Simultaneous Position and Stiffness Control for an Inflatable Soft Robot

Morgan T. Gillespie Charles M. Best Marc D. Killpack

***Abstract*— Soft robot research has led to the development of platforms that should allow for better performance when working in uncertain or dynamic environments. The potential improvement in performance of these platforms ranges from mechanical robustness to high forces, to applying lower incidental contact forces in uncertain situations. However, the promise of these platforms is limited by the difficulty of controlling them. In this paper, we present preliminary results on simultaneously controlling stiffness and position for a pneumatically actuated soft robot. Improving on our prior work, we show that by including the pressure in our soft robot actuation chambers as state variables we can improve our average rise time by up to 137%, settling time by 119%, and overshoot by 853%. In addition to these improvements, we can now control both joint position and stiffness simultaneously. This performance improvement comes from using Model Predictive Control running at 300 Hz with improved dynamic models of the soft robot. High performance control of soft robot joints, such as the joint presented in this paper, will enable a wide range of robot applications that were previously difficult or impossible due to the rigid nature of traditional robot linkages and actuation schemes.**

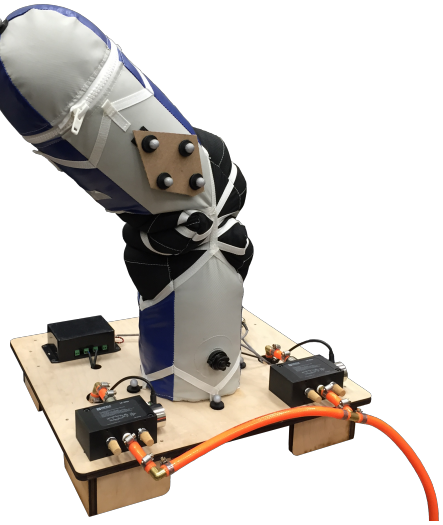

Fig. 1: Single joint test hardware that is called a Grub.

I. Introduction

Control and actuation of traditional robotic platforms are accomplished through electrical motors acting on inelastic joints and rigid links. While these stiff platforms excel at accurate position control, they are fundamentally different than biological organisms such as humans. Compliant human joint actuation is achieved through antagonistic muscles pulling on bones through tendons. The ability to co-contract these antagonistic muscles allows independent control over both joint stiffness and joint position. A soft robot actuator with compliant antagonistic elements similar to the human muscular system should also allow for simultaneous control of joint stiffness and position. This type of actuator, when compared to a rigid robot, will be naturally more resistant to environmental damage upon impact and safer to operate around humans or delicate environments. Additionally, variable stiffness will enable a soft robot to take advantage of stored energy within the system to move explosively, mitigate unexpected contact forces, and vary end effector stiffness passively. The cost for these advantages is a complexity in design and control of the actuator. Much research has been done on the design and control of robots that use significant compliance in meaningful ways [1]. This paper presents a control method and the dynamic models necessary for control of both position and stiffness of a pneumatically actuated inflatable robot joint.

Our specific contributions include the following:

- Development of a linearized state space model that predicts the state of an antagonistic, pneumatically actuated, soft robot.
- First application of Model Predictive Control to soft robots that allows simultaneous control of joint position and stiffness.
- A controller which shows significant improvement in terms of overshoot, rise time and settling time of joint control as compared to past work [2].

The rest of this paper is organized as follows, in Section II we discuss past related research. In Section III we describe the soft robot joint we use for testing. Section IV contains the formulation of our dynamic models and the model predictive controller. Section V contains results and discussion for joint control and Section VI has results on stiffness tracking while maintaining and varying a commanded angle.

II. Related Work

In this section, we describe related research in terms of compliant and variable stiffness robot control, soft robot platforms and control, and past applications with model predictive control.

Significant research, both in terms of the amount and impact, has been completed in the area of adding compliance to common robotic platforms. A system such as the DLR lightweight robot [3], [4] can simulate compliance with an impedance model by utilizing high-rate feedback and low-level torque control [5]. Simulating compliance is limited by control bandwidth and torque capabilities while it can also introduce instabilities in the base controller. In addition, this advanced control and specialized hardware can result in a

more expensive platform. Hardware designs seeking to avoid these limitations have introduced elastic elements between the actuator and load. The most well known method is the use of series elastic actuators (SEA) as described in [6]. SEAs are limited in that the compliance of the actuator is passive and the mechanical passive compliance cannot be changed. Any change in compliance must be done in low-level torque control which is bandwidth and force limited due to the spring. Designs for actuators that have variable compliance can be seen in [7] and [8]. These designs typically use variable spring lengths, non-linear springs, or pneumatic actuators to affect stiffness. Optimal methods to vary compliance in [9] were shown to maximize impact forces from hammering and in [10] unintentional contact forces with a human were minimized with controlled stiffness. In [10] the Head Injury Criterion (HIC) is identified as a limit for serious head injury upon robot impact, where inertia of the robot is a major driving contributor. This indicates with all operating parameters constant, a reduction in robot inertia, such as for our fabric-based soft robot joint, directly reduces the HIC rating.

The variable compliance actuation method most relevant to the platform used in this work is the method proposed by Laurin-Kovitz [11] which utilizes two motors pulling on antagonistic nonlinear springs to actuate the same rigid rotary joint. The systems Laurin-Kovitz designed allow for concurrent control over both joint torque and stiffness, but required two motors and significant additional hardware per joint. A different actuation approach utilizes McKibben Artificial Muscles, which emulate human muscles, was proposed in [12] and [13]. More recently, work has been completed on rotary elastic chamber actuators such as in [14] and [15], where two antagonistic bellows impart torque on an armature rotating about a rigid rotary joint. In all of these cases however, the safety and robustness introduced by compliance in the joints is mitigated by higher inertia motors or rigid links between the joints.

Our paper presents an improved model and controller for an alternative actuation design utilizing inflatable fabric-based joints for simultaneous position and stiffness control. While comparable to the bellow type actuators previously described, this system applies a torque to the joint through expanding chambers pressing upon an soft structure rather than a rigid armature. For this system, the joint and link structure come from inflatable polymer bladders within non-rigid fabrics, resulting in an entirely compliant platform. The model proposed in this work expands on the model used in [2] in order to account for the pressures in each bladder. Doing so allows for control of both position and stiffness. A fabric-based platform of this nature is intrinsically lightweight, compact, and compliant due to the physical properties of the air used for both actuation and structure.

The unique advantages of soft, inflatable robots over rigid robots for specific applications are what motivate this research. Applications include health care, living assistance, space exploration, search and rescue, orthotics, and prosthetics. It has already been shown that contact forces from inflatable links can be controlled in [16] and [17] with cable driven actuators. In [18], [19], and [20] it was found that planning is viable for elastomer actuators using dynamic and constant curvature kinematic models. Designs for rotary, fabric-based, pneumatic actuated joints were proposed in [21] and our actuators used in this work are based on these designs. Work in [22] characterized different models for traditional rigid servo-pneumatic actuators which made different constant temperature assumptions and in [23] these assumptions were used to control force and stiffness for a linear pneumatic actuator. The goal of our work is to use a modified linear model and optimal control methods to control a rotary joint with an inflatable structure instead. The lack of literature on the control of inflatable structures where there is a wide range of applications suggests a viable new area for research.

The proposed control approach utilizes Model Predictive Control (MPC) which is an optimal control method originating from the process industry in chemical production and oil refinement [24] and more recently in robotics [25] and UAV research [26], [27], [28]. Recent advances in computing power and optimization techniques such as those proposed in [29], [30] have made this possible. MPC allows for the minimization of configurable costs across a discrete finite time horizon given specific constraints. The antagonistic compliant system we described, benefits significantly from a model-based control method in order to achieve dynamic stiffness and accurate position control. MPC not only allows for model-based control but also the consideration of real-world constraints such as joint limits, pneumatic valve actuator limitations, and hardware pressure limits. This paper presents the dynamic model and controller to achieve simultaneous stiffness and position control of a soft pneumatic platform and presents experimental data demonstrating performance.

III. Robot Platform Description

The inflatable robot platform used in this research was both designed and constructed by Pneubotics, an affiliate company of Otherlab. A single degree of freedom pneumatic testing platform called a Grub seen in Figure 1 was used for the purpose of this paper. One end of the fabric structure is secured to a wooden platform where the valves are affixed. A single pressurized body bladder extends up the entire length of the Grub, pressing against a non-rigid outer fabric shell to provide structural links. The body bladder is bent at its midsection during actuation, bisecting the central bladder into two separate links at the same pressure. Two flexible inextensible pockets are built into the outer shell, seen as black fabric in Figure 1, where the actuation bladders are inserted. The actuation bladders sit opposite and lateral to the body bladder. Bidirectional bending is achieved by pressurizing the actuation bladders, causing them to expand and lengthen. The bladder pocket fabric contains this expansion and translates extension into joint rotation.

Pressure and angle sensing are read by a pressure sensor and inertial measurement unit (IMU) board affixed to the

outside of the top body link. Joint angles are determined by the IMU angle relative to the gravity vector normal to base platform (see Section IV-B for more detail). The internal body bladder is inflated to 1-2 pounds per square inch gauge (psig) pressure whereas the actuation bladders are filled from 0-20 psig. The source pressure is provided by a compressor regulated to 25 psig.

Air flow is controlled from the pressure source and vented to atmosphere through Enfield LS-v25 five port spool valves. For the purpose of the paper, only one output port of the Enfield valves is used, making these valves act as three port spool valves. As seen in Figure 2, each bladder has an individual valve for control, whilst sharing the same pressure source.

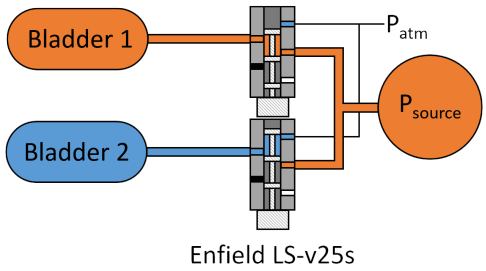

Fig. 2: Representative figure of valve and bladder configuration.

Communication between the sensors and controller is achieved using the Robot Operating System (ROS) operating in non-realtime on an Ubuntu workstation. Sensor values for the IMU and pressure are read and published at approximately 1 kHz. Pressure for each bladder is controlled by an underlying PID controller also operating at approximately 1 kHz. Desired pressures are published over ROS and the valves are actuated to achieve commanded pressure values.

IV. Single Joint Dynamics and Control

A. Dynamics

Previous work has been completed by Best et al [2] on the single joint control of this platform. While effective at angle tracking, the work only utilized change in angle and angular velocity as states and did not permit the modulation of stiffness. In the previous work, angle was related to pressure through a single pressure profile curve. A single pressure P was determined using a desired angle from an experimentally determined fit of the data. This approach effectively reduced the degrees of freedom for control to one instead of the two actual pressures that can be controlled for this platform. This past work provided the desired joint angle control, but was limited in both its operational range and performance capabilities as it restricted operational pressures and did not consider the dynamics of the underlying pressure controllers for each actuation chamber.

In [2], a steady state joint angle response for a range of different pressure inputs was produced. This plot, seen in Figure 3, produced a general trend to relate angle and pressures in both bladders but control accuracy was limited

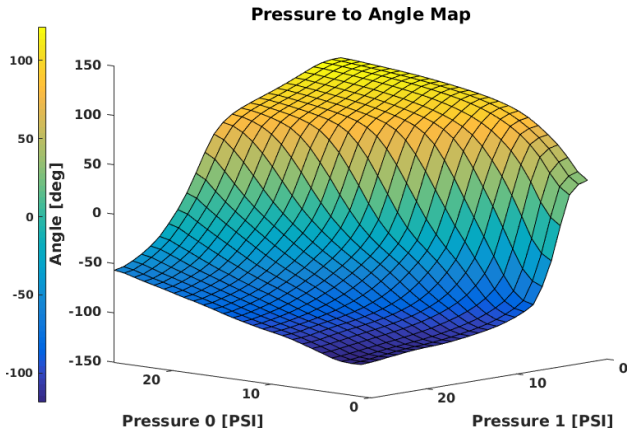

Fig. 3: Plot of steady state mapping used to convert equilibrium angles to chamber pressures.

by hysteresis in the real platform. In practice, different initial conditions of pressure and angle produced a different steady state angle.

Instead of using the mapping directly, we use it to identify an equilibrium angle. This angle is the angle the system would settle to based upon given input pressures, if the chambers were drained and then filled to the input pressures. Unlike the past work in [2], this equilibrium angle was then used in an impedance model to describe torque at the joint as a torsional spring:

$$\tau = K_s(\theta_e - \theta) \tag{1}$$

where K_s is a fixed stiffness constant, θ_e is the equilibrium angle determined by the pressure to angle map, and θ is the current state angle. For our dynamic models, we assume that the platform has rigid links relative to its compliant joints and that a constant K_s is a sufficient model.

The basic dynamic model of the rigid link is the same as used in [2], where the link itself was described as an inverted pendulum:

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

where I is the calculated moment of inertia rotating about the approximate joint center, K_d is a damping constant, m describes the mass of the link, g describes gravity, and L is the approximate distance from the joint center to the center of mass.

K_d, K_s, and m were found in [2] through system identification, where a step input was applied and parameters were optimized to fit collected data. The optimal value for m was near zero and 3 orders of magnitude below similarly optimized values for K_d and K_s. Due to our lightweight inflatable hardware, the influence of mass on the system performance was identified to be negligible and removed from the model for a single joint. The simplified dynamic model becomes:

$$I\ddot{\theta} + K_d\dot{\theta} = \tau \tag{3}$$

Future work with multiple links will need to include this removed term to accurately describe gravitational effects.

Unlike the models used in [2], dynamics of the underlying pressure PID controller are captured by a simple first order model:

$$\dot{P} = -aP + bP_D \tag{4}$$

where P is the pressure in a corresponding bladder, P_D is the desired pressure, and a and b are parameters fit to collected step performance data. For the purpose of this paper fill and drain rates are assumed to be equal. Accurately modelling the true difference in fill and drain rates will require better gas dynamic models.

Instead of using the original nonlinear equilibrium mapping between P_0, P_1 and θ shown in Figure 3, we use a tangent plane estimation at each time step. This planar equation is constructed as a function of the pressures in the two antagonistic bladders, P_0 and P_1:

$$\theta = \alpha_1 + P_0\alpha_2 + P_1\alpha_3 \tag{5}$$

where α_1, α_2, and α_3 are three parametric constants needed to describe a plane, P_0 is the pressure for one bladder, and P_1 is the pressure for the opposing bladder. The steady state data surface was smoothed and refined using cubic spline interpolation.

The variable θ_e is also approximated using this same planar relation, using commanded pressures sent to the underlying pressure PID controller as inputs:

$$\theta_e = \alpha_1 + P_{D,0}\alpha_2 + P_{D,1}\alpha_3 \tag{6}$$

where $P_{D,0}$, represents the desired pressure for bladder 0. Commanded pressures are used to describe θ_e since they are the pressure values that the system will reach at steady state.

By combining equations 1, 3, 5, and 6 we are able to construct a differential equation which describes the body dynamics as a function of current and commanded pressures.

$$I\ddot{\theta} + K_d\dot{\theta} = K_s(\alpha_2(P_{D,0} - P_0) + \alpha_3(P_{D,1} - P_0)) \tag{7}$$

where α_1 describes the vertical offset of the planar surface approximated by the steady state angle map in figure 3. By including α_1 in both θ and θ_e, it is removed from the system dynamics. The removal of α_1 serves to dismiss the influence of measurement error and hysteresis when comparing the steady state angle to current system performance.

The rigid body dynamics along with pressure dynamics in equation 4 can then be placed in linearized state space form:

$$\begin{bmatrix} \ddot{\theta} \\ \dot{\theta} \\ \dot{P}_0 \\ \dot{P}_1 \end{bmatrix} = \boldsymbol{A} \begin{bmatrix} \dot{\theta} \\ \theta \\ P_0 \\ P_1 \end{bmatrix} + \boldsymbol{B} \begin{bmatrix} P_{D,0} \\ P_{D,1} \end{bmatrix} \tag{8}$$

where:

$$\boldsymbol{A} = \begin{bmatrix} \frac{-K_d}{I} & 0 & \frac{-K_s\alpha_2}{I} & \frac{-K_s\alpha_3}{I} \\ 1 & 0 & 0 & 0 \\ 0 & 0 & -a & 0 \\ 0 & 0 & 0 & -a \end{bmatrix} \tag{9}$$

$$\boldsymbol{B} = \begin{bmatrix} \frac{-K_s\alpha_2}{I} & \frac{-K_s\alpha_3}{I} \\ 0 & 0 \\ b & 0 \\ 0 & b \end{bmatrix} \tag{10}$$

In this formulation, the matrix $\boldsymbol{A}$ is singular. As such, the matrix exponential method that was used in [2] cannot be used for discretization. However, the controllability matrix formed from this system is full rank for $\boldsymbol{A}$ and $\boldsymbol{B}$ in the form of Equations 9 and 10 which means the system is controllable. These state space equations are instead transformed from the continuous time-domain to discrete state space equations using the bilinear transform method.

This transformation gives the discrete state space equations:

$$\begin{bmatrix} \dot{\theta}[k+1] \\ \theta[k+1] \\ P_0[k+1] \\ P_1[k+1] \end{bmatrix} = \boldsymbol{A_d} \begin{bmatrix} \dot{\theta}[k] \\ \theta[k] \\ P_0[k] \\ P_1[k] \end{bmatrix} + \boldsymbol{B_d} \begin{bmatrix} P_{D,0}[k] \\ P_{D,1}[k] \end{bmatrix} \tag{11}$$

These discrete-time state space equations can now be used for predicting the future states of the system given pressure inputs. This is an essential part of the model predictive controller described in Section IV-C.

B. Sensing

An IMU sits above the center of mass on the top link section of the outer shell. An angle of 0 degrees is assumed to be vertical, where the central body bladder is straight and unbent. Angle is read as a deflection from vertical in a plane perpendicular to the joint axis. Utilizing 3-axis accelerometers and gyroscopes, we calculate angle by measuring the acceleration vector due to gravity and smoothing it using a Kalman filter designed for an inverted pendulum.

C. Control

The model predictive controller solves an optimization, simulating the predicted states over the horizon T by varying the pressure inputs to produce the trajectory incurring the least cost subject to constraints. The discretized matrices A_d and B_d, the current states $\dot{\theta}[k]$, $\theta[k]$, $P_0[k]$, and, $P_1[k]$, the previous inputs $P_{D,0}[k-1]$ and $P_{D,1}[k-1]$, the final goal angle θ_{goal}, target pressures $P_{T,0}$ and $P_{T,1}$, and the model constraints and weights are fed into an MPC solver at every time step. Pressure target values, $P_{T,0}$ and $P_{T,1}$, are low weight cost value parameters that allow us to set a desired pressure value that correlates to stiffness. The cost function minimized across the horizon T is:

$$\begin{aligned} minimize \sum_{k=0}^{T} \Big(&\|\theta_{goal} - \theta[k]\|_Q^2 + \|\dot{\theta}[k]\|_R^2 \\ &+ \|P_0[k] - P_{T,0}\|_S^2 + \|P_1[k] - P_{T,1}\|_S^2 \Big) \end{aligned} \tag{12}$$

subject to the system model as constraints, as defined in equation 11, as well as the following additional constraints:

$$\begin{aligned} |\theta| &\leq \theta_{max} \\ P_{min} \leq P_D &\leq P_{max} \\ |\Delta P_D| &\leq \Delta P_{max} \end{aligned} \tag{13}$$

where Q, R, S are scalar weights manually tuned for superior empirical performance, the same value of S is used for both pressure target costs, θ_{min} and θ_{max} are joint limits, P_{min} and P_{max} are pressure bladder limits, ΔP_D is the change in desired pressure from the previous time step and ΔP_{max} is the maximum change in desired pressure per time step permitted. Because we use simplified pressure dynamics, the slew rate restrictions serve to prevent valve chattering.

In order to achieve angles away from the center position, different values of the pressure target must be sent to individual bladders. When bent to one side, the bladder on the inside of the bend must be at a lower pressure than the bladder on the outside of the bend. Individual pressure target values, P_T, are calculated based on the set stiffness pressure, P_S, according to the following equations:

$$\begin{aligned} P_T &= M\,|\theta_{goal}| + P_S \\ &\textbf{if } (\theta_{goal} \geq 0): \\ &\qquad P_{T,0} = P_S \\ &\qquad P_{T,1} = P_T \\ &\textbf{else}: \\ &\qquad P_{T,0} = P_T \\ &\qquad P_{T,1} = P_S \end{aligned} \tag{14}$$

where M is a constant which decreases target pressure as a function of θ_{goal}, P_T is a calculated target pressure for the bladder on the inside edge of a bend, P_S is a target pressure applied to the bladder on the outside edge of the bend as well as the operating point for P_T function. P_S can be varied during operation from one actual time step to the next, but is constant over an entire MPC horizon when the optimization is running.

We generated an efficient solver for our MPC problem using CVXGEN (see [31]), a web-based tool for developing convex optimization solvers. The optimization solver written in C and subsequent Python code that calls the solver can be run at 300 Hz predicting a trajectory horizon of T = 20 time steps (meaning it is predicting 0.067 seconds into the future).

Once solved, the first time step from the optimized trajectory of desired pressures is applied to the system and published over ROS. These desired pressures are received by the underlying pressure PID controller and valve position commands are then sent to the individual valves.

As described in Figure 4, the current pressure states are read and fed back into the pressure controller, while the angle states are fed into a Kalman filter. Both the current filtered angle states and pressure states are fed into the MPC controller.

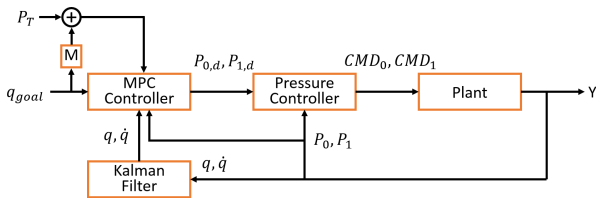

Fig. 4: MPC control diagram for joint and stiffness control.

V. Results and Discussion for Single Joint

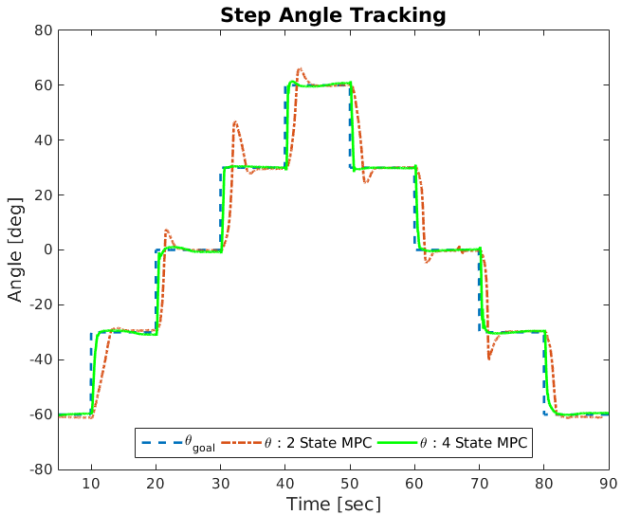

Fig. 5: Single Joint Control Response

A series of 30 degree step angle commands ranging from -60 to 60, changing in increments of 10 seconds, were sent to both the previous 2-state MPC controller used in [2], as well as our new 4-state controller described in this paper. The resultant θ of both controllers and commanded θ_{goal} are plotted over time in figure 5. Compared to the previous 2-state controller with only θ and $\dot{\theta}$, adding pressure states in the 4-state MPC controller significantly improved overall performance. As described in Table I, the new controller produced remarkably faster 90% rise time and 5% settling time while vastly decreasing percent overshoot. With angular step commands of only 30 degrees, the 2-state controller saw an average % overshoot of 24.408% or 7.3 degrees. The introduction of pressure dynamics reduced average % overshoot to just 2.561%, less than 0.77 degrees.

TABLE I: Performance comparison between previous controller and new controller

	Avg. Rise Time	Avg. Settling Time	Avg. % Overshoot
2 State MPC	1.319 sec	2.985 sec	24.408%
4 State MPC	0.556 sec	1.357 sec	2.561%
Improvement	137.24%	119.89%	853.21%

VI. Variable Stiffness Tracking

By varying the P_S value, one can effectively alter the stiffness during operation. While the exact relationship between

pressure and stiffness is currently unknown for our platform, pneumatic spring stiffness equations can be found in standard literature [32] as follows:

$$K = \frac{nPA^2}{V} \tag{15}$$

where K is the pneumatic spring stiffness, n is the polytropic exponent, P is pressure behind the diaphragm or piston, A is the cross sectional area, and V is the volume of the fluid. With the Grub platform as seen in Figure 1, once bladders have been filled, the non-expanding envelope prevents additional changes in volumes when angle remains constant. During operation, n and A are also assumed to be constant. Through these assumptions, an increase in P results in a direct linear increase in stiffness. How stiffness is changing is less clear when angle is also changing which we also demonstrate in this paper.

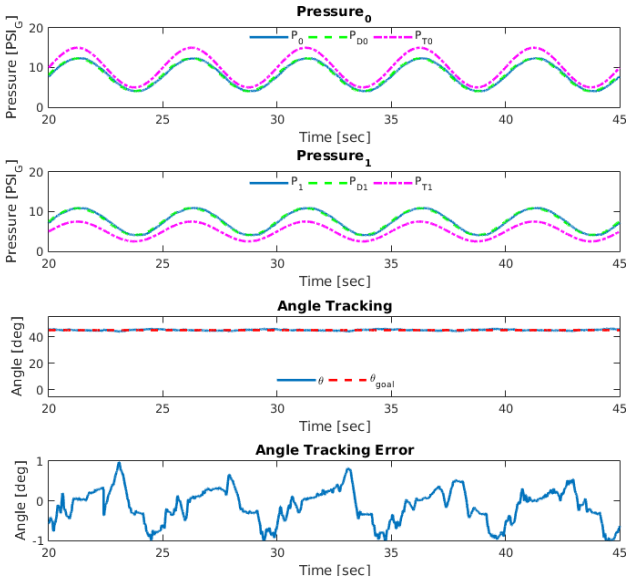

Fig. 6: Results for Grub holding a constant angle with a sinusoidal pressure command which is related to stiffness.

Given a constant θ_{goal} and a sinusoidal goal for pressure P_S, Figure 6 shows that the joint angle θ can be maintained while changing stiffness. P_S is a sinusoid with a frequency of 0.2 Hz and an amplitude of 5 psig operating around 10 psig. As seen in Figure 6, angle tracking error remains within 1 degree during changes of up to 200% in pressure in both bladders. As described in equation 15, changes in pressure whilst maintaining a constant angle indicates a direct change in stiffness. This test demonstrates the controller's ability to alter joint stiffness by as much 200% at a rate of 0.2 Hz while holding the joint angle relatively constant.

Even when varying the P_S value sinusoidally, θ_{goal} step commands can still be sent and followed since angle is given a higher weighting in our cost function. Figure 7 describes the resultant pressures and angle as θ_{goal} is stepped between 30 and -30 deg at a rate of 0.2 Hz and the P_S as a sinusoid at 0.13 Hz with an amplitude of 2.5 psig operating around 10 psig. Once the commanded angle has been reached, tracking

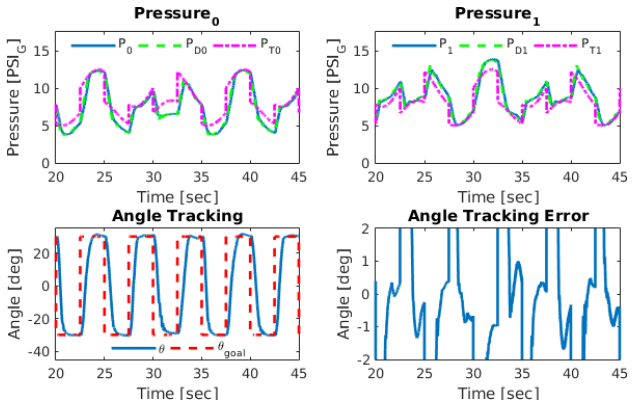

Fig. 7: Results for a series of step commands in joint angle with a commanded sinusoidal change in pressure.

error remains less than 2 degrees. The large steps in P_S are a function of changes in θ_{goal} as described in equation 14.

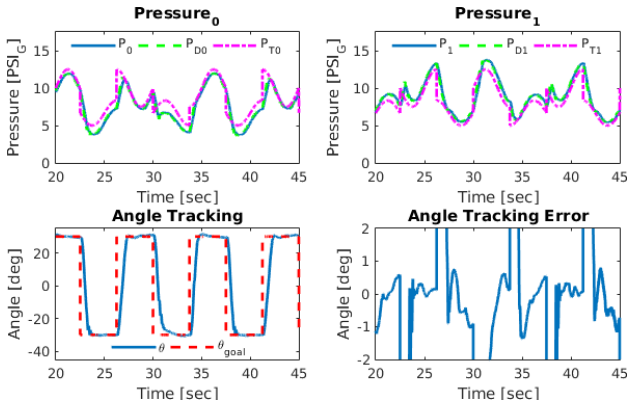

Fig. 8: Same results as Figure 7 but with frequencies reversed.

Figure 8 describes the same situation as Figure 7 with the commanded frequencies reversed, where θ_{goal} is stepped between 30 and -30 deg at 0.13 Hz and the P_S a sinusoid at 0.2 Hz with an amplitude of 2.5 psig operating around 10 psig. As in the previous case, non-transient tracking error remains less than 2 degrees despite large changes in pressures.

The tests shown in 7 and 8 describe the controller's ability to alter joint stiffness at variable rates while still achieving large changes in commanded angle. These tests demonstrate the viability of this controller and joint for use in applications such as those described in [9], where a variable impedance joint is used to maximize the effect of a robot controlled hammer impact, while still maintaining high overall system compliance needed for safety as described in [10].

VII. Conclusions

This paper describes a method for the modelling and control of an antagonistic inflatable soft robot platform while also demonstrating superior angle tracking performance compared to previous work and introducing variable stiffness

control. Soft robot platforms are inherently compliant, often with low inertia making them suitable for operation in environments where a rigid system would be dangerous to its surroundings. Through the methodologies described, an inflatable soft robot can effectively be controlled in both position and stiffness simultaneously. We expect that the work in this paper will allow robots equipped with similar soft joints to perform tasks safely with human-like performance in future research.

ACKNOWLEDGEMENT

We gratefully acknowledge support from the NASA ECF grant NNX14A051G, which has made this research possible. We also wish to acknowledge the hard work of Kevin Albert and his team at Pneubotics for collaboration in regards to our soft robot hardware platform.

REFERENCES

[1] D. Rus and M. T. Tolley, "Design, fabrication and control of soft robots," *Nature*, vol. 521, no. 7553, pp. 467–475, 2015.

[2] C. M. Best, J. P. Wilson, and M. D. Killpack, "Control of a pneumatically actuated, fully inflatable, fabric-based humanoid robot," in *Humanoids, 2015 IEEE-RAS International Conference on*. IEEE, 2015 (accepted).

[3] C. Loughlin, A. Albu-Schäffer, S. Haddadin, C. Ott, A. Stemmer, T. Wimböck, and G. Hirzinger, "The dlr lightweight robot: design and control concepts for robots in human environments," *Industrial Robot: an international journal*, vol. 34, no. 5, pp. 376–385, 2007.

[4] R. Bischoff, J. Kurth, G. Schreiber, R. Koeppe, A. Albu-Schäffer, A. Beyer, O. Eiberger, S. Haddadin, A. Stemmer, G. Grunwald, *et al.*, "The kuka-dlr lightweight robot arm-a new reference platform for robotics research and manufacturing," in *Robotics (ISR), 2010 41st international symposium on and 2010 6th German conference on robotics (ROBOTIK)*. VDE, 2010, pp. 1–8.

[5] A. Albu-Schäffer, C. Ott, and G. Hirzinger, "A unified passivity-based control framework for position, torque and impedance control of flexible joint robots," *The International Journal of Robotics Research*, vol. 26, no. 1, pp. 23–39, 2007.

[6] G. Pratt and M. Williamson, "Series elastic actuators," in *Intelligent Robots and Systems 95. 'Human Robot Interaction and Cooperative Robots', Proceedings. 1995 IEEE/RSJ International Conference on*, vol. 1, Aug 1995, pp. 399–406 vol.1.

[7] R. Ham, T. Sugar, B. Vanderborght, K. Hollander, and D. Lefeber, "Compliant actuator designs," *Robotics Automation Magazine, IEEE*, vol. 16, no. 3, pp. 81–94, September 2009.

[8] B. Vanderborght, A. Albu-Schäffer, A. Bicchi, E. Burdet, D. G. Caldwell, R. Carloni, M. Catalano, O. Eiberger, W. Friedl, G. Ganesh, *et al.*, "Variable impedance actuators: A review," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1601–1614, 2013.

[9] M. Garabini, A. Passaglia, F. Belo, P. Salaris, and A. Bicchi, "Optimality principles in variable stiffness control: The vsa hammer," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, Sept 2011, pp. 3770–3775.

[10] A. Bicchi and G. Tonietti, "Fast and "soft-arm" tactics [robot arm design]," *Robotics Automation Magazine, IEEE*, vol. 11, no. 2, pp. 22–33, June 2004.

[11] K. F. Laurin-Kovitz, J. E. Colgate, and S. D. Carnes, "Design of components for programmable passive impedance," in *Robotics and Automation, 1991. Proceedings., 1991 IEEE International Conference on*. IEEE, 1991, pp. 1476–1481.

[12] A. Bicchi, S. L. Rizzini, and G. Tonietti, "Compliant design for intrinsic safety: General issues and preliminary design," in *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, vol. 4. IEEE, 2001, pp. 1864–1869.

[13] G. Tonietti and A. Bicchi, "Adaptive simultaneous position and stiffness control for a soft robot arm," in *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, vol. 2. IEEE, 2002, pp. 1992–1997.

[14] O. Ivlev, "Soft fluidic actuators of rotary type for safe physical human-machine interaction," *2009 IEEE International Conference on Rehabilitation Robotics, ICORR 2009*, vol. 28359, pp. 1–5, 2009.

[15] I. Gaiser, R. Wiegand, O. Ivlev, a. Andres, H. Breitwieser, S. Schulz, and G. Bretthauer, "Compliant Robotics and Automation with Flexible Fluidic Actuators and Inflatable Structures," *Smart Actuation and Sensing SystemsRecent Advances and Future Challenges*, pp. 567–608, 2014.

[16] S. Sanan, J. Moidel, and C. Atkeson, "Robots with inflatable links," in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, Oct 2009, pp. 4331–4336.

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

[18] A. D. Marchese, K. Komorowski, C. D. Onal, and D. Rus, "Design and control of a soft and continuously deformable 2d robotic manipulation system," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 2189–2196.

[19] A. D. Marchese, R. K. Katzschmann, and D. Rus, "Whole arm planning for a soft and highly compliant 2d robotic manipulator," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*. IEEE, 2014, pp. 554–560.

[20] A. D. Marchese, R. Tedrake, and D. Rus, "Dynamics and trajectory optimization for a soft spatial fluidic elastomer manipulator," in *2015 IEEE International Conference on Robotics and Automation (ICRA). IEEE, New York*, 2015.

[21] S. Sanan, P. S. Lynn, and S. T. Griffith, "Pneumatic torsional actuators for inflatable robots," *Journal of Mechanisms and Robotics*, vol. 6, no. 3, p. 031003, 2014.

[22] J. F. Carneiro and F. G. de Almeida, "Reduced-order thermodynamic models for servo-pneumatic actuator chambers," *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, vol. 220, no. 4, pp. 301–314, 2006.

[23] X. Shen and M. Goldfarb, "Simultaneous force and stiffness control of a pneumatic actuator," *Journal of Dynamic Systems, Measurement, and Control*, vol. 129, no. 4, pp. 425–434, 2007.

[24] S. J. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, July 2003.

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

[26] D. H. Shim, H. J. Kim, and S. Sastry, "Decentralized nonlinear model predictive control of multiple flying robots," in *Proceedings. 42nd IEEE Conference on Decision and Control*, vol. 4. IEEE, 2003, pp. 3621–3626.

[27] C. Leung, S. Huang, N. Kwok, and G. Dissanayake, "Planning under uncertainty using model predictive control for information gathering," *Robotics and Autonomous Systems*, vol. 54, no. 11, pp. 898–910, July 2006.

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

[29] Y. Wang and S. Boyd, "Fast model predictive control using online optimization," *IEEE Transactions on Control Systems Technology*, vol. 18, no. 2, pp. 267–278, March 2010.

[30] A. Domahidi, A. U. Zgraggen, M. N. Zeilinger, M. Morari, and C. N. Jones, "Efficient interior point methods for multistage problems arising in receding horizon control," in *Proceedings of the 51st IEEE Conference on Decision and Control*, no. EPFL-CONF-181938, 2012.

[31] J. Mattingley and S. Boyd, "Cvxgen: a code generator for embedded convex optimization," *Optimization and Engineering*, vol. 13, no. 1, pp. 1–27, 2012. [Online]. Available: http://dx.doi.org/10.1007/s11081-011-9176-9

[32] B. W. Anderson, *The Analysis and Design of Pneumatic Systems*. Krieger Publishing Company, 3 2001.