuming and in many applications such as space exploration or search and rescue operations, it may be infeasible or at the very least must be performed at a later time. As a result methods that prolong the utility of the robot system before physical intervention is necessary to fix the leak are needed.

The robot platform that we use in this work (see Figure 1.1) relies on air pressure both for actuation and for structure. Each joint in the system has two antagonistic joints with the structural chamber running the entire length of the robot arm or link. This unique design means that there is a potential for leaks to occur in both the structural and actuation chambers. The goals of this research were as follows:

- Develop methods to accurately detect the presence of leaks in both the structural and actuation chambers of the robots.

- Develop an algorithm to vary stiffness in order to extend longevity of the robot when a structural leak has occurred.
- Demonstrate the effect of varying joint stiffness on leak rate in structural and actuation chambers of the robots.

The inflatable structure of our robots makes them safe for human interaction because impacts are low-energy due to the low mass of the arm. However, the inflatable structure can leak, which results in the arm losing rigidity and results in a high link compliance that can prevent the robot from actuating accurately. Figure 1.2 shows a comparison between a commanded set of joint angles when the structure is empty and full. The empty arm has a large amount of error and is unable to actuate correctly due to the low pressure in the structural chamber that has occurred as a result of the leak. Structural leaks typically result from small punctures in the structural bladder that occur as a result of pinching and friction between the bladder and the hoses and instrumentation during actuation.

Because of how the system is designed actuator leaks do not typically affect system performance negatively (like structural leaks do) provided the leak is small and all instrumentation is still correctly connected (leaks may occur due to pressure sensing hoses disconnecting). However, actuator leaks can grow over time so detecting even small leaks is valuable. Additionally many future uses for these robots such as space exploration would rely on a limited air supply that would be recycled by the actuation system. If air is lost to the environment due to an actuator leak then it cannot be recycled and the system will eventually lose enough air to severely limit its mobility and functionality. In our current hardware actuator leaks typically occur as a result of loose or leaky fittings and are rarely the result of a puncture in the actuation bladder.

Due to the loss of limited system resources and performance degradation that happens as a result of actuator and structural leaks we have developed methods that detect the presence of a leak in the actuation or structural chambers of our robots. We then demonstrate an algorithm to slow the leak rate from a structural leak while maintaining a commanded level of accuracy. Finally we demonstrate the trade-off for actuator leaks between leak rate and accuracy during actuation. We show that acceptable performance can be maintained while using a lower joint actuation pressure (which is correlated with joint stiffness) to slow the mass flow rate through the leak orifice to the

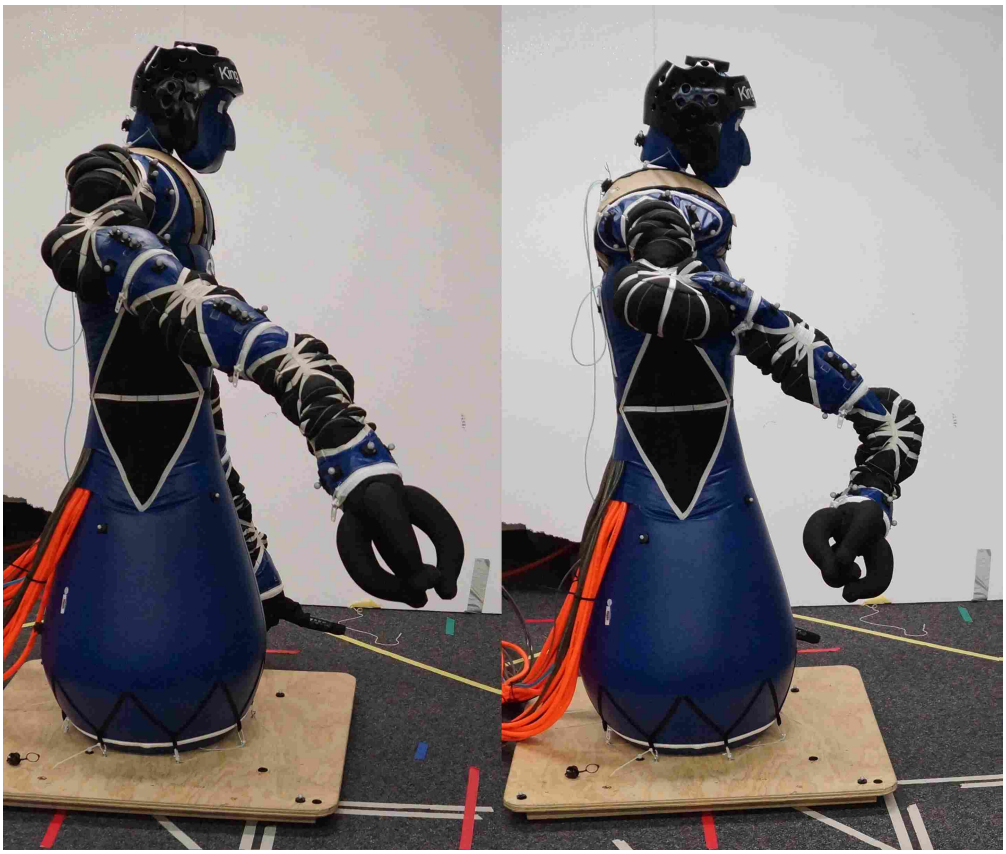

Figure 1.2: On the left is a commanded movement while the arm has a normal structural pressure. On the right is the same motion when the structural chamber has been leaking and is nearly empty. The arm is sagging significantly. The error at the wrist is caused by the motion capture system being unable to track the marker set at the end effector because the arm has deformed due to the low structural pressure.

environment. In this thesis performance and leak data are shown for both the structural chamber and actuation chambers of an inflatable robot but the algorithm developed was only tested on structural leaks. The results for actuation leaks are limited to leak detection and identifying the trade-off between actuation accuracy and average leak rate as joint stiffness is varied.

The following chapters outline the methods we have developed and implemented. First in Chapter 2 we discuss other relevant work in the field and how the work in this thesis is different and adds a unique contribution to the field. Then we discuss the robots themselves as well as their design and control in Chapter 3. Chapter 4 describes the methods we have developed to detect leaks in the structural chamber and the algorithm we have developed to slow structural leaks. Finally Chapter 5 contains the method we developed to detect leaks in the actuation chambers of our robots as well as

an analysis of the trade-off between leak rate from an actuator leak and system performance as the target pressure is changed.

CHAPTER 2. RELATED WORK/LITERATURE REVIEW

Previous work has been done in many areas related to our work. In this chapter we discuss the work similar or relevant to what we have done. We discuss prior work in soft robotics, pneumatic actuation, failure detection, MPC, adaptive control, variable-stiffness actuation, and SVM.

Soft/Inflatable robotics and pneumatic actuation work is relevant because that is the area of robotics to which we are directly contributing and from which we have drawn inspiration. Previous work in inflatable robotics is particularly relevant because the hardware and system behaviour in other inflatable robots can be similar to the hardware and behaviour of our systems. Previous work in adaptive control and failure detection is relevant because adaptive control is used to adapt systems to changing conditions such as failure. This work produces methods that detect leaking (a form of system failure) and (in the case of structural leaks) adjust control inputs to mitigate leaking so naturally other work that is similar is relevant. Prior work in MPC is relevant because our robots are controlled using MPC and prior work has established the viability and effectiveness of MPC for controlling robots. Variable-stiffness actuation work is relevant because our robots have variable-stiffness actuators which are utilized in the work presented here. Prior work in SVM is relevant because SVM libraries developed by others were used in our work. The most relevant area to the work presented here is the work in inflatable robotics and adaptation to failure in robots. Interestingly these are also the areas where we have found the least amount of relevant prior work.

2.1 Soft/Inflatable Robotics

The field of soft robot manipulation and control is a new and emerging area. Work in soft robotics is typically centered around materials and construction methods such as the work done by [1], [2], and [3]. Some work on control and modeling is also being done such as the work by [4], [5], and [6] and recent research in particular has begun to demonstrate the potential benefits of soft robots [1, 7, 8]. Benefits of some classes of soft robots include their high compliance and low

inertia that allow them to operate relatively safely around people [1] or in delicate environments. These characteristics of low inertia and significant compliance due to the pneumatic actuation and inflatable structure also make modeling and control of soft robots more difficult than many traditional robotic systems because of the difficulty in modeling the soft robot dynamics and controlling robots with significant compliance [9, 10].

An excellent review of the breadth of current work in soft robotics is presented in [8]. Work at BYU in the RaD Lab has focused primarily on the control and modeling of inflatable, fabric-based robots which has been demonstrated to be useful in allowing a soft robot to repeatably perform a specific task such as a pick and place operation [7, 11].

2.2 Pneumatic Actuation

Much of the work done in robotics with pneumatic actuation has been done using Pneumatic Muscle Actuators (PMAs) which typically seek to imitate human muscles in some way. A commonly used variety of PMA is the McKibben tendon. McKibben tendons are based on the idea that a flexible cylinder when filled with pressurized gas will expand radially and contract laterally. The shortening that occurs in McKibben tendons when they are filled can be used to perform actuation due to the tensile force created when it contracts laterally. An example of this type of actuation is shown in Figure 2.1. Examples of work with PMAs can be found in [4, 12, 13]. Drawbacks of many PMAs, especially McKibben tendons, is that they require high pressures and the range of motion can be very limited since the maximum decrease in length is limited.

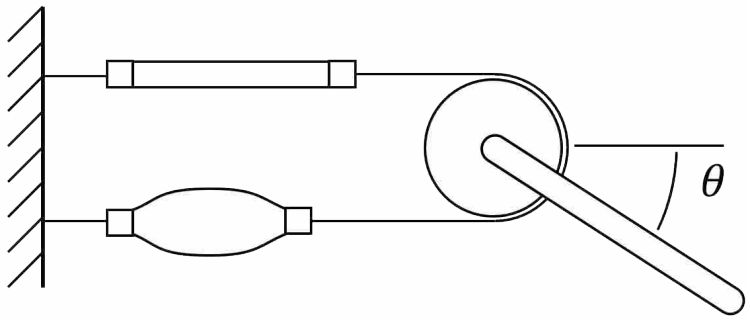

Figure 2.1: Actuation of a McKibben tendon. As the lower bladder is pressurized it expands outward which results in a decrease in length which in turn exerts a force on the joint resulting in the angle θ.

Other pneumatic actuators typically use the pressure of the air to expand and use a compressive force for actuation (in contrast to the tensile force from McKibben tendons). Examples of this include the actuators on our robots (see Figure 3.2) as well as the flexible rotary joints described in [14] and the rotary elastic chambers described in [15, 16].

Because the actuation system in our robots is very different from common rigid fluidic systems we have not discussed actuation systems that involve a rigid cylinder. However, rigid systems can be considered a precursor to more novel pneumatic systems similar to our robot actuators.

2.3 Failure Detection and Mitigation

Work in inflatable soft robotics is centered around control and modeling; failure has not been heavily considered. Some research into failure of pneumatic robots has been done by [13] but their research centered around McKibben artificial muscle actuators failing due to fatigue. In contrast our actuators are more similar to the rotary elastic chambers described in [15, 16]. The failure modes we analyze are more likely to occur from a leaky fitting or a pinhole leak which in our experience are more common for this type of robot than fatigue failure.

Research has been demonstrated by [17] to minimize air consumption in pneumatic actuation in robots but this work is primarily concerned with valve control systems that prevent excess air from being lost and does not deal with leakage or with changes in joint stiffness. Some work on adaptive control of antagonistic pneumatic systems was done by [12] but this work is centered around position control rather than failure. Research has been done by [18] and [19] into using stiffness control in adaptive control but this work focused on improving system performance and stability rather than adjusting system inputs to mitigate system failure. The robot systems used in other work are also very different from our system. Some work on stiffness control for pneumatic systems was demonstrated in [20] using a pneumatic system consisting of rigid pneumatic chambers rather than the soft inflatable systems used in our research.

Additional work has been done related to failure detection and mitigation in robotics or other fields. Much work in failure detection in robotics centers around software or algorithm failure as in the work done in [21]. Outside of robotics there is work on identifying hardware failure in industrial systems such as that shown in [22]. Other failure detection work for industrial systems

such as the gearbox failure analysis shown in [23] is relatively common although such research is often kept internal by the corporation funding the study. Studies such as these are very different from the work presented here because of the difference in the system. We have been unable to find any prior failure detection or mitigation work related to pneumatic or inflatable robotic systems similar to ours and feel confident that this work is both novel and useful.

2.4 Model Predictive Control

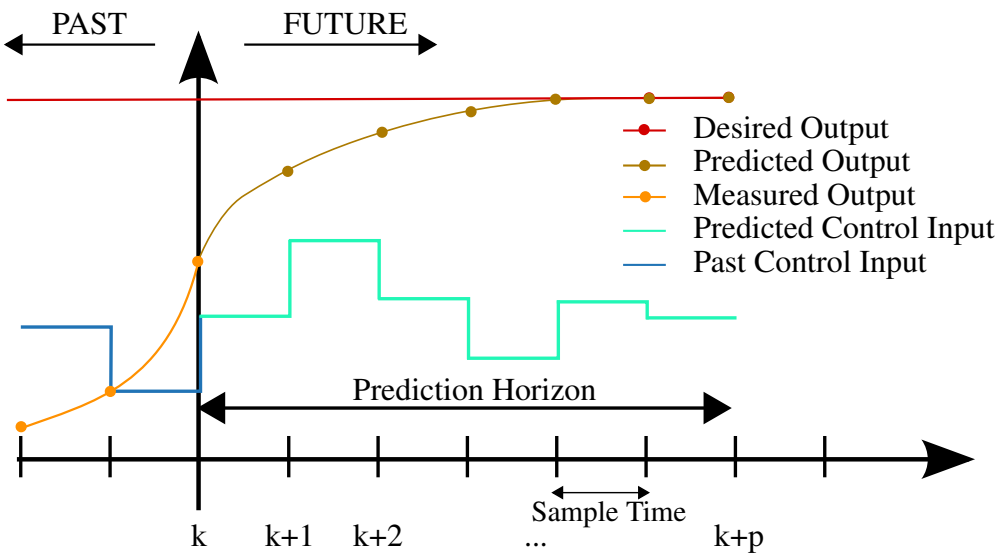

Figure 2.2: A figure showing the basic flow for a Model Predictive Controller. Figure By Martin Behrendt - Own work, CC BY-SA 3.0, Linked here.

Model Predictive Control (MPC) is a method of control wherein future states are predicted using a system model and a series of future inputs is determined by optimizing to minimize a cost function. The first calculated input is then used and the process is repeated. An example of MPC inputs and outputs can be seen in Figure 2.2. As the figure shows predicted outputs are determined from the system model and future inputs are optimized to bring the predicted output as close to the desired output as possible.

The inflatable robots used for this research are controlled using a MPC described in [24]. The MPC cost function has two main cost terms. The first cost is on the accuracy of commanded joint positions and decreases as the measured joint angles approach the commanded values. The

second cost is incurred when the commanded actuation pressure in both actuation chambers for a joint are not at the commanded target pressure. The target pressure allows us to control stiffness. A commanded angle can be reached with a wide range of target pressures with a higher target pressure resulting in a higher stiffness at the joint. In this thesis we will use the term target pressure (P_{nom}) because that is the value we actually control but it is important to understand that the target pressure we command to our controller is directly correlated with joint stiffness. A more in-depth discussion of joint stiffness in these robots is presented in [24, 25].

MPC has traditionally been used primarily in the chemical processing industry [26] but with modern computing speeds and new methods such as those demonstrated in [27, 28] a greater body of research is now using MPC as a high bandwidth real-time control method for robotics [7, 9, 29–31].

2.5 Adaptive Control

Adaptive control as a field has many standard methods, many of which are well documented in [32]. More recently work has been done to integrate various adaptive control methods with a model predictive controller [33, 34]. Some work has been done with adaptive control in the field of pneumatic robotics such as the work done by [12] and the failure analysis work done by [13] but this has primarily related to Mckibben actuators which, while pneumatic, behave very differently from the actuation used in our research. Adaptive position control work has also been done by [6] and while the actuation used is very similar to that in our research they do not use stiffness variation for adaptation (although they do have that capability in their robot arm). In addition the other research discussed here does not use robots with an inflatable structure which also differentiates our research from theirs in terms of decreased overall mass, increased compliance, and possible degrees of freedom. The work presented in this thesis most closely related to adaptive control is in Chapter 4 where a control input is adjusted online to mitigate a leak. Since the control is not directly adjusted the work presented here for structural leak mitigation is not in the strictest sense adaptive control but is similar to adaptive control methods in the end result.

2.6 Variable-Stiffness Actuation

Work has also been done to integrate variable-stiffness actuators and adaptive control techniques. The most closely related work is the research done by [19] on adaptive control of soft robotic manipulators with variable stiffness. The work done by Tonietti et al. differs from our research in that they use Mckibben actuators and use adaptive control for position control rather than adaptation to mitigate failure due to losing limited resources. More recently work has been done to use stiffness variation for adaptive control such as the work done by [18] but there is little to no literature on attempting to use variation in joint stiffness to increase time until system failure or slow the rate of a leak in pneumatic systems.

2.7 Support Vector Machines

Support Vector Machines (SVMs) are used in fields as diverse as computer vision, natural language processing, and data classification. SVMs are widely accepted and used in classifying data by determining where data belongs between multiple distinct data sets. In this thesis we use them for classification of leak detection. More information about SVMs and the library used in this research (libSVM) can be found in [35]. A discussion of how SVM is used in this thesis is included in Chapter 5.

2.8 Contribution of This Work

None of the past research described here attempts to address either the problem of increasing time until complete failure in robots after initial failure has occurred or the question of whether there exists a trade-off between longevity and performance when initial failure has occurred. It was necessary to perform this work not only to establish how adjusting P_{nom} affects leaks and performance in our particular robots but also to introduce the idea of adapting to failure and trading performance for longevity of limited resources into the field of robotics. The problem of leak detection in inflatable robots and pneumatic systems has also not been addressed in the literature. The work presented in this thesis demonstrates methods for leak detection and an algorithm for leak mitigation in the structural chamber of an inflatable robot. Additionally, leak detection is demonstrated for the pneumatic actuators of our robots and a viable trade-off is shown between leak

rate and accuracy of the system. All of these contributions are new to the field of soft and inflatable robotics.

CHAPTER 3. HARDWARE PLATFORM AND BASIC CONTROL

Our robots were designed and manufactured by Pneubotics which is an affiliate of Otherlab. The robots are fabric based and are designed to be low inertia. Two robots were used in this research and are pictured in Figure 1.1. A 1-DoF robot called a Grub was used for initial testing and validation of concepts and a 5-DoF robot called King Louie was the final test and validation platform for this research.

3.1 Robot Design

The inflatable robots we use consist of an internal structural bladder that runs the length of the robot arm linkage and two additional opposing pneumatic chambers at each joint which provide a bellows type actuation for that joint (see Figure 3.1). The actuators are similar to the fluidic actuators described in [2, 14, 16]. The pressures in the actuation chambers and the structural chambers are measured using pressure sensors mounted inside the links of the robots. Pressure readings are communicated to a PID pressure controller using the Robot Operating System (ROS). Commanded pressures for the actuation chambers are set by a model predictive controller (MPC) that receives estimated joint angles from either a motion capture system (for King Louie) or an IMU (for the Grub) and commanded angles from the user. A diagram outlining the control system can be seen in Figure 3.3. The pressure controller runs at a rate of 1000 Hz while the MPC runs at a rate of 300 Hz. For the methods described in this thesis all data was recorded and commands sent at a rate of 300 Hz.

Our robots use antagonistic pressure chambers for actuation where the pressure difference between the two chambers in a joint determines the resultant torque generated in the joint and as a result the acceleration and position achieved during actuation. The robots are shown in Figure 1.1 and the configuration for actuation of a single joint is shown in Figure 3.2. Since the torque produced by the joint is determined by the pressure difference between the two chambers, a variety

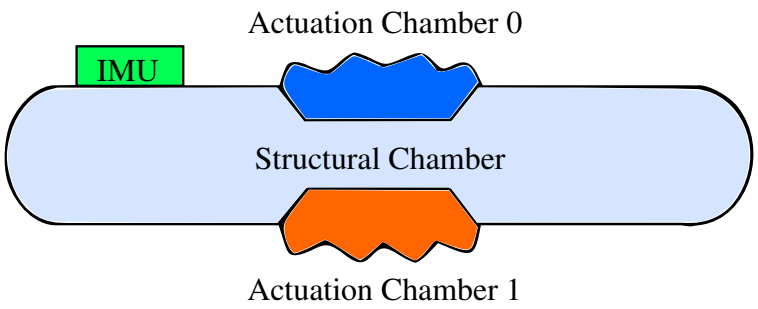

Figure 3.1: Cross sectional view of the Grub showing how the actuator chambers interact with the structural chambers. The actuator chambers expand into the structural chamber as they are filled with air which decreases the volume in the structural chamber and increases structural chamber pressure.

of pressures can result in the same torque since the pressure in both joints may be increased or decreased simultaneously while maintaining the same torque. In addition to the desired position the MPC is also given a desired or nominal pressure (P_{nom}). Lowering P_{nom} effectively lowers joint stiffness since lower pressures in the actuation bladders result in a higher compliance at the joint. Conversely raising P_{nom} will increase joint stiffness. Work done by Charles Best in [25] shows the relationship between the actuation pressures and joint stiffness. This method of stiffness control is similar to the work described in [15] and uses the models and control described in [24].

Each link (separated by at most two actively controlled joints) contains an instrumentation board with an IMU and pressure sensors. The instrumentation board returns measured pressures for each actuation chamber as well as accelerometer, gyroscope, and magnetometer data.

3.2 System Models and Control

Models for the inflatable robots have been and continue to be under development. Many of the members of the RaD Lab have worked on developing models and improving control of these robots and the work is still in progress. As a result multiple models and control systems were used during this research. Different models were used to control King Louie and the Grub and we used newer models in this research as they were developed. When we collected data we used the most advanced model that the lab had developed and tested up to that point. Since testing was performed on the Grub before King Louie, a more advanced model had been developed and tested by the time we took more data and the more advanced model was used for data collection on King Louie. The

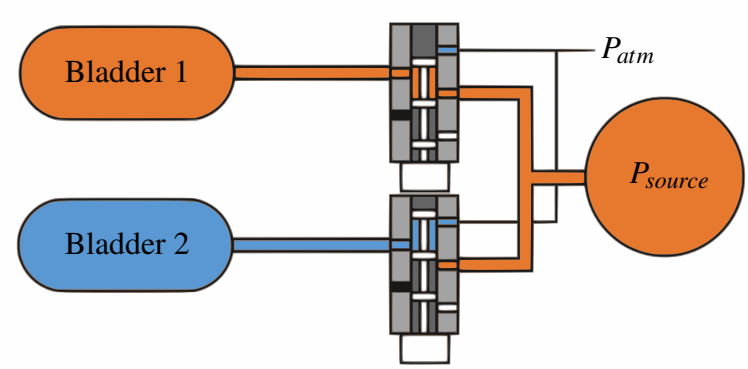

Figure 3.2: Actuation setup for the pneumatic actuation of our robots. The two bladders shown provide antagonistic force which allows for variable stiffness in addition to actuation. Figure used with permission from Morgan Gillespie.

two models that were used for control are included below. Both models perform adequately and since in this thesis we only compare trends between data taken on the Grub and King Louie it does not affect our results that different models were used for the two different robots. It is also of note that the methods described in this thesis operate outside the MPC which means that an improved model can be substituted at any time and it will not affect the results or methods presented in this thesis. The ability to upgrade the model and controller at any time means this work can easily be updated and tested further in the future with different robotic platforms or improved models and control as long as we can still control P_{nom} in the controller formulation.

3.2.1 4-State Model

The first model we used is a 4-state model (Equation 3.1) based on the commanded pressures for each joint as well as the angle and angular velocity of the joint. This model was developed by Morgan Gillespie and Charles Best and is described in detail in [24]. This 4-state model was the primary model used to control the Grub.

$$\begin{bmatrix} \ddot{\theta} \\ \dot{\theta} \\ \dot{P}_0 \\ \dot{P}_1 \end{bmatrix} = \begin{bmatrix} \frac{-K_d}{I} & 0 & \frac{-K_s\alpha_2}{I} & \frac{-K_s\alpha_3}{I} \\ 1 & 0 & 0 & 0 \\ 0 & 0 & -a & 0 \\ 0 & 0 & 0 & -a \end{bmatrix} \begin{bmatrix} \dot{\theta} \\ \theta \\ P_0 \\ P_1 \end{bmatrix} + \begin{bmatrix} \frac{-K_s\alpha_2}{I} & \frac{-K_s\alpha_3}{I} \\ 0 & 0 \\ b & 0 \\ 0 & b \end{bmatrix} \begin{bmatrix} P_{D,0} \\ P_{D,1} \end{bmatrix} \tag{3.1}$$

$$I\ddot{\theta} + K_d\dot{\theta} + mg\frac{L}{2}sin(\theta) = K_s(\theta_e - \theta) \tag{3.2}$$

$$\theta = \alpha_1 + P_0\alpha_2 + P_1\alpha_3 \tag{3.3}$$

$$\theta_e = \alpha_1 + P_{D,0}\alpha_2 + P_{D,1}\alpha_3 \tag{3.4}$$

The 4-state MPC model given in Equation 3.1 is fairly straightforward and is based on the model given in Equation 3.2 where we neglect the gravity term and θ and θ_e are replaced by Equation 3.3 and Equation 3.4 respectively and α_1 has been incorporated into the θ terms. I is the moment of inertia of the link about the joint center, K_s is the stiffness constant of the structure, θ is the measured joint angle, θ_e is the equilibrium angle obtained from a mapping between joint angle and actuation pressures, P is a vector of pressures in the actuation chambers with P_D indicating desired pressure, K_d is a damping constant, m, g, and L are joint mass, gravity, and distance from joint center to the center of mass. The other parameters a, b, and α are parameters determined using collected data to fit the model to actual measurements taken on the system. For a more in depth explanation regarding this model see [24].

3.2.2 Torque Model

The second model we used is also a 4-state model but is based on a model of the torque at each individual joint (Equations 3.5-3.8). This model was also developed by members of the RaD Lab and greater detail regarding this model is given in [11]. This was the primary model we used for controlling King Louie.

$$I\ddot{\theta} + K_d\dot{\theta} + mg\frac{L}{2}sin(\theta) = \tau_a \tag{3.5}$$

$$\tau_0 - \tau_1 = \gamma_0 P_0 - \gamma_1 P_1 \tag{3.6}$$

$$\tau_{stiffness} = K_s \theta \tag{3.7}$$

$$\tau_a = \tau_0 - \tau_1 + \tau_{stiffness} \tag{3.8}$$

The torque model has two components. The first component is the resultant torque due to the pressures in the two antagonistic chambers (see Figure 3.1) of a joint and the equation for that torque is given in Equation 3.6. The second component (Equation 3.7) is the torque due to the stiffness of the structural bladder which tends to bring the joint back to equilibrium as it exerts some torque at the joint. Equation 3.8 combines these two components to give the resultant torque at the joint as a function of angle and actuation pressures. For a more in depth look at the torque model see [11]. When put in state space form the resultant torque model used for control looks very similar to Equation 3.1. For control both of the models described here are discretized.

3.2.3 Advanced Control

The models described above were used in a MPC paradigm to control the movements of the robots and allowed for accurate angle control and relatively high precision ($\approx$1 cm error on average) at the end effector for repeated movements. The MPC has a multi-objective cost function shown in Equation 3.9 where θ and θ_{goal} are the measured and desired joint angle respectively, P_0 and P_1 are the pressures in the actuation bladders, and the subscripts Q, R, and S are the weightings on the different terms of the cost function. This cost function has cost associated with error in the commanded joint angles as well as P_{nom}. The MPC tries to reach the desired position with actuation pressures as near to the commanded P_{nom} as possible. As a result we can vary the stiffness of a joint by adjusting the nominal pressure (P_{nom}) commanded to the joints. The cost function given in Equation 3.9 is also subject to the constraints given in Equation 3.10.

Cost Function:

$$minimize \sum_{k=0}^{T} \left(\|\theta_{goal} - \theta[k]\|_Q^2 + \|\dot{\theta}[k]\|_R^2 + \|P_0[k] - P_{nom}\|_S^2 + \|P_1[k] - P_{nom}\|_S^2 \right) \tag{3.9}$$

$$\begin{aligned} \dot{x} &= Ax + Bu \\ |\theta| &\leq \theta_{max} \\ P_{min} \leq P_D &\leq P_{max} \\ |\Delta P_D| &\leq \Delta P_{max} \end{aligned} \tag{3.10}$$

The MPC controller is given a goal angle and angular velocity (which currently must be zero) and returns desired pressures for each chamber of the actuator. Pressure values are sent to a pressure controller which uses the desired pressures and measured pressures in a PID controller to send electrical current commands to individual valves connected to each actuation chamber. Angles are measured using either motion capture (for King Louie) or IMU data (for the Grub). IMU data is run through a Kalman filter which then returns filtered angle and angular velocity data to the MPC controller to close the loop. The control diagram is depicted in Figure 3.3.

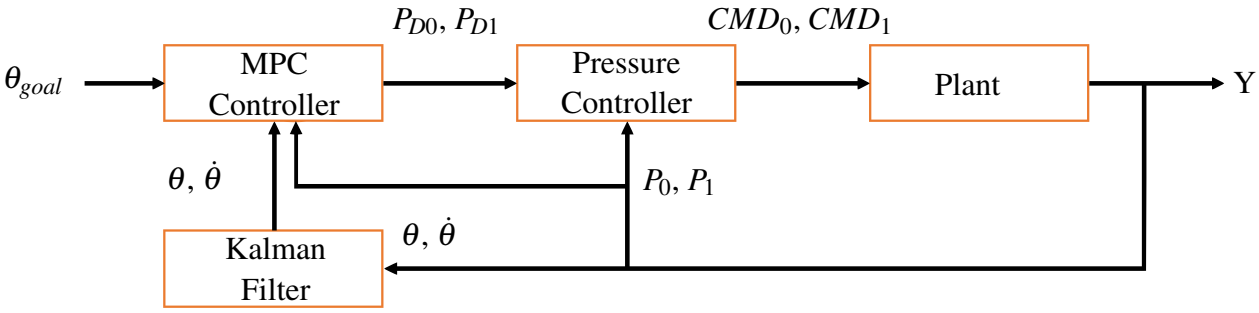

Figure 3.3: Flow chart showing the control flow used to control the inflatable robots.

The information in this chapter has been included to give the reader a greater understanding of the robots used in our research as well as how they operate. Such an understanding is necessary to fully appreciate the remainder of this thesis. For example the interaction between the structural and actuation chambers is used to slow structural leaks and if it is not clear to the reader how the robots are designed this understanding would be lacking. Because of the unique nature of the robots and their control systems we are able to utilize P_{nom} to show how leaks can be slowed in both the structural and actuation chambers of our robots. An understanding of the control is useful

particularly because an understanding of the cost function used in the MPC is vital to understanding how changes in P_{nom} actually affect the pressures used for actuation.

CHAPTER 4. STRUCTURAL LEAKS

The robot platform that we used in this work relies on air pressure, not only for actuation, but also for structure. Each joint in the system has two antagonistic joints with the structural bladder running the entire length of the robot arm or link (see Figure 3.1). The lack of a rigid structure helps with safety because the inertia is very low but it does exacerbate problems caused by low structural pressures due to a leak. In addition to loss of limited system resources (air in this case) a leak in the structural component of the robot will typically result in degraded system performance. When a leak is present in the system, the robot gradually loses stiffness of its limbs and is unable to perform even the most basic task because the joints cannot accurately actuate links with extremely high compliance (i.e. floppy links). Figure 4.4 shows a comparison between a commanded set of joint angles when the structure is full and empty. The empty arm has a large amount of error when compared with the commanded position. In this chapter we first discuss the method used to detect leaks in the structural chamber in Section 4.1 and then we discuss the algorithm we have developed to slow leaks in the structural chamber of our robots in Section 4.2.

4.1 Structural Leak Detection and Estimation

We use two distinct methods to detect when a leak has occurred in the structural chamber of our robots. The first method is for detection of fast leaks and the second is for the detection of slow leaks. For detecting fast leaks the algorithm measures the average structural pressure every 0.1 seconds after the measured angle has reached steady state and determines that a leak has occurred if the structural pressure has decreased for every measurement for 1.6 seconds. The number of data points was tuned empirically to obtain a threshold of 16 data points (or 1.6 seconds of data) of decreasing pressure as the point at which a leak is detected because we found that fewer than 16 points was not robust to avoiding false leak detection due to noise in the data. Measurements must

occur at steady-state because structural pressure fluctuates significantly during actuation due to bending of the chamber and changing actuation pressures that interact with the structural chamber.

The second method uses an internal (structural) pressure map such as the one shown in Figure 4.1 which was generated with a known mass of air inside the robot to determine the relationship between structural pressure, angle, and target pressure (P_{nom}). The map was obtained through measurement of structural pressure as angle and target pressure were varied. Using the internal pressure from Figure 4.1 we can use Equations 4.1 and 4.2 to calculate the volume for a given target pressure and angle. From the calculated volume and the measured internal pressure the mass of air inside the structural chamber can be calculated using fundamental gas equations (included here as Equations 4.1 and 4.2). During initialization of our controller, an average mass is calculated from the first 30 seconds of data by averaging the estimate of internal mass calculated from Equation 4.2 at each time step. The average mass is then calculated every 0.1 seconds and a leak is flagged if the measured mass drops below 90% of the initial value.

$$P_i = \rho_i R T \tag{4.1}$$

$$m_i = \rho_i V \tag{4.2}$$

4.1.1 Validation of Mass Estimation in Structural Leak Detection

To validate our model of the leak based on internal pressure, we used the first principles model given in Equation 4.3. This allowed us to simulate a leak of known size to demonstrate that the model we use for relating pressure to internal mass was accurate. We then calculated the rate of mass change as shown in Equation 4.3. By using a leak of known size we know what the leak rate will be and can show that our simulated model of mass in the structural chamber during a leak matches what we would expect from the leak we induced. Figure 4.2 shows the results.

$$\dot{m} = \alpha \rho_a A \left(\frac{2 P_{ig}}{\rho_a} \right)^{0.5} \tag{4.3}$$

For the simulation of the leak a starting pressure of 10.34 kPa (1.5 psig) was used and the starting density (ρ_i) was calculated using Equation 4.1 where R, T, P_i, and ρ_i are the specific

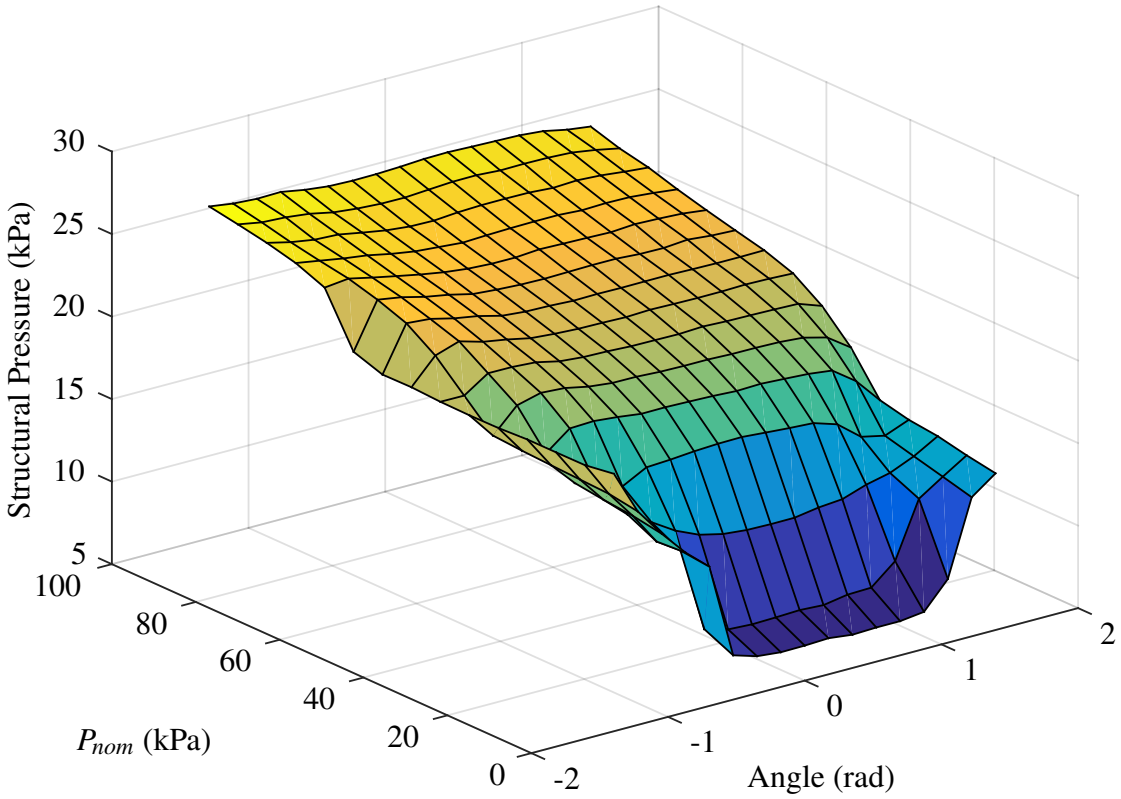

Figure 4.1: This plot shows the relation between internal (structural) pressure, target pressure (stiffness), and angle for the Grub. The surface is used to determine the internal structure volume of the Grub for given operating parameters.

gas constant for air, the temperature, the absolute pressure, and the density of air in the structural chamber respectively. The result from Equation 4.1 was then used in Equation 4.2 to calculate the starting mass of air (m_i) in the structural chamber. Equation 4.3 was then used to calculate $\frac{dm}{dt}$. At the next time step the mass lost ($\Delta T \cdot \frac{dm}{dt}$) is subtracted from the previous mass of air and the new mass is used in Equation 4.1 to calculate a new pressure which in turn is used in Equation 4.3 to calculate a new mass flow and the process then continues for the next time step. In this way the pressure in the structural chamber is simulated with a small ΔT.

If our model of internal mass is correct we expect the simulated structural pressure to closely match the measured pressure during a test on the real system. As can be seen in Figure 4.2 this is the case and demonstrates the accuracy of our estimate of mass of air in the structural chamber of

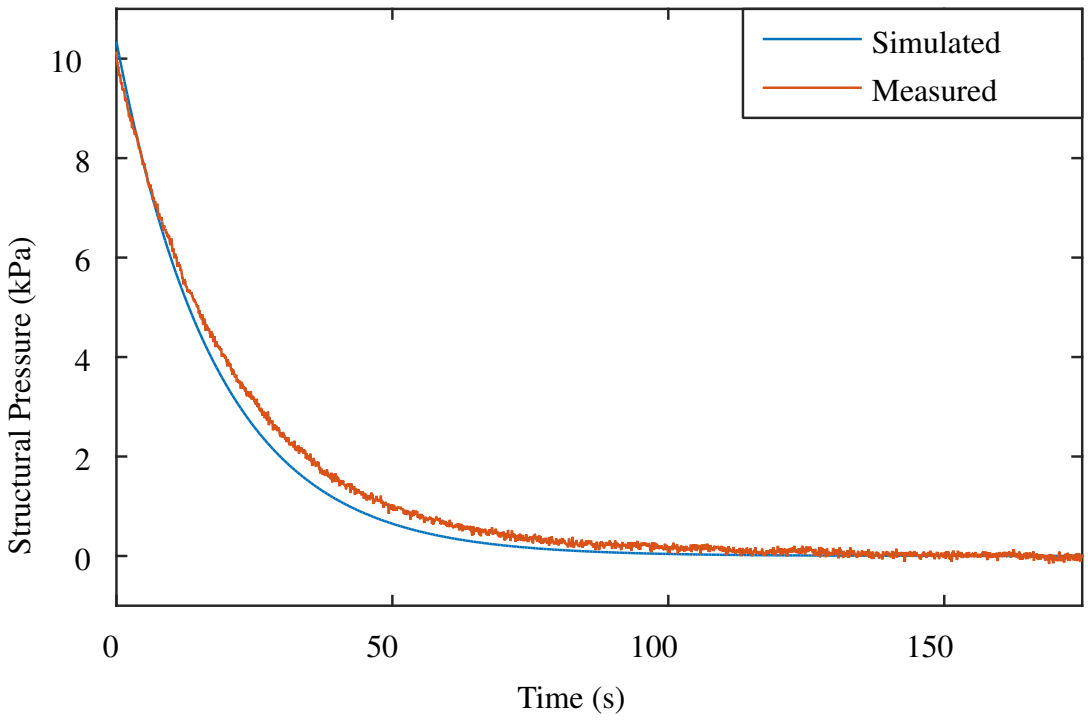

Figure 4.2: Simulated leak for a 1.59 mm ($\frac{1}{16}$ inch) diameter leak compared against actual measured internal pressure data.

our inflatable robots at steady state. For the purposes of this test the robot was not actuated since actuation would alter the structural volume (V). Mass flow was not directly measured for validation of the model because the mass flow rate for our trials is multiple orders of magnitude smaller than the range which most mass flow sensors are capable of measuring.

The data shown in Figure 4.3 show the internal mass estimate of air in the structural chamber of the Grub for two of the trials we performed where a series of angles were sent as step commands. The first case (upper line in blue) shows a data set for which there is no leak. This is included to demonstrate that the mass calculation is consistent over time and for a variety of angles during actuation. The calculated mass varies by less than 10% during any given actuation cycle. The second data set shown on this plot shows the calculated internal mass of air for a case in which a leak is introduced 60 seconds into the actuation cycle.

In the case of the 1-DoF inflatable robot the maximum volume was estimated directly by measuring the external dimensions since its shape makes volume measurement fairly straightforward.

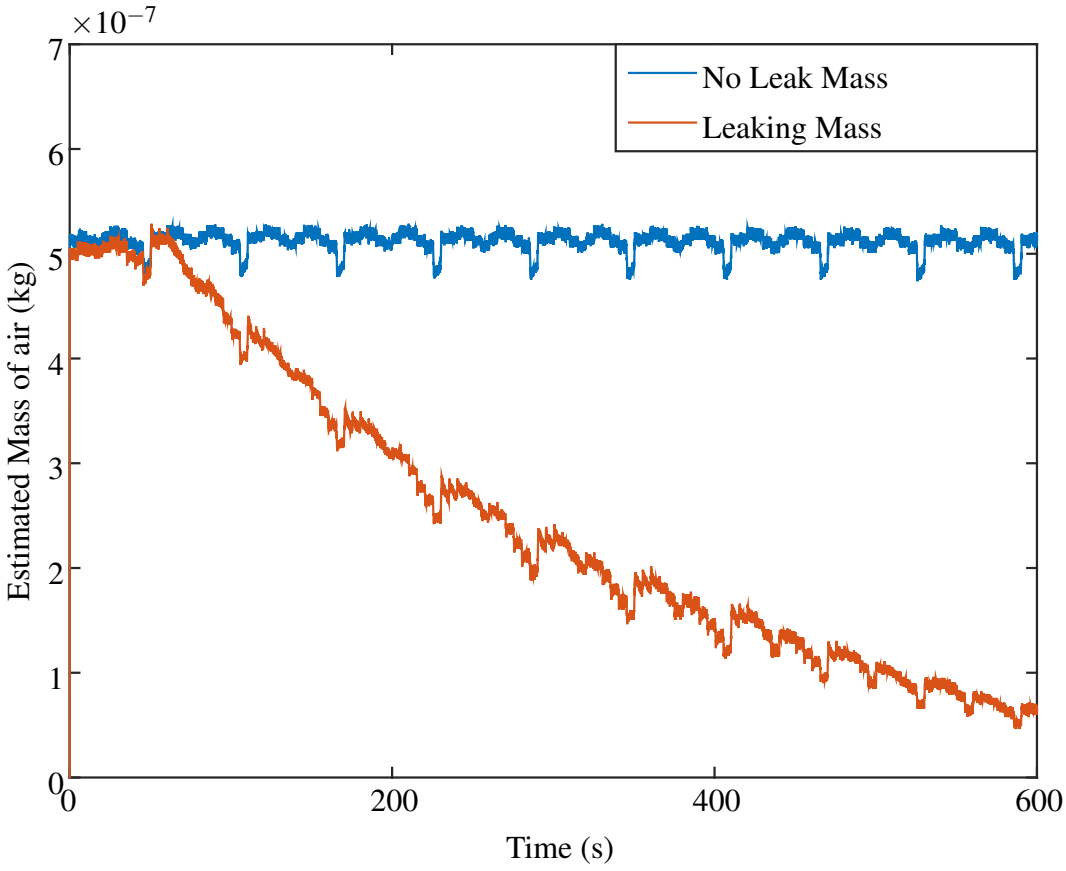

Figure 4.3: Calculated internal mass of air in the structural chamber of the Grub. Data is shown for a case with no leak and another case where a leak was introduced at time = 60 seconds.

In the case of the humanoid inflatable robot the maximum volume of the structural chamber in the arm cannot be easily measured and instead was calculated from data taken during a leak. The accuracy of the simulated leak model (based on the results shown in Figure 4.2) for the 1-DoF robot allows us to infer an accurate volume for the structural chamber of the 5-DoF arm from a leak test similar to the one shown in Figure 4.2 where an optimization was run to find V in Equation 4.2 such that the simulated pressure matched the measured pressure during the leak.

To get consistent results for testing different sized holes, we used artificially induced leaks with two orifice sizes. A one meter long tube was attached to the internal structural chamber of the robot and holes were drilled in end plugs to control the size of the leak. For what we call "fast leaks" in this thesis, a 1.59 mm ($\frac{1}{16}$ inch) hole was used and for what we call "slow leaks" a 0.508 mm ($\frac{1}{50}$ inch) hole was used. This method allowed easy control of the size (and thereby the rate) of the

leak which allowed us to compare leaks from multiple runs consistently. For each run the initial fill pressure of the structural chamber was 10.34 kPa (1.5 psig) before actuation began (the actuation chambers were at atmospheric pressure).

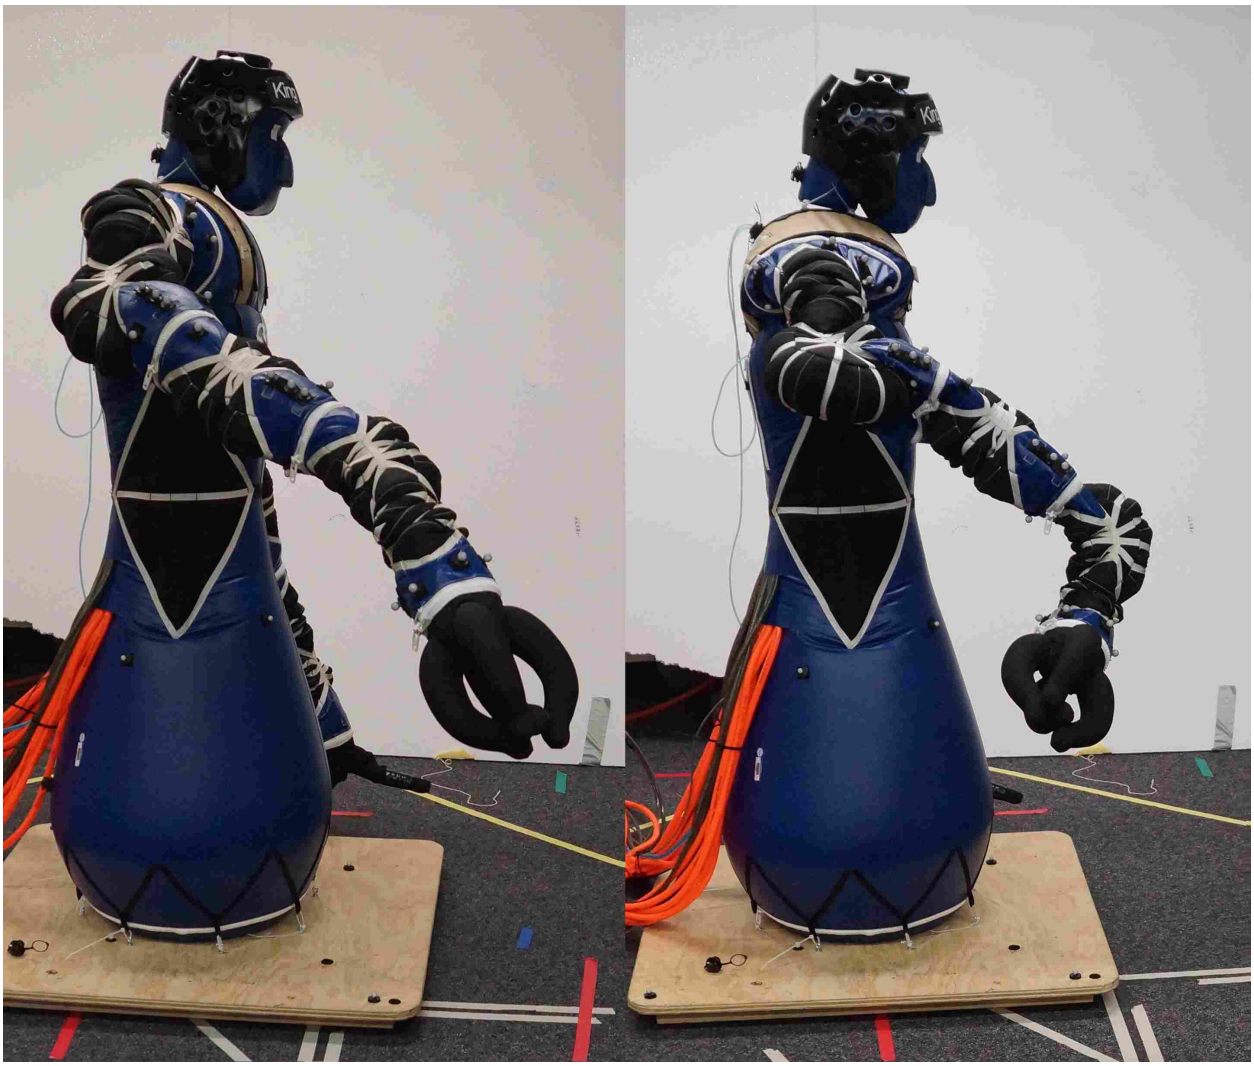

Figure 4.4: On the left is a commanded position while the arm has a normal structural pressure. On the right is the same position when the structural chamber has lost most of its air through a leak.

4.2 Methods for Structural Leak Mitigation

To reduce the rate mass is lost through a leak, the internal pressure in the structure of the inflatable robot must be reduced. This reduction in mass flow can be seen from first principles

gas models shown in Equation 4.3 which is the equation for mass flow rate from a high pressure body (the robot structural bladder) to a low pressure body (the environment) through a hole of fixed cross-sectional area. The variable $\dot{m}$ is the mass flow rate, A is the area of the leak, P_{ig} is the internal gage pressure, ρ_a is the density of the ambient air, and α is a loss coefficient to account for frictional and other losses. As can be clearly seen the only factor that can be easily controlled to affect the mass flow rate out of the leak is the internal gage pressure of the robot structure which can be lowered by lowering the target pressure for the joints which interact with and compress the structural chamber as shown by Figures 3.1 and 4.1.

If the structural pressure is lowered too much, as happens when a leak occurs, then the performance and accuracy of the system decrease due to a very low stiffness in the body as displayed in Figure 4.4. So the ideal method would be to use the lowest possible stiffness that will still result in a high enough structural pressure to give acceptable performance during continued actuation. This would allow the system to operate with acceptable performance for as long as possible before the leak decreases the internal mass of air in the structure to the point where acceptable performance is no longer possible.

As discussed previously the required performance of a soft robot system can vary depending on the task being performed. For example pointing a camera, sweeping dust off solar cells, and many non-prehensile manipulation tasks do not require high accuracy at the end effector whereas insertion of a soil sample into analysis instrumentation would require more accurate manipulation. We implemented an algorithm to monitor system performance based on step response for 1-DoF systems and steady-state error at the end effector in multi-DoF systems and adjust the nominal pressure (stiffness) to maintain the end effector position of the robot to be below a user specified threshold for steady-state error. Steady-state error is defined in this thesis as the sum-squared error in Cartesian space between the current position of the end effector and the achieved position of the end effector during normal operating conditions for the same series of commanded angles. End effector position was measured using a motion capture system.

The result is that for tasks that do not require high precision or for positions that can be reached with low actuation pressures the stiffness will be lowered to slow the leak and extend the life of the system. For higher precision tasks the nominal pressure will be raised by a simple integral controller until the accuracy is within the user specified threshold. In this way the system

can quickly adjust to the requirements of the task being performed while simultaneously extending the life of the system as long as possible by slowing the leak. The completed algorithm is shown in Algorithm 1. For all our tests 4 cm of error was used for the desired level of accuracy. Error of 4 cm was used because it is slightly worse than typical performance under ideal conditions which generally result in 1-3 cm of error (typical performance is discussed in [11]) but is a small enough error that many useful tasks could still be performed.

4.2.1 Algorithm Development and Tests on Grub

The basic philosophy behind our algorithm is straightforward. Lower actuation pressures will result in a lower structural pressure and slow the rate of a leak in the structural chamber. Our initial trials simply lowered P_{nom} which would slow the leak initially but did not lead to improved performance long term because the low stiffness resulted in a poor step response on the Grub. We decided to develop an algorithm that adjusts P_{nom} similar to an integral controller for each command depending on the required performance as specified by the user. The resulting algorithm is shown here as Algorithm 1.

Algorithm 1 Algorithm To Adjust P_{nom}

```
1:  procedure ADJUST P_nom(error, threshold)
2:      while Running do
3:          if error > threshold then
4:              P_nom = P_nom+
5:          else
6:              P_nom = P_nom−
7:      if θ ≠ θ_prev then
8:          if θ is in dictionary then
9:              P_nom = getFromDictionary(θ, P_nom)
10:         if θ_prev is not in dictionary then
11:             appendToDictionary(θ_prev, P_nom)
12:     return P_nom
```

Algorithm 1 is fairly straightforward. The function is given the steady-state error at the end effector as well as the desired threshold for how much error is allowable. If the error is less than the threshold then P_{nom} (the stiffness) is decreased which will slow a leak. If the error is greater than

the threshold then P_{nom} is increased to decrease error at the cost of a faster leak rate as the structural pressure and leak rate will increase with increasing actuation pressure (P_{nom}). P_{nom} is allowed to change at a rate such that it can change from one extreme to the other in 5 seconds. This rate was chosen because faster changes resulted in oscillations as P_{nom} changed faster than the system responded and P_{nom} would overshoot the ideal value. Each time a new position is commanded the algorithm checks if it has been to that position before (if θ is in its dictionary) and starts at the P_{nom} previously found for that position. If the previous position is not in the dictionary then it is added. An example of how P_{nom} changes over time during a run on King Louie when this algorithm is used and a structural leak is present is shown in Figure 4.5. In this example two sets of joint angle commands are used with values of $\theta = [0,0,0,0,0]$ degrees and $\theta = [-30,45,0,-30,0]$. The stiffness varies for each of the two positions and in general is lower early in the run as the arm has not been leaking for long. As the leak continues and the structural chamber loses air (and becomes less rigid) the required stiffness trends upward until the maximum stiffness of 58.6 kPa (8.5 psig) is maintained continuously.

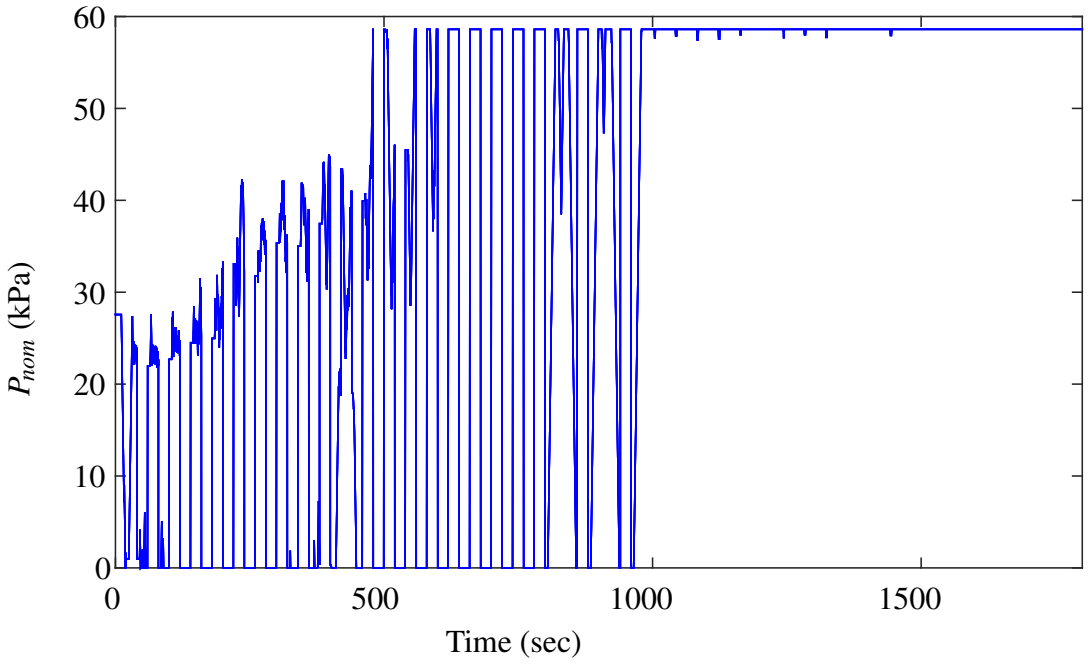

Figure 4.5: This plot shows how P_{nom} varies over the course of a leak as the robot actuates between two different positions.

4.3 Results

Initially tests were run on the Grub but because the Grub is essentially an inverted pendulum where the actuator is a significant portion of the structure all cases performed similarly and never reached a point of failure even when the joint was loaded. The desire to test to failure motivated trials on a multi-DoF system to test the expected benefits of our algorithm.

4.3.1 New Control Flow

The standard control used for controlling our inflatable robots can be seen in Figure 4.6. A more in depth discussion of the control system used and the MPC can be found in [24] and [11]. The methods described in this thesis modify the basic control structure by adding the top two boxes (in green) in the control diagram. Output from the pressure controller and the measured state of the robot are used to detect the presence of a leak in the structure (or actuator as seen in Chapter 5) and then the algorithm shown as Algorithm 1 is used to adjust the desired stiffness (P_{nom}) that is sent to the MPC. Note that the algorithm has only been tested for structural leaks and not actuator leaks. By adjusting P_{nom} during operation we are able to slow a leak.

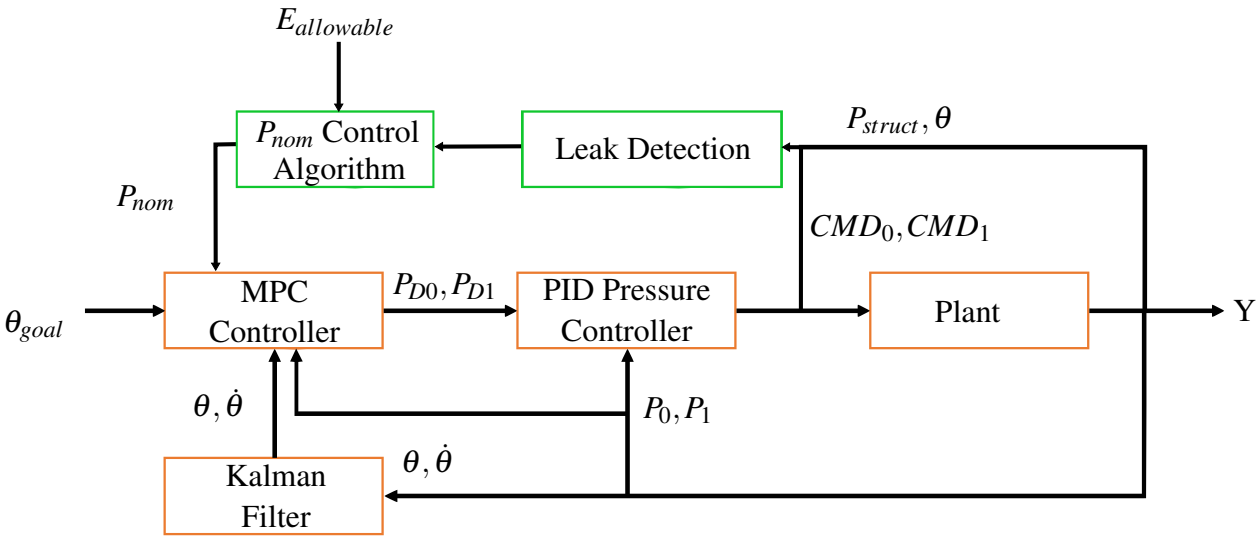

Figure 4.6: Control structure for position control and leak mitigation of a single joint of one of our inflatable robots. P are the measured pressures, P_D are the desired pressures, CMD are the commanded currents sent to the valves, P_{nom} is the nominal pressure sent to the MPC, θ and $\dot{\theta}$ are measured angle and angular velocity respectively, θ_{goal} is the desired joint angle, and $E_{allowable}$ is the amount of error that is acceptable for the task being performed.

4.3.2 Tests on Humanoid Robot Arm

For the tests performed on the 5-DoF arm (King Louie, not including his gripper) the arm was commanded to move between two different sets of joint angles at 20 second intervals. The arm was commanded to move between $[0,0,0,0,0]$ degrees and $[-30,45,0,-30,0]$ degrees which correspond to global end effector positions of $[0.66,0.15,0.85]$ meters and $[0.29,-0.285,0.36]$ meters respectively. The desired end effector position (the "true" position) was measured by motion capture and was based on the average of multiple reaches for the non-leaking case when the joints had reached the commanded angles. Steady-state error at the end effector is taken as the sum-squared difference between the actual position during actuation and the desired position measured using motion capture (the values given above). The arm in the commanded positions is shown in Figure 4.7. Note the difference between Figure 4.7 and Figure 4.4 which demonstrates the error introduced over time due to the leak.

The resulting data are plotted in Figure 4.8 and shows data for three different types of stiffness settings. Each data set is the average of five distinct runs and the error bars show one standard deviation above and below the mean. The first case shows the results when P_{nom} is maintained at 58.6 kPa (8.5 psig) during the leak. The second case shows the simplest form of adaptation we considered which involves simply lowering P_{nom} (in this case to 13.79 kPa (2.0 psig)) and leaving it low during the leak in an attempt to slow the leak rate. The final case shows the run while using our algorithm (see Algorithm 1) to adjust P_{nom} during operation. In the case where our algorithm was used P_{nom} was changed as described in Algorithm 1 to attempt to keep error at 0.04 meters for each of the positions resulting from the commanded angles. Steady-state error in meters at several times during the run is also included in Table 4.1 and shows that the results when using Algorithm 1 result in improved performance compared to the nominal case for nearly the entire run. In terms of sum-squared error the results from using the algorithm outperform the nominal case by over 40% at times and by over 20% for the majority of the run. The results for the algorithm do show decreased accuracy initially in Table 4.1 because the controller is lowering stiffness to increase error to 4 cm. By doing this the leak is slowed and some accuracy is traded for longevity.

From Figure 4.8 and Table 4.1 we can see that the method described in Algorithm 1 to control stiffness is superior to simply lowering P_{nom} (which results in the worst performance over the entire time out of all three cases) and also quickly outperforms the base case where P_{nom} is

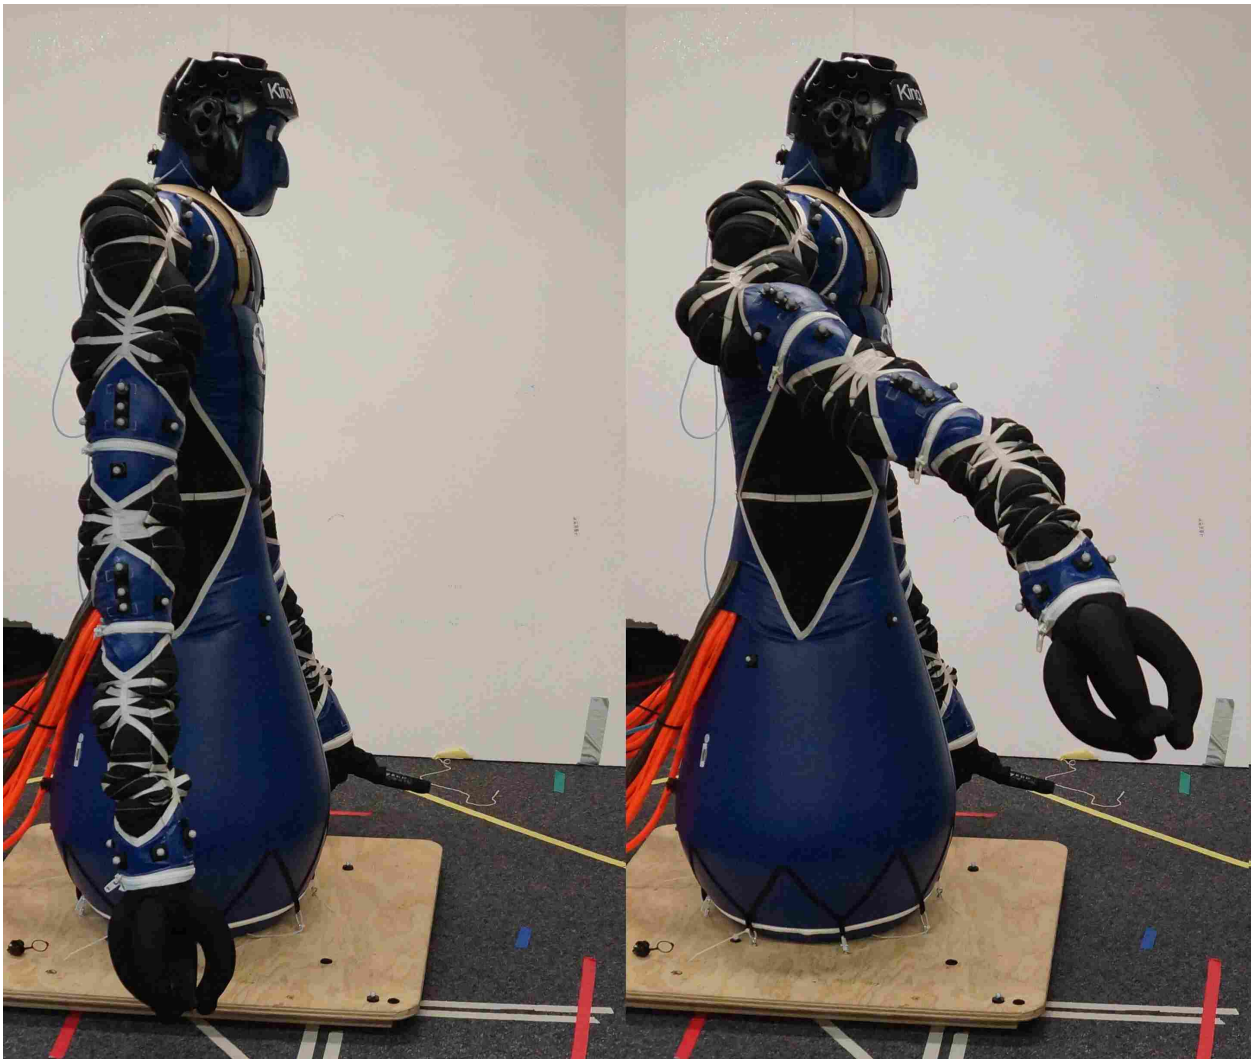

Figure 4.7: Actuation positions for King Louie. The left image shows commanded joint angles of $[0,0,0,0,0]$ degrees. The right image shows commanded joint angles of $[-30,45,0,-30,0]$ degrees.

left at 58.6 kPa (8.5 psig). The improvement in the case where our algorithm is used is largely because for much of the run high accuracy can be obtained at the resting position ($[0,0,0,0,0]$ degrees) with very low pressures and thus can drastically slow the leak. This shows that the amount of improvement that that we can achieve with our algorithm is task and configuration dependant. The trials with our algorithm running have as much as 10 cm less error at times and on average outperform the regular trials after 440 seconds for this task.

Based on the data we have collected we can conclude that our algorithm is effective at prolonging the operational life of an inflatable robot. When a leak in the body chamber that gives

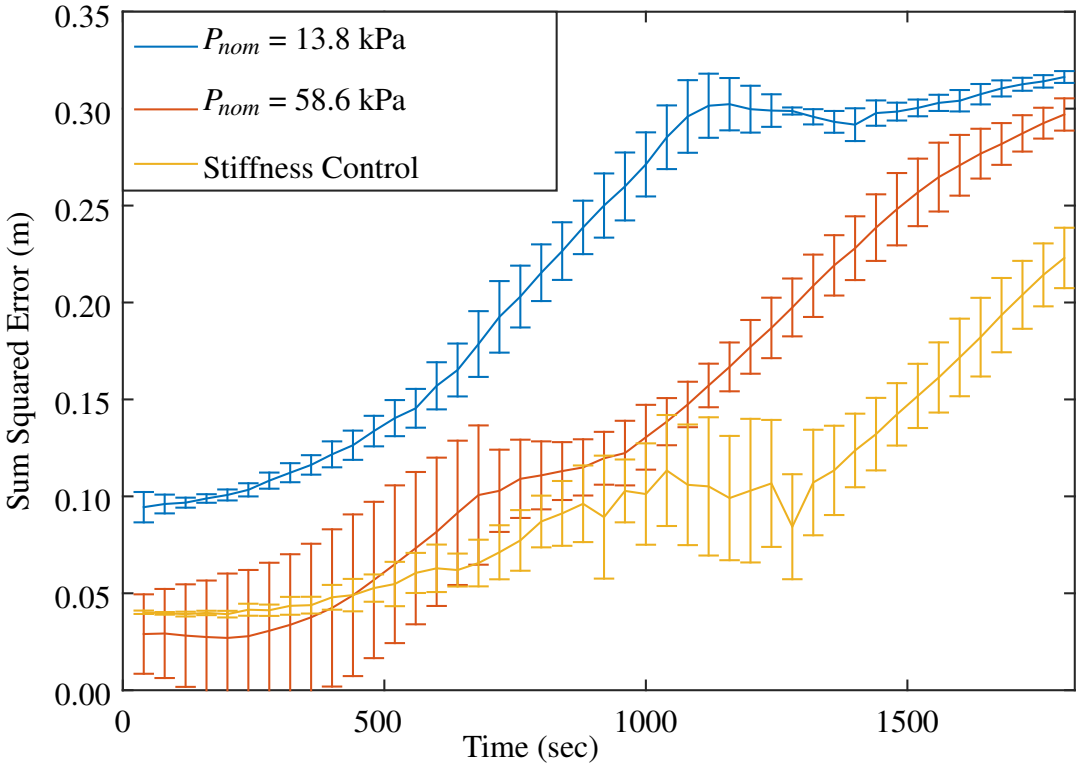

Figure 4.8: Distance from the end effector at steady-state to the actual desired position. Distance is the square root of the sum of the squares of the x, y, and z error. Each line is the average of five runs. Error bars show one standard deviation in either direction.

structure to the robot is present our algorithm can decrease long term error by as much as 50%. We also expect these results to be extensible to other types of leaking in inflatable robots. For example a lower target pressure will result in a lower average actuation pressure which would slow the rate of a leak in an actuation chamber as well. This would be particularly important in circumstances where loss of excess compressed air to the environment would be very undesirable such as space exploration or other applications using a closed loop system to recover vented gas.

Additionally we expect the idea of trading accuracy for longevity to be applicable to other types of resource limited systems. Robots running on battery power could decrease performance and accuracy by actuating less forcefully and accurately to increase battery life. A similar principle is already in use in cell phone location data. To preserve battery life less accurate location information from wifi or cellular triangulation can be used instead of operating power hungry GPS chips. In

Table 4.1: Accuracy data for manipulation of the 5-DoF arm of the inflatable robot during a structural leak.

Sum-Squared Error (m) at	40 sec.	400 sec.	800 sec.	1200 sec.	1600 sec.	1800 sec.
Low: P_{nom}= 13.8 kPa	0.0942	0.1217	0.2154	0.2997	0.3041	0.316
Nominal: P_{nom}= 58.6 kPa	0.0289	0.04245	0.1108	0.1771	0.2708	0.297
Algorithm Controlled P_{nom}	0.0402	0.04793	0.08703	0.1029	0.1717	0.223
% Improvement	-38.75	-12.91	21.45	41.90	36.60	24.92

this way accuracy is sacrificed to preserve battery life. Similar principles will certainly be put into practice as untethered robots become more prevalent. This work demonstrates the viability of such methods for extending the operational life of robots operating with limited resources.

CHAPTER 5. ACTUATOR LEAKS

One of the problems unique to inflatable and fluid actuated systems such as our pneumatic actuation system (Figure 3.2) is dealing with leakage in the actuation chambers. One of the use cases being investigated for this type of robot is space exploration and extra-planetary sample retrieval in which air would be recycled from actuation to be reused rather than vented to the atmosphere as is currently done. In this case compressed air would be a limited resource and leakage in the actuation would result in loss of a limited resource that could not be easily replaced in most space exploration environments due to a thin or non-existent atmosphere around the robot. In addition to the current use cases for these robots it is notable that all untethered robots are resource limited and in general any robot operating in a remote or dangerous environment would likely need to preserve limited resources such as electrical power or fuel. The results in this chapter demonstrate a method for detecting actuator leaks and show a viable and useful trade-off between resource use and system performance that with some modification would be applicable to other robotic platforms to preserve limited resources other than just compressed air. The methods and results in this chapter are for leaks in the actuation chambers of our inflatable robots.

5.1 Actuator Leak Detection

Detecting a leak in an actuation chamber is relatively straightforward due to the nature of the system and our pressure controller. Because the system attempts to maintain a commanded pressure in each actuation chamber, when that pressure has been reached the electrical current commanded to the valves (to open or close) is such that their is no airflow in or out of the chamber. When a leak occurs, however, the valve must remain slightly open to compensate for the air lost through the leak and maintain the commanded pressure in the chamber. By monitoring the average current command when the robot has reached steady state (when the joint angles are no longer changing) we can easily detect when a leak has occurred.

To have robust leak detection that does not need to be hand-tuned for each chamber we utilized the SVM classification library libSVM [35]. Leaks were induced shortly down line from the valve along the air delivery line to obtain leaking data for training the SVM. To be able to detect a wide variety of leak sizes, leaks were induced ranging from 0.5 mm (0.02 in) to 3.175 mm (0.125 in). Data was then given to the SVM along with the status of that data (leaking or non-leaking) and the SVM determined the optimal dividing line between the leaking and non-leaking data. Once the training was finished new data could be given to the SVM for classification and it would return a value indicating whether or not a leak was present. A sampling of the training data along with the dividing line found by the SVM is shown in Figure 5.1. Because the current data is inherently noisy the average of the most recent 300 current commands was used for each data point for testing and training the SVM. The data was also zeroed so that when no leak is present the average current command for each joint at steady-state is 0. This was necessary since each actuator valve has a unique offset and zeroing the data removes the offset. Additionally only data taken while the system was at steady state (pressures and angles were not changing) was used for training and leak detection.

For leak sizes larger than about 1 mm (0.039 in) leak detection worked well. For the smaller leaks we tested with a hole size of 0.5 mm (0.02 in) leak detection was not consistent enough to be reliable. The ability of the SVM to detect a leak is related to the average current command sent to the valve which is in turn directly related to the rate at which air is lost from the bladder. The 1 mm hole (0.039 in) drains the 1-DoF Grub bladder at a rate of about 5% of the total mass of air in the bladder each second. Smaller leaks resulted in a change in valve command that was too small to reliably detect a leak using the SVM. These results should be generalizable and we expect that for a bladder of any size that the SVM should successfully detect leaks in any bladder with similar valves as long as the rate of air loss each second is greater than about 5% of the mass of air in the bladder. It is also worth noting that the system began to have trouble actuating for hole sizes that would result in a loss each second of about 50% of the mass of air in the bladder as the system could no longer fill the bladder in the manner needed for actuation due to the high rate of leakage.

After training the SVM we were able to obtain over 99% percent accuracy at detecting whether a leak larger than 0.5 mm (0.02 in) was present. In total about 20 minutes of data was taken with a leak present and 15 minutes without a leak. Half the data was used for training and

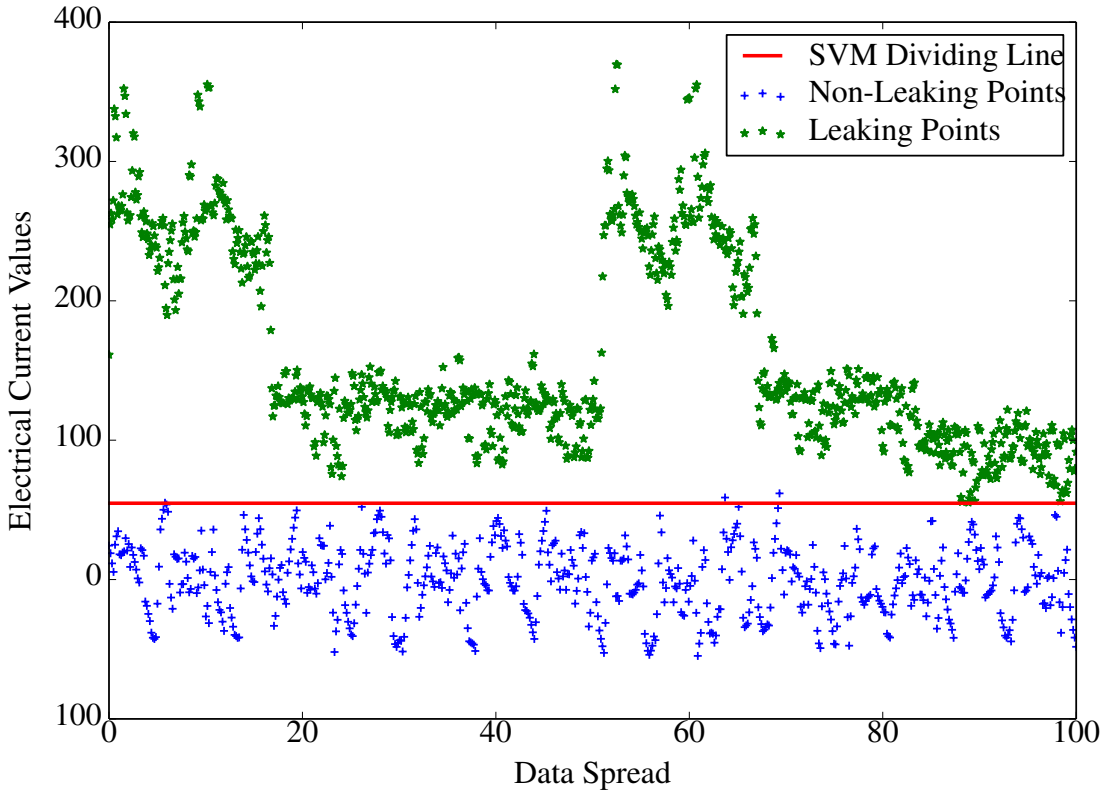

Figure 5.1: A sampling of data used by the SVM to determine the optimal dividing line between leaking and non-leaking current data in the valves. Data for leaks of three different sizes is shown.

the remaining half was used for cross-validation. The data sets were then switched to verify that the training and detection are robust to different data sets. The leak detection performed equally well regardless of which portions of the data were used for training. We also tested the SVM on chambers not used in the training and were able to achieve over 99% accuracy on three additional chambers without using any training data from those chambers which shows that the SVM is general to this type of actuator and not specific to a single chamber. An example of the data given to the SVM and the output is shown in Figure 5.2. As models of the system improve these methods could likely be further improved by utilizing dynamic response data for large step inputs to determine changes in dynamic response that could occur due to a leak which may also enable identification of smaller leaks.

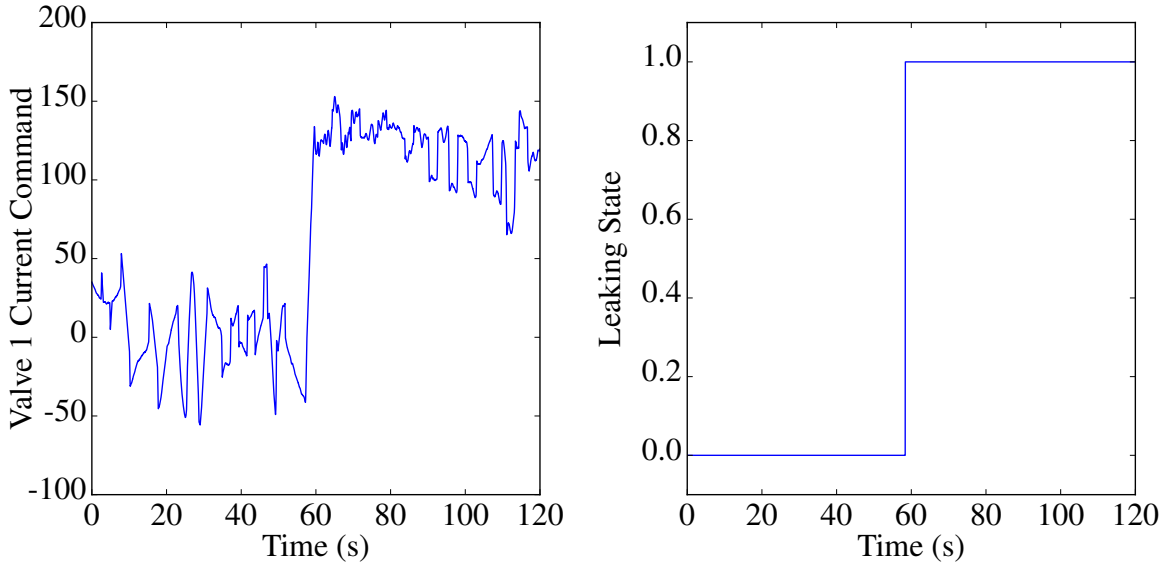

Figure 5.2: The plot on the left shows the current commanded to the valve which shows a fairly obvious jump at 60 seconds when a leak was introduced. The plot at right shows the output of the SVM given the data on the left as input. The leak is flagged immediately at 60 seconds when it was introduced.

5.2 Slowing Actuation Leaks

The following results demonstrate that a leak can be slowed by lowering the target pressure (P_{nom}) in the joint that is leaking but that lowering P_{nom} will also have an effect on the performance and accuracy of the system during actuation. We expect the results here to be generalizable to any system with pneumatic actuation since the leak rate and force output for any pneumatic actuator are both directly correlated with the pressure in the actuator. As a result although the efforts shown here are specific to our platform they are generalizable to any robot using pneumatic actuation as well as many non-robotic pneumatic actuation systems where leaks may occur and in which it would be desirable to preserve compressed air. In particular any system with antagonistic actuation could apply these results.

5.2.1 Grub Performance vs. Leak-Rate Tradeoff

Since we are using stiffness variation to determine how leaks that occur in the actuation chambers may be slowed we took extensive data to determine how system performance changes when stiffness (by virtue of P_{nom}) is adjusted. Figure 5.3 shows how performance changes for the

1-DoF Grub for the range of stiffnesses we use in this thesis. Performance for the Grub is measured using a series of step commands and takes into account rise time (T_r), settling time (T_s), overshoot (PO), steady state error (E_{ss}), and the number of oscillations about the commanded value (OC). The form of the cost function is given in Equation 5.1 where α is a vector of weighting factors that were tuned manually to $\alpha = [1,5,1,4,1]$ in order to emphasize settling time and steady-state error since we were primarily focused on measuring accuracy in this work.

$$cost = \alpha_0 T_r + \alpha_1 T_s + \alpha_2 PO + \alpha_3 E_{ss} + \alpha_4 OC \tag{5.1}$$

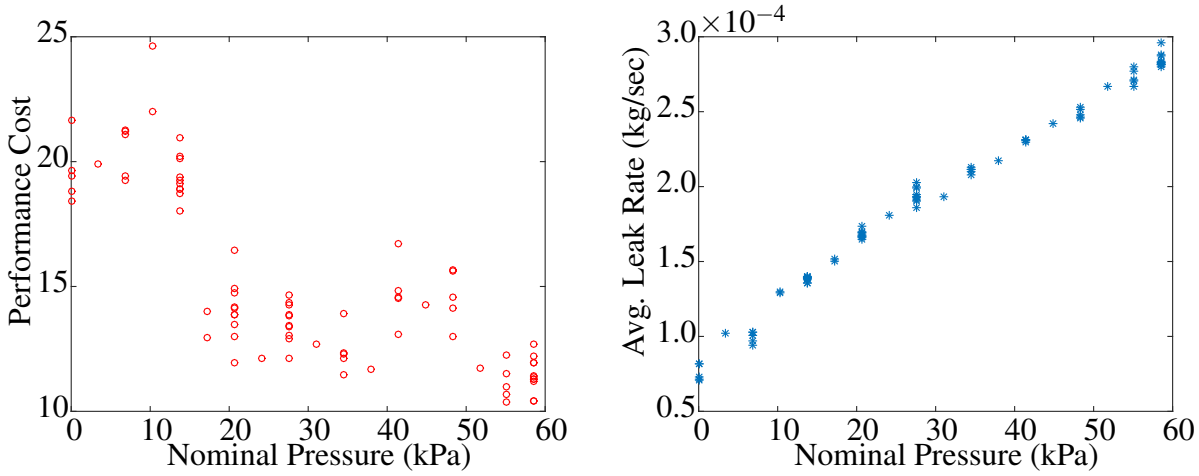

Figure 5.3: Figure at left shows Grub performance at various nominal pressures. The Figure at the right shows the average leak rate for each run at various nominal pressures. The equivalent data for two joints of King Louie is included in Section 4.3 in Figures 5.5 and 5.6.

For all of the trials in this thesis other than for developing the leak detection the leak size was 1.59 mm (0.0625 in). The same leak size was used for all our trials to provide consistency throughout the data and make it possible to easily compare leak and performance data from different joints and trials. For the initial work with the 1-DoF robot a series of step commands were performed repeatedly until the end of the trial.

Baseline performance data was determined by running trials with no leak. We then took data while the system ran with a 1.59 mm (0.0625 in) hole in one actuator. Data was taken with and without a leak for various target pressures (P_{nom}) up to 10 times for each pressure. The Grub was actuated with a 37.8 liter (10 gal.) tank as a pressure source while and the time for the pressure

source to drop from 482.6 - 275.8 kPa (70-40 psi)[1] was measured. For applications where loss of compressed air would be a significant problem (such as space or planetary exploration) an air recycling system would be used to recover air used for actuation and the only air lost would be through the leak. Because we are considering actuation air to be recovered it was necessary to isolate the air lost through the induced leak as this is the only air that would be lost to the environment in a recycled air system. To determine the leak rate each trial required a run with the leak and a run without the leak. The average mass flow rate for the leaking and non-leaking case could each be determined using equation 5.2 which is basically an approximation of the time derivative of the ideal gas law with a large time step where ΔP is the change in pressure during the time interval taken, V is the volume of the air reservoir, R is the gas constant for air, T is the temperature of the air, and $\Delta Time$ is the time elapsed to achieve the desired pressure change. For our tests both volume and temperature were assumed to be constant. The difference between the leaking and non-leaking mass flow was taken as the average leak rate for that test. Each test as described here results in a single point plotted in Figure 5.3 which shows performance as well as average leak rate as a function of P_{nom}.

$$\dot{m} = \frac{\Delta P \cdot V}{R \cdot T \cdot \Delta Time} \tag{5.2}$$

The results in Figure 5.3 show a clear relationship between P_{nom} and average leak rate for the Grub. In the plot on the right of Figure 5.3 it can be clearly seen that as P_{nom} decreases so does the average leak rate. In the lot on the left it can be seen that as P_{nom} decreases the average performance cost for the Grub to actuate increases. Figure 5.3 shows the trade-off between average leak rate and performance for the Grub where by changing P_{nom} the average leak rate can be slowed at the cost of worse system performance. In the next section we discuss the trade-off for King Louie.

5.2.2 King Louie Performance vs. Leak-Rate Trade-off

To adequately demonstrate the trade-off between leak rate and accuracy we took performance and leak rate data from the joint of a 1-DoF system as well as two distinct joints of a 5-DoF system.

[1] Trials actually ran with the pressure source dropping from 551.6 - 206.8 kPa (80-30 psi) but the beginning and end of the data were removed so as to avoid variation that can occur due to refilling of the tank (at the beginning of the data) and variations due to pressure drops during actuation which can cause the trial to end as much as 27.6 kPa (4 psi) prematurely.

This data clearly shows the relationship between desired actuation pressure (P_{nom}), leak rate, and system performance and accuracy at the end effector. The trends shown and discussed here could be used with an algorithm such as the one presented in Chapter 4 or other algorithms that could adjust P_{nom} to slow the actuator leak rate. The testing of an algorithm that would use the results presented here is left as future work.

The methods described in section 5.2.1 were applied on our 5-DoF inflatable robot (Figure 5.4) to determine the affects of stiffness changes on a leak in the actuation bladders. Because King Louie has five joints we gathered data for a leak in the shoulder and a leak in the wrist (joints labeled RSPM and RWPA in Figure 5.4). Only the P_{nom} of the leaking joint was varied while the P_{nom} for the rest of the joints was left unchanged. The shoulder generally requires higher pressures than other joints because it must act on the full weight of the arm and will require higher pressures to reach commanded angles which in turn results in a faster leak on average. Because of this, a leak in the actuator for the shoulder joint is the worst possible location for a leak to occur. The wrist can achieve many desired angles with relatively low pressures (depending on the load) and would result in a slower leak. The wrist is the least catastrophic leak because it allows for a wide range of stiffnesses with minimal change in performance. The shoulder joint tested was the proximal medial joint of the right arm and is designated as RSPM. The wrist joint tested was the proximal anterior joint of the right arm and is designated RWPA (joints are shown labeled on King Louie in Figure 5.4).

The resulting average leak rate (calculated as described in Section 5.2.1) for various commanded nominal pressures is shown in Figure 5.5 and the performance for each run is shown in Figure 5.6. Performance is shown as sum-squared error at the end effector for the commanded joint angles where the desired position is taken from the baseline case where $P_{nom} = 58.6$ kPa (8.5 psi) when the angles have reached their commanded values and there is no leak present. Joint angles and end effector position were measured using a motion capture system where link orientations were obtained from motion capture used to calculate the joint angles[2] and end effector position is measured in the global coordinate system of the motion capture area. For the trials with the 5-DoF robot the pressure reservoir was a 113.6 liter (30 gallon) tank which was necessary because the

[2]Calculating joint angles from joint orientation is necessary because our robots do not have a means of directly measuring joint angle.

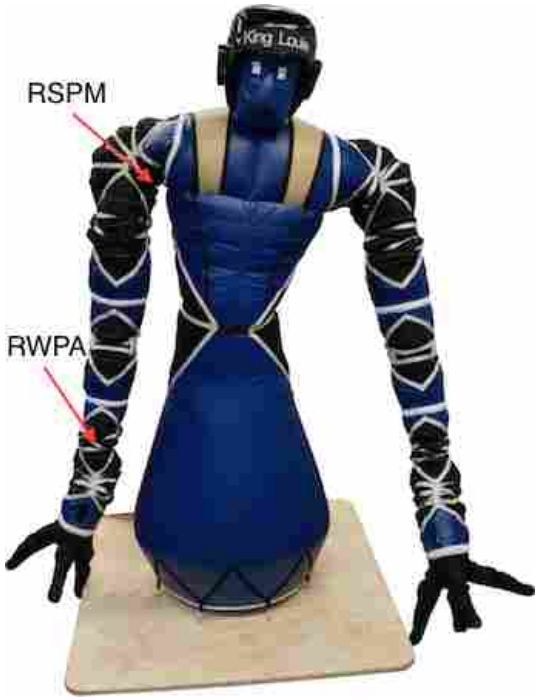

Figure 5.4: Joints of King Louie used for actuation leak testing. Both a shoulder and a wrist were used because the shoulder and wrist behave differently.

5-DoF robot uses significantly more air to actuate than the 1-DoF and the larger tank allows the effects from the leak to be better seen because the trials are longer.

For both joints the data in Figure 5.5 clearly shows that by lowering the commanded nominal pressure (and thereby decreasing stiffness) we can effectively slow the rate at which air is lost to the environment. The spread in the data for each pressure in Figures 5.5 and 5.6 is due to hysteresis in the system as the system uses slightly different pressures and has a somewhat different response for every movement due to the compliance of the system. Table 5.1 lists the times for a 37.8 liter (10 gal.) tank with an initial pressure of 689.48 kPa (100 psi) to lose enough air due to an actuator leak to drop to a pressure of 206.84 kPa (30 psi) which is the typical source pressure for actuation of the inflatable robots. Below 207 kPa actuation would be severely impeded because the robot could not use pressures high enough to actuate properly. The purpose of this work was to develop and implement a method for finding the trade-off between accuracy and run time (due to leak rate) for our robots. Now that this trade-off is known future work could implement an algorithm similar to Algorithm 1 and adjust $E_{acceptable}$ (see Figure 4.6) to operate at the lowest acceptable accuracy to extend life as much as possible.

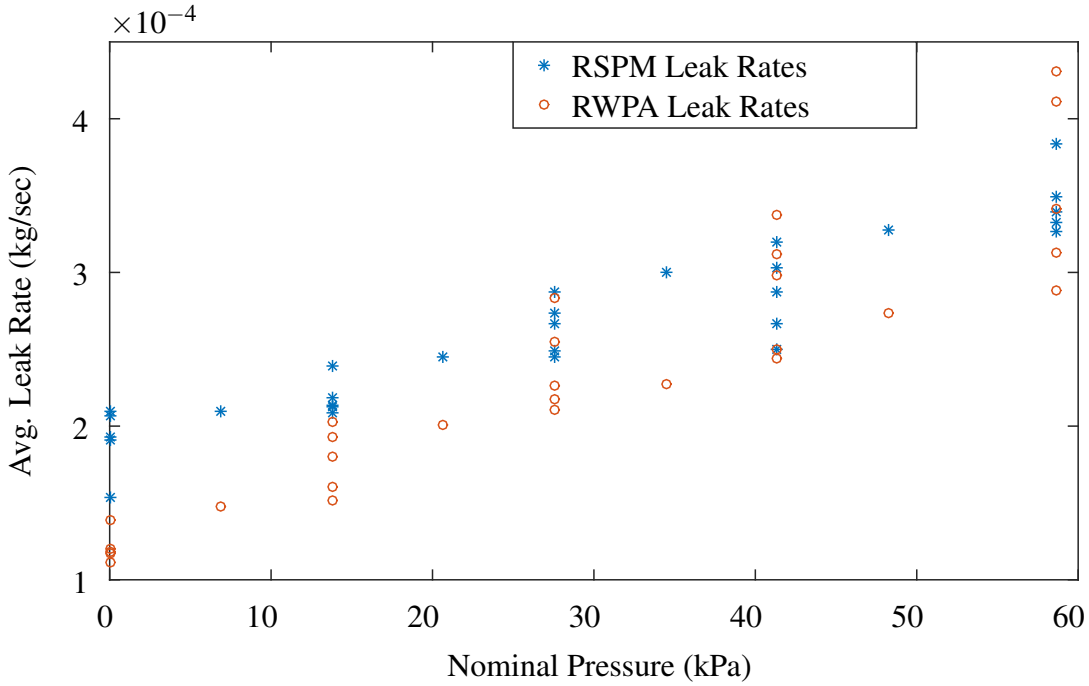

Figure 5.5: Leak rates from induced actuator leak in $\frac{kg}{sec}$ for two joints of King Louie at various commanded nominal pressures. Only the nominal pressure for the leaking joint is varied while the rest remain at $P_{nom} = 58.6$ kPa.

Table 5.1: Accuracy and run data for leaks in two actuators of the 5-DoF arm and the Grub. The times shown are the times for a 37.8 liter (10 gal.) tank to drop in pressure from 689.48 kPa (100 psi) to 206.8 kPa (30 psi) for the average leak rate measured at the commanded P_{nom}.

P_{nom} (kPa)	0.0	13.8	27.6	41.4	58.6
Grub Time (sec)	2887.6	1577.8	1124.3	946.2	768.4
RSPM Time (sec)	1145.6	998.8	827.1	765.1	630.9
RSPM ESS (m)	0.1009	0.0711	0.0272	0.0073	0.0082
RWPA Time (sec)	1800.2	1229.9	915.9	757.7	612.1
RWPA ESS (m)	0.0214	0.0192	0.0218	0.0229	0.0137

As can be seen in Figure 5.6, depending on which joint is leaking and how the nominal pressure is adjusted, performance will be affected. In the case of the wrist joint (RWPA) performance is affected very little and is reasonably consistent throughout the entire regime of pressures shown due to the fact that the wrist is able to achieve its commanded location with a wide range of pressures. The shoulder joint (RSPM) in contrast to the wrist, is unable to fully reach the commanded angle at the lower nominal pressures and the end effector thus has a greater amount of error. A trade-off

must be considered between longevity of the system (as given in Table 5.1) and system performance (which can vary depending on the assigned task). For many non-prehensile tasks, a higher error in robot movements may be acceptable while for a precise manipulation task accuracy may be more important than longevity.

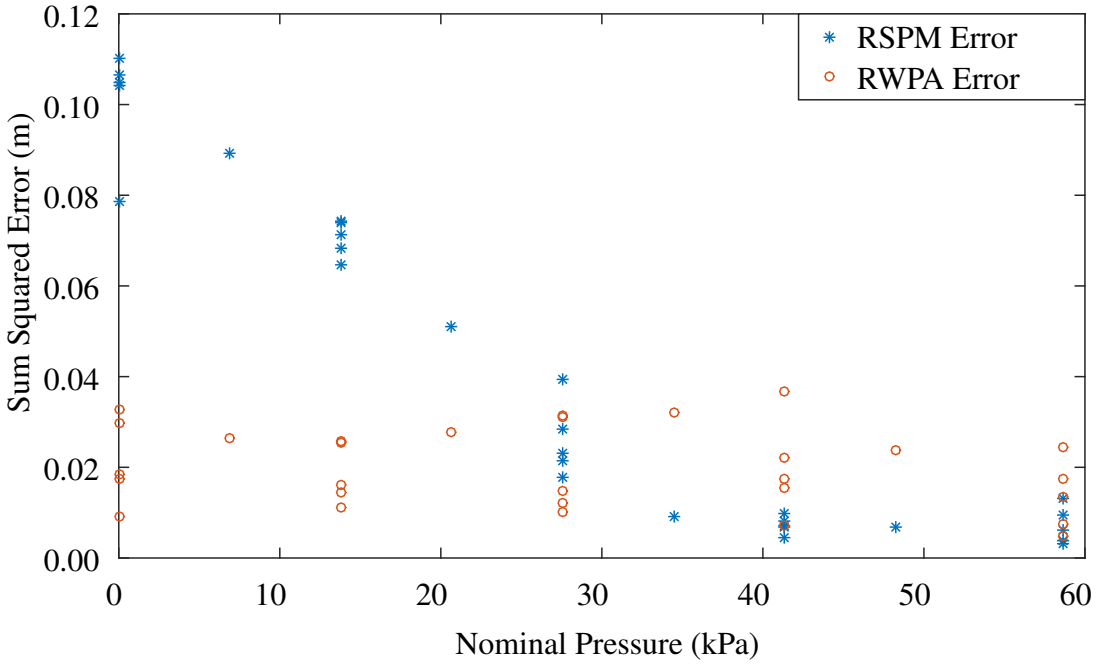

Figure 5.6: Steady state error at the end effector in meters for two joints of King Louie at various commanded nominal pressures. Only the nominal pressure for the leaking joint is varied while the rest remain at $P_{nom} = 58.6$ kPa.

The results shown here for actuation leaks demonstrate that a viable trade-off exists between performance and longevity in our robots. As an example we may imagine the task of sweeping dust off solar cells on Mars. Such a task would require maneuverability but not high accuracy since a simple brushing motion is all that is required. If a leak were present in a shoulder joint the task could likely be performed with a commanded P_{nom} of 0.0 kPa which would result in very low actuation pressures being used (note that 0.0 kPa is the goal but higher pressures would be necessary to lower error). The joint would still actuate and use pressure (and leak) but would attempt to minimize pressure as much as possible. This would result in sum-squared error of 10 cm on average which is likely acceptable for a brushing motion and would result in 80% longer operating time.

Alternatively delivery of a soil sample to a receptacle with a 4.0 cm radius circular opening could be achieved with a P_{nom} of 27.6 kPa which would result in less than 3.0 cm on average which would allow successful delivery of a soil sample while allowing the robot to operate for 30% longer while performing this task. If the leak were in the wrist instead of the shoulder then a commanded P_{nom} of 0.0 kPa could be used in the wrist and both tasks described previously could be performed with less than 2.2 cm of sum-squared error on average and the robot would be able to operate for 194% longer while performing either the solar cell cleaning task or the soil delivery task described.

CHAPTER 6. CONCLUSION

In conclusion we have demonstrated the usefulness and viability of a method and algorithm to adjust control inputs to slow leaks in the structural chamber of an inflatable robot. We have demonstrated that we can effectively detect leaks and have shown that by controlling the stiffness of the joints in an inflatable robot we can slow the rate at which a structural chamber leaks while maintaining an acceptable level of usefulness. We have shown that the operational life of an inflatable robot can be extended and long term error decreased by using our algorithm to adjust target pressures once a leak has occurred. We expect that with work currently under way to improve models and control the algorithm we have developed will perform even better in future iterations.

One of the main limitations to the work presented here is that both leak detection (for structural and actuator leaks) and the algorithm shown in Chapter 4 to slow structural leaks depend on steady-state data from the robot. If the robot is always moving and does not reach a steady-state value for any appreciable amount of time then the methods outlined in this thesis cannot be used. Future work could improve these methods so that leak detection and mitigation can be performed during dynamic motion as well as at steady-state. We also expect that by collecting more data with these kinds of platforms, future work can be done to develop a better estimator for when our system is leaking air to improve detection time and reliability. For inflatable robots to be a viable option in many real world applications such as space exploration and search and rescue, this kind of work will be necessary to prolong operation and avoid mission failure.

We have also demonstrated the usefulness of varying the target pressure to slow the leak rate from a leak in the actuation chamber of a pneumatically actuated robot. We have demonstrated the ability to accurately detect leaks in pneumatic actuators. We have shown that there exists a trade-off between the target pressure (or stiffness) of a joint and the steady-state error at the end effector for a commanded angle. The trade-off varies for each individual joint depending on loading and location along the arm but the trends are the same for all joints. Using this data robot operators could

effectively determine how to set the target pressure during operation of a pneumatic robot depending on the accuracy and longevity requirements of the mission. These trends and this work may be applied to other robotic systems with pneumatic actuation with minimal adjustment. Additionally the core idea of utilizing a trade-off between resource consumption and system performance is an idea that could be applied to more traditional robots with electric motor actuation in resource limited environments. The work done here has relied on experimental data to determine the trade-off between performance and leak rate. Future work will hopefully allow a closed form equation for the trade-off as we continue to develop better models of these robots and how they operate. Future work can also include implementing the algorithm shown in this thesis for structural leaks to adjust P_{nom} during operation to slow an actuator leak while maintaining an error threshold.

REFERENCES

[1] Sanan, S., Ornstein, M. H., and Atkeson, C. G., 2011. "Physical human interaction for an inflatable manipulator." In *Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE*, IEEE, pp. 7401–7404. 6, 7

[2] Gaiser, I., Wiegand, R., Ivlev, O., Andres, A., Breitweiser, H., Schulz, S., and Bretthauer, G., 2012. *Smart Actuation and Sensing Systems - Recent Advances and Future Challenges*. October, ch. Compliant Robotics and Automation with Flexible Fluidic Actuators and Inflatable Structures, pp. 567 – 608. 6, 13

[3] Hawkes, E. W., Christensen, D. L., and Okamura, A. M., 2016. "Design and implementation of a 300% strain soft artificial muscle." In *2016 International Conference on Robotics and Automation (ICRA)*, IEEE, pp. 4022–4029. 6

[4] Caldwell, D. G., Medrano-Cerda, G. A., and Goodwin, M., 1995. "Control of pneumatic muscle actuators." *IEEE Control Systems,* ***15***(1), pp. 40–48. 6, 7

[5] Post, B. K., 2013. "Robust state estimation for the control of flexible robotic manipulators." PhD thesis, Georgia Institute of Technology, December. 6

[6] Tahia, J., Wilkening, A. W., and Ivlev, O., 2013. "Adaptive position control of fluidic soft-robots working with unknown loads." In *2013 IEEE/ASME International on Advanced Intelligent Mechatronics*, IEEE/ASME, pp. 1121–1126. 6, 10

[7] Best, C. M., Wilson, J. P., and Killpack, M. D., 2015. "Control of a pneumatically actuated, fully inflatable, fabric-based humanoid robot." In *Humanoids, 2015 IEEE-RAS International Conference on*, IEEE. 6, 7, 10

[8] Rus, D., and Tolley, M. T., 2015. "Design, fabrication and control of soft robots." *Nature,* ***521***(7553), pp. 467–475. 6, 7

[9] Rupert, L., Hyatt, P., and Killpack, M. D., 2015. "Comparing model predictive control and input shaping for improved response of low-impedance robots." In *Humanoids, 2015 IEEE-RAS International Conference on*, IEEE. 7, 10

[10] Bicchi, A., Rizzini, S. L., and Tonietti, G., 2001. "Compliant design for intrinsic safety: General issues and preliminary design." In *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, pp. 1864 – 1869. 7

[11] Best, C. M., Gillespie, M. T., Hyatt, P., Killpack, M. D., Rupert, L., and Sherrod, V., 2015. "Model predictive control for pneumatically actuated soft robots." Accepted to IEEE Transactions on Robotics and Automation. 7, 16, 17, 27, 29

[12] Medrano-Cerda, G. A., Bowler, C. J., and Caldwell, D. G., 1995. "Adaptive position control of antagonistic pneumatic muscle actuators." In *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*, Vol. 1, IEEE, pp. 378–383. 7, 8, 10

[13] Klute, G. K., and Hannaford, B., 1998. "Fatigue characteristics of mckibben artificial muscle actuators." Vol. 3, IEEE, pp. 1776–1781. 7, 8, 10

[14] Bao, G., Yang, Q., Li, S., and Zhang, L., 2009. "Static model of pneumatic flexible rotary joint." In *Proceedings of the 2009 IEEE International Conference on Robotics and Biomimetics*, IEEE, pp. 2163 – 2167. 8, 13

[15] Mihajlov, M., Hübner, M., Ivlev, O., and Gräser, A., 2006. "Modeling and control of fluidic robotic joints with natural compliance." In *IEEE International Conference on Control Applications*, IEEE, pp. 2498–2503. 8, 14

[16] Ivlev, O., 2009. "Soft fluidic actuators of rotary type for safe physical human-machine interaction." In *2009 IEEE 11th International Conference on Rehabilitation Robotics*, IEEE. 8, 13

[17] Granosik, G., and Borenstein, J., 2004. "Minimizing air consumption of pneumatic actuators in mobile robots." In *IEEE International Conference on Robotics and Automation*, IEEE, pp. 3634–3639. 8

[18] Ozawa, R., Kobayashi, H., and Ishibashi, R., 2015. "Adaptive impedance control of a variable stiffness actuator." *Advanced Robotics,* ***29***(4), pp. 273–286. 8, 11

[19] Tonietti, G., and Bicchi, A., 2002. "Adaptive simultaneous position and stiffness control for a soft robot arm." In *Proceedings of the 2002 IEEE Intl. Conference on Intelligent Robots and Systems*. 8, 11

[20] Shen, X., and Goldfarb, M., 2007. "Simultaneous force and stiffness control of a pneumatic actuator." *Journal of Dynamic Systems, Measurement, and Control,* ***129***, July, pp. 425–434. 8

[21] Plagemann, C., Stachniss, C., and Burgard, W., 2006. *Efficient Failure Detection for Mobile Robots Using Mixed-Abstraction Particle Filters*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 93–107. 8

[22] Simani, S., 2003. *Model-based fault diagnosis in dynamic systems using identification techniques*. Springer, London; New York. 8

[23] Sheng, S., 2014. *Gearbox typical failure modes, detection and mitigation methods*. National Renewable Energy Laboratory, National Wind Technology Center, Golden, Colorado. 9

[24] Gillespie, M. T., Best, C. M., and Killpack, M. D., 2016. "Simultaneous position and stiffness control for an inflatable soft robot." In *2016 International Conference on Robotics and Automation*, IEEE. 9, 10, 14, 15, 16, 29

[25] Best, C., 2016. "Position and stiffness control of inflatable robotic links using rotary pneumatic actuation." Master's thesis, Brigham Young University, August. 10, 14

[26] Qin, S. J., and Badgwell, T. A., 2003. "A survey of industrial model predictive control technology." *Control Engineering Practice,* ***11***(7), July, pp. 733–764. 10

[27] Wang, Y., and Boyd, S., 2010. "Fast model predictive control using online optimization." *IEEE Transactions on Control Systems Technology,* ***18***(2), March, pp. 267–278. 10

[28] Mattingley, J., and Boyd, S., 2012. "Cvxgen: a code generator for embedded convex optimization." *Optimization and Engineering,* ***13***, pp. 1–27. 10

[29] Jain, A., Killpack, M. D., Edsinger, A., and Kemp, C. C., 2013. "Reaching in clutter with whole-arm tactile sensing." *The International Journal of Robotics Research*. 10

[30] Killpack, M. D., Kapusta, A., and Kemp, C. C., 2015. "Model predictive control for fast reaching in clutter." *Autonomous Robots*, pp. 1–24. 10

[31] Erez, T., Lowrey, K., Tassa, Y., Kumar, V., Kolev, S., and Todorov, E., 2013. "An integrated system for real-time model predictive control of humanoid robots." In *13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, IEEE, pp. 292–299. 10

[32] Ström, K. J. A., and Wittenmark, B., 2008. *Adaptive Control*., 2nd ed. Dover Publications, Inc. 10

[33] Annamalai, A. S. K., Sutton, R., Yang, C., Culverhouse, P., and Sharma, S., 2014. "Robust adaptive control of an uninhabited surface vehicle." *Journal of Intelligent & Robotic Systems*, pp. 1–20. 10

[34] Chowdhary, G., Mühlegg, M., How, J. P., and Holzapfel, F., 2013. "Concurrent learning adaptive model predictive control." In *Advances in Aerospace Guidance, Navigation and Control*. Springer, pp. 29–47. 10

[35] Chang, C.-C., and Lin, C.-J., 2013. *LIBSVM: A Library for Support Vector Machines*. Department of Computer Science National Taiwan University, Taipei, Taiwan, March. 11, 35