

High performance pneumatics using Model Predictive Control

Vikash Kumar, Visak CV, Emanuel Todorov

Abstract— High force density, compact form factor and muscle like compliance properties makes pneumatics actuators quite desirable for robotic applications. However, the compressibility of the air not only throttles their bandwidth but also makes them harder to control. In this work, we leverage online trajectory optimization techniques to design high performance controllers for pneumatic actuators of ADROIT manipulation system [1]. These controllers use future predictions to act in anticipation. We show that the proposed controllers prepares the system ahead of time in order to achieve better performance and use less controls while doing so. Hardware results are presented on the pneumatic actuators of the ADROIT platform. Controller’s robustness is evaluated using hardware variants.

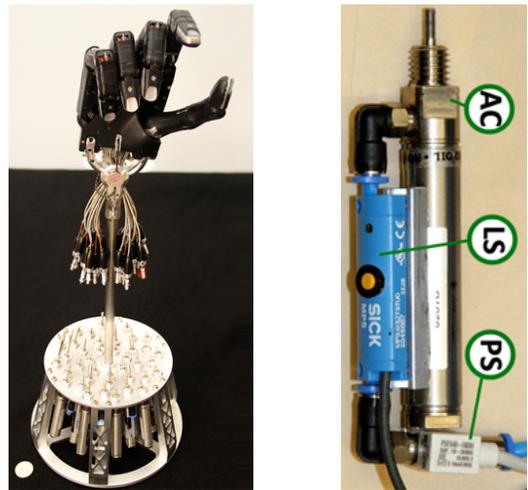
I. INTRODUCTION

Search for better actuators continues as robotic devices move from being hard and stiff, to being soft, nimble, agile and compliant. Substantial improvements are still needed (spanning safety, form factor, price point etc.) to kick-start the era of personal robotics. Fluid based actuators have several desirable properties relevant to the present and upcoming needs of such devices. They are inexpensive, mechanically simple with few moving parts and light weight with high force to weight ratio. Due to their compact form factor and high force density, they can be mounted directly on the moving degree of freedom (DoF), eliminating the need of gears and transmissions. Robustness, low friction and direct mounted on a DoF makes them ideal for force control applications.

However, the benefits of fluid based actuators have been overshadowed by the complexity of the control techniques they require. For applications where the dynamics of the fluid doesn’t need excitation, deployment of simple linear controllers have made these devices widely successful. They are the defacto actuators for industrial automation needs, commercial heavy weight equipments, load bearing and transmission mechanisms, power tools etc. For applications such as robotics, where the dynamics of the fluid needs to be accounted, controller design still remains a challenge. As the bandwidth of fluid based actuators (specially hydraulics) is improving, with the advancements in the value technology (MOOG), they are increasingly being used in the fast dynamics applications – Spot, Atlas [2], HyQ quadruped [3], Cheetah [4].

Bandwidth of the pneumatics actuators are lower than that of its hydraulic counterpart due to compressibility of the air, resulting in timescales of the order of $100ms$. On the other hand pneumatics is cleaner, lighter, quieter and easier to operate. They have properties similar to biological

muscles which make them quite desirable for biological applications seeking compliance – (1) they are back drivable and compliant at the mechanism level; (2) they have internal activation state (air pressure in the case of pneumatics, calcium concentration in the case of muscles) whose dynamics makes the entire system 3rd-order; (3) Pressure dynamics effectively introduces a low pass filter between the command and the force with timescales similar to that of biological muscles; (4) Much like biological muscles, co-contraction can be achieved by tuning the stiffness of the antagonistic counterparts. Compressibility is often considered a liability, but it is quite desirable at the same time. Spring dampers have long been used in mechanisms design for desirable passive behaviors. Pneumatics is the most generic form of an actuated spring damper with tunable gains. There is no doubt that – if efficient controllers are realised, pneumatics actuators (fluid actuators in general) will gain strong traction, specially in robotic applications.



(a) ADROIT hand mounted (tendons disconnected) on the pneumatic muscle assembly (b) Pneumatic muscle [AC: Actuator, LS: Length Sensor, PS: Pressure Sensor]

Fig. 1: Hardware setup

Our interest in synthesising complex behaviors [5] [6] [7] for biological systems and the desirable properties of pneumatic actuators led us into the development of ADROIT manipulation platform [1] – which is a modified Shadow Hand with custom pneumatic actuation (Figure 1). Our specific motivation for designing pneumatic controllers roots in the need of high performing low level controllers for ADROIT. In order to extract performance out of a system with large timescale and low bandwidth, effective planning

This work was supported by the US National Science Foundation. The authors are with the Departments of Computer Science & Engineering and Applied Mathematics, University of Washington, WA 98195, USA
E-mail: {vikash, todorov}@cs.washington.edu, visakc@uw.edu.

through the dynamics of overall system is required. Pneumatic dynamics modelling has received significant attention in the past – parametric [8] as well as physical models [9] have been developed. We build and improve on these pressure dynamics models as we exploit ideas from the field of model based trajectory optimization to design a high performance pneumatic controller for our system. This work is a natural continuation of our previous works in [10] and [8].

II. SYSTEM

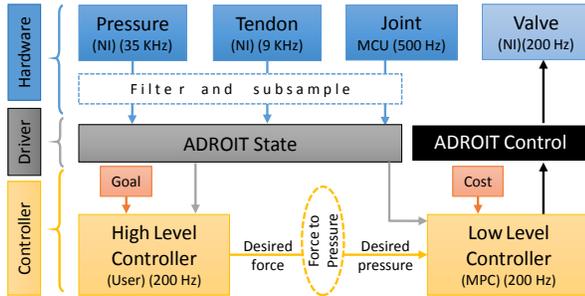


Fig. 2: System architecture

We first describe our system (Figure 2) with relevant technical details. Later, we present experimental observations that highlight properties relevant for controller design.

A. Hardware

ADROIT platform houses a 24 DoF bio-mimetic hand (Figure 1a). 20 Dofs are independently actuated using 40 antagonistic *pneumatic muscle actuation units* (Figure 1b) while the rest (finger’s distal joints) are coupled. ADROIT’s control-loop runs are 200 Hz, sampling 135 sensors and commanding 48 actuators using a 9 slot NATIONAL INSTRUMENTS 3U PXIe-1078 chassis¹(NI). The chassis is configured from the *computational unit*² using a PXIe-PCIe 798MB/s bandwidth data channel.

A low friction, low stiction double acting pneumatic cylinder (*AIRPEL M9 37.5NM*, *AIRPORT Corporation*, [11]) forms the muscle actuation unit. As muscles can only pull, the rear chamber of the cylinder is left passive and open to the atmosphere. Each unit has stroke length of 37.5mm, can produce 42N of force at 100 PSI and weighs about 37.5 grams. *Airpel* replaces traditional pneumatic seals with “air seals” in order to achieve low friction and low stiction conditions. While air seals are desirable for smooth movements, the constant leakage from the seals makes the non linear pneumatic dynamics further chaotic and hard to model.

The muscle actuation unit is observed using two sensors. The pressure inside a muscle unit is observed using a solid state (*SMC PSE540-IM5H*³) pressure sensor. The muscle excursion is measured as piston stroke length using a magnetic length sensor (*SICK MPS-032TSTU0*⁴). The pressure

¹The chassis has a data rate of 1Gb/s, 250 Mb/s bandwidth per slot.

²12 cores 3.47GHz Intel(R) Xeon(R) processor with 12GB memory running Windows x64

³The pressure sensor can measure up to 10^6 pascal with < 2% resolution, < $\pm 0.7\%$ linearity, < $\pm 0.2\%$ repeatability and weighs 4.6 grams

⁴The length sensors can measure up to 32 mm excursions with 0.05mm (about 2mm in practice) resolution, < 0.3mm linearity, < 0.1mm repeatability, The sensor has sampling time of 1ms and can sense movements up to a max speed of 3m/s.

sensors are sampled at 32KHz and the length sensors are sampled at 9KHz. High frequency components of the sensor readings are filtered out, using low pass filters, before they are made available for use. High sampling rate allows us to perform data filtering without introducing significant delays. Joint angle sensors are sampled using an embedded micro-controller(MCU) at 500Hz.

A high flow 5/3 festo proportional valve is used to drive the front chamber of the muscle actuation unit. The proportional valve (MPYE-5-M5-010-B from *FESTO*) has a flow rate of 100 liters/min at 87 PSI, bandwidth of 125 Hz and weighs 290gms.

B. Hardware exploration

We explored the asymptotic pressure response (Figure 3) of the ADROIT’s pneumatic muscle unit by subjecting the volume locked assembly to voltage step changes from either extremes. The input voltage affects the flow through the valve. Ideally, the asymptotic response of the valve should be a step function. Near the zero point (center of the input range) of the valve, it is partially connected to both the source (compressor) and the sink (atmosphere), resulting in intermediate asymptotic values. The pressure dynamics is significantly slower around the zero point as illustrated by the time required to reach the asymptotic pressure curve. Impulse response obtained from the step changes, for two different tube lengths, are illustrated in Figure 4. Higher latencies for the longer tubes can be attribute to the time required for pressure wave to travel through them.

The steady state responses (Figure 5) of the muscle unit was obtained as slow varying voltage signals sweep the input range of a volume locked assembly. Normalization wrt the compressor pressure reveals its dependence on the compressor pressure, which changes as we use the system. The fat belly around the center illustrates the nondeterministic nature of the spool (hence the pressure) around the valve’s zero position. A pneumatic valve spends most of its time around the nominal positions. The nondeterminism around the nominal position might pose significant challenge for controller design.

The muscle assembly was subjected to chirp signals to study the frequency response of the over-all pneumatic circuit. The critical frequency was found to be around 25Hz.

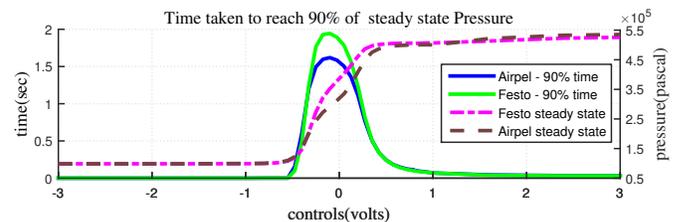


Fig. 3: Asymptotic pressure response of Adroit’s actuators

III. PNEUMATICS

A. Cylinder

A cylinder (Figure 7) is a device with two chambers, separated by a moving bore. Each chamber has an orifice, called port, that connects it to a pressure reservoir. The reservoir with pressure higher (usually air compressor) than that of the chamber is called the source and the corresponding

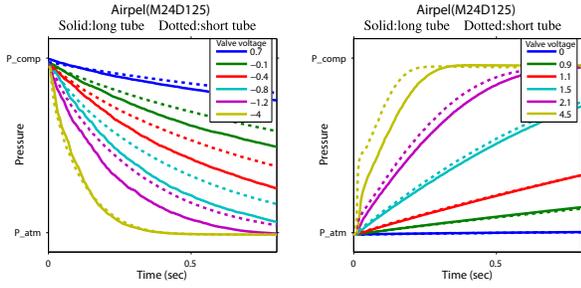
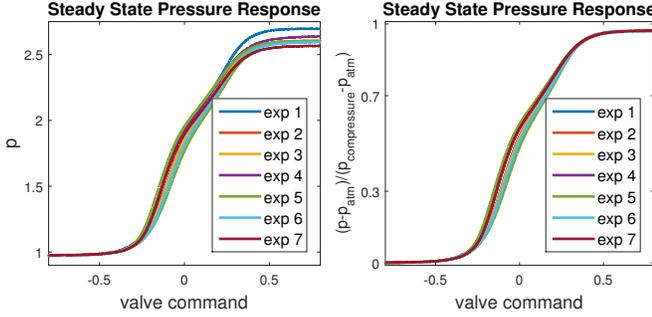


Fig. 4: Impulse responses for different tube lengths



(a) Raw pressure measurements (b) Normalized relative-pressure

Fig. 5: Steady state pressure response

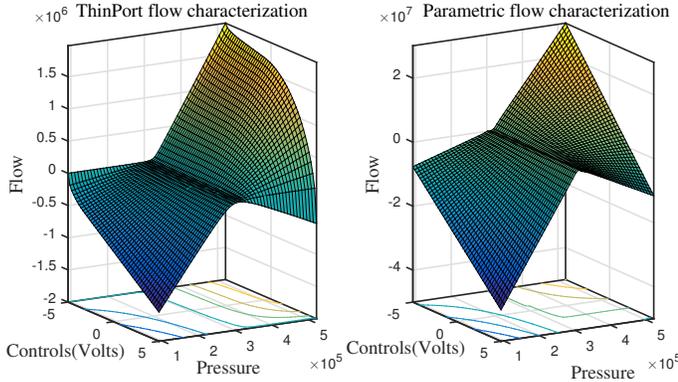


Fig. 6: Flow characterization of ADROIT's muscle unit

port is called the supply port. The reservoir with pressure lower (usually atmosphere) than that of the chamber is called the sink and the corresponding port is called the exhaust port.

B. Valve

The pressure inside a chamber can be regulated by connecting a valve to it's port. A valve is a mechanical device that connects a chamber to multiple reservoirs (usually two - a source and a sink) and regulates the cross-sectional area opening between them. Two commonly used types of valves are (a) binary on/off valves which use Pulse Width Modulation scheme for area modulation and (b) proportional solenoid valves. Unlike binary valves, proportional valves are expensive but provide fine grained control over the port area resulting in smooth operation. A solenoid actuated spool moves inside the valve in response to the input command thereby smoothly changing the port's cross-sectional area (Figure 8).

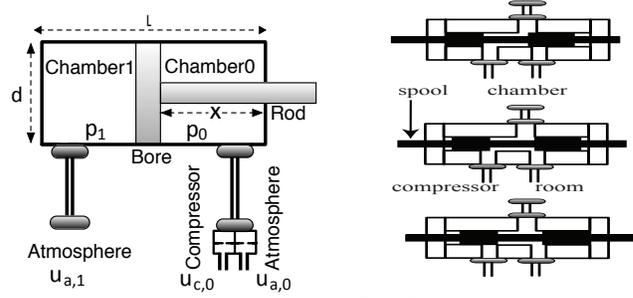


Fig. 7: Schematic representation of a pneumatic cylinder with a port controlled by a valve and the other port open to the atmosphere

Fig. 8: Schematic representation of a proportional pneumatic valve. (Top: supply port open. Center: Both exhaust & supply port partially open. Bottom: Exhaust port open)

C. Thin Plate Port pressure dynamics model

The thin plate port model (Figure 9) explains the flow of fluid (air) through a port as a function of the area of the orifice, the upstream pressure p_u and the downstream pressure p_d . Assumption being that the plate connecting the chambers is thin, the port area is small, fluid in use is a perfect gas, both the chambers are at the same temperature and the flow is isentropic. Thin plate port model is common in the pneumatics control literature. [9] is an excellent article that reviews and builds on previous work in the light of real time control applications. Modifications, with justified assumptions, have been proposed to handle computational complexity while accounting for prediction accuracy necessary for such applications. The model is briefly summarized below for completeness.

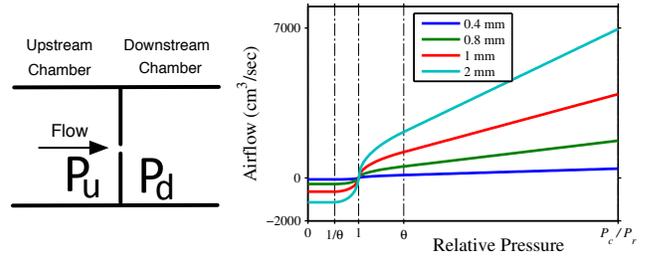


Fig. 9: (a)Thin Plate. (b)Thin Plate Flow Function. The 4 curves show the flow rate for 4 orifice diameters. The downstream pressure p_d is kept at room pressure $P_r \approx 100kPa$, and p_u is varied from 0 to the compressor pressure $P_c \approx 620kPa \approx 90psi$. The linear regime is outside the dotted vertical lines.

1) *Thin plate flow function*: Thin plate flow function $\phi(p_u, p_d)$ describes the mass flow \dot{m} through a port per unit port area a .

$$\dot{m} = a \cdot \phi(p_u, p_d) \quad (1a)$$

$$\phi(p_u, p_d) = \begin{cases} z(p_u, p_d) & \text{if } p_u \geq p_d \\ -z(p_d, p_u) & \text{if } p_u < p_d \end{cases} \quad (1b)$$

$$z(p_u, p_d) = \begin{cases} \alpha p_u \sqrt{\left(\frac{p_d}{p_u}\right)^{\frac{2}{\kappa}} - \left(\frac{p_d}{p_u}\right)^{\frac{\kappa+1}{\kappa}}} & \text{for } p_u/p_d \leq \theta \\ \beta p_u & \text{for } p_u/p_d > \theta \end{cases} \quad (1c)$$

The physical constants κ , α , β and θ are described in the Appendix (Section VII). Note that the flow function (Figure 9) is continuously differentiable and is linear in the upstream pressure p_u for $p_u > \theta p_d$.

2) *Two port chamber*: The net mass flow into a chamber with two ports can be described as

$$\dot{m}(p, a_c, a_r, p_c, p_r) = a_c \phi(p_c, p) - a_r \phi(p, p_r) \quad (2)$$

where a_c, a_r are the orifice areas connecting the chamber to the compressor and room respectively, and p_c, p_r are the respective constant pressures. Figure 10 shows this function to be monotonically decreasing, which corresponds to stable dynamics that converge to a steady-state pressure p_{ss} given by $a_c \phi(p_c, p_{ss}) = a_r \phi(p_{ss}, p_r)$.

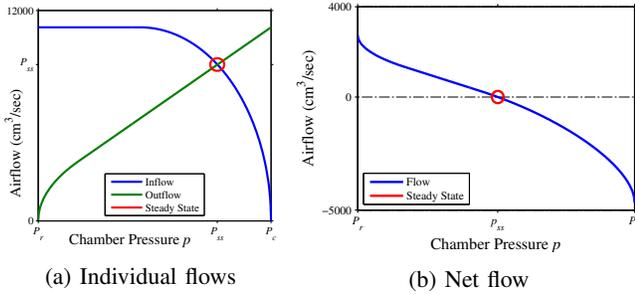


Fig. 10: Two port chamber

3) *Chamber pressure dynamics model*: Using ideal gas laws, two port net flux (Equation (2)) and polytropic process with isothermal constraints, pressure dynamics of a single chamber can be written as

$$\dot{p}(p, u, v, \dot{v}, p_c, p_r) = \kappa \frac{RT}{v} \dot{m} - \kappa \frac{\dot{v}}{v} p \quad (3a)$$

$$\dot{m}(p, u, p_c, p_r) = a_c(u) \phi(p_c, p) - a_r(u) \phi(p, p_r) \quad (3b)$$

where v is the volume of the chamber, \dot{v} is the rate of change of that volume and κ , R and T are physical constants (see Appendix (Section VII)). The first term in the pressure dynamics equation (Equation (3a)) captures the dynamics dictated by the net flow into the chamber, while the second explains the dynamics due to the changing chamber volume.

D. Parametric pressure dynamics model

The over all time scale of a pneumatic system depends on the valve dynamics, the delays in the pneumatic circuit, and the chamber volume (Equation (3a)). While the pneumatic is slow in general, a fully retracted cylinder with a low chamber volume can have a time constant of the order of microseconds. As a result, a generic parametric form of pressure dynamics will need extremely small time step or a variable time integrator in order to integrate the dynamics forward.

$$\dot{p}(p, u, v, \dot{v})|c = (s(u, v, \dot{v})|c - p) \cdot r(u, v, \dot{v})|c \quad (4)$$

Equation (4) summarises the parametric model that we zeroed in our prior work [8]. Note that it is linear with respect to p , allowing allows us to perform analytical integration. As a result it is stable for arbitrary time step, and computationally inexpensive to evaluate. The function $s()$ has a unit of pressure and represents the steady state pressure. The $r() > 0$ has a unit of inverse of time and represents the rate of change. The system is differentiable if $s()$ and $r()$ are differentiable.

Here we present a slightly modified parametric model (note that the parametric form is still the same) that is more general and delivers better predictions.

$$\dot{p}(p, u, v, \dot{v}, p_c, p_r)|c = (c_b + c_s \text{sigmoid}(c_3 \hat{u} + c_4 \hat{u}^3) - p) \frac{c_7 (\text{sabs}(\hat{u}) + c_8 \hat{u}) + c_9}{\tau} + \frac{c_5 \dot{v} p}{\tau} \quad (5)$$

Where $c_b = (p_c + p_a)/2$ represents the mid point pressure, $c_s = (p_c - p_a)/2$ represents the pressure range, $\hat{u} = u - c_2$ is centered control voltage, $\tau = 1 + c_6$, $\text{sigmoid}(z) = z/\sqrt{1+z^2}$ and $\text{sabs}(z) = \sqrt{z^2 + c_7^2} - c_7$. Note that c_s and c_b are no longer fixed constants based on maximum and minimum pressure (as treated in [8]). They now are inputs to the model that need to be updated using real time pressure readings.

IV. MODEL IDENTIFICATION

A. Thin port model

The dependence of port area on input voltage $a_c(\mathbf{u})$, $a_r(\mathbf{u})$ is not known and is not possible to measure directly. We obtain this dependency using numerical optimization. A volume locked chamber ($\dot{v} = 0$) was subjected to step voltage changes starting from each voltage extreme. We solve a small optimization problem to recover $\{a_c^{u_i}, a_r^{u_i}\}$ pair for each voltage step. Standard curve fitting techniques was used to fit a parametric form (gompertz function 6) to the resulting area pairs. Refer [12] for details.

$$\begin{aligned} \{a_c^{u_i}, a_r^{u_i}, v_o\} = \underset{a_c, a_r, v_o}{\text{argmin}} \{ & \dot{P}_{measured}^{u_i} - \\ & \kappa \frac{RT}{v + v_o} (a_c^{u_i} f(P_c, p) - a_r^{u_i} f(p, P_r)) \} \\ a(\mathbf{u})|c = & c_0 + c_1 e^{-c_2 e^{-c_3(\hat{u} + c_4)}} \end{aligned} \quad (6)$$

Where c_0 is the minimal leak area of the port, $\hat{u} = (\mathbf{u} - c_1)$ is the centered control voltage. Figure 11c outlines the recovered area pairs and the parametric area fits. Figure 11a and 11b presents the \mathbf{p} and $\dot{\mathbf{p}}$ predictions using the identified model.

B. Parametric model

1) *Steady state response*: In order to capture the steady state response $s(u, v, \dot{v}; c)$ of our pneumatic system, a volume locked chamber was subjected to slow varying control input. The slow varying control input drags the system equilibrium along as it sweeps the entire input range. We fit a steady state model of the form $p_{steady\ state} = c_b + c_s \text{sigmoid}(c_3 \hat{u} + c_4 \hat{u}^3)$ to the resulting data. Figure 13 summarizes the system response and the model predictions wrt time(left) and input voltage (right). Unlike the predictions (black curves) from our old model [8] with fixed c_s and c_b , the predictions from the

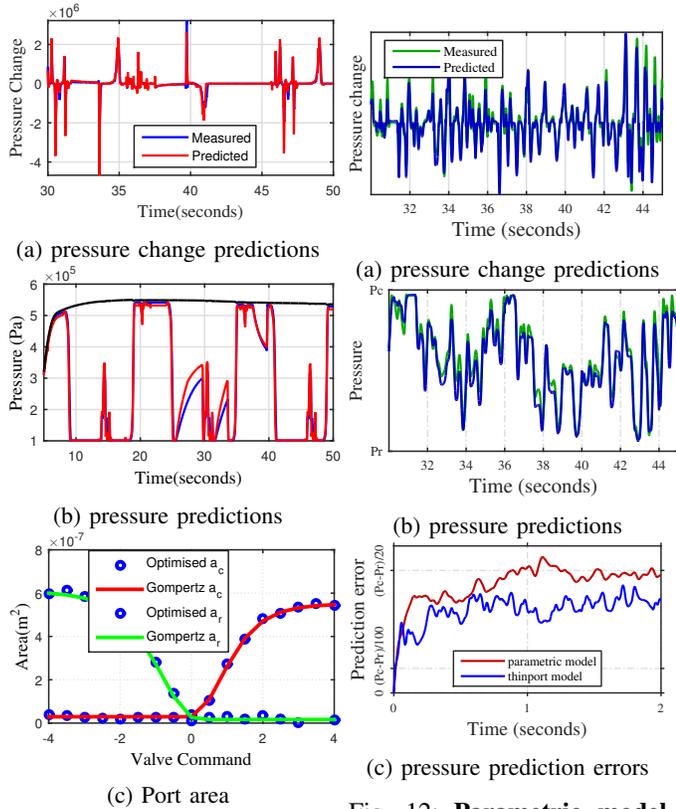


Fig. 11: **Thin Port model identification**

model using real time c_b and c_s readings nicely captures the dependence of the steady state pressure on the compressor pressure.

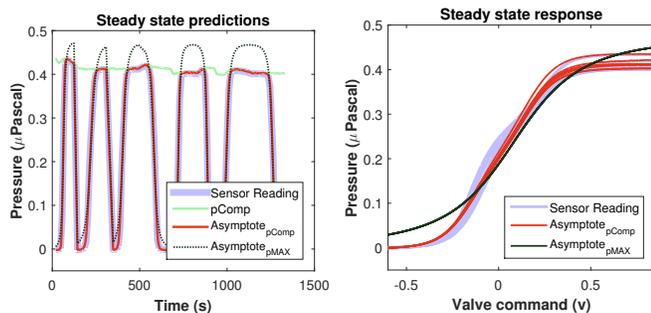


Fig. 13: **Steady state predictions**. Unlike the predictions $Asymptote_{pMAX}$ from model using constant c_s and c_b , the predictions $Asymptote_{pComp}$ from model using real time c_s and c_b nicely captures the dependency of the steady state pressure on compressor pressure.

2) *Rate*: For data diversity, the cylinder was overdriven by applying external forces to the cylinder’s piston (in order to diversify the chamber’s volume v and the rate of change of volume \dot{v}) while it was subjected to random signals for rate $r(u, v, \dot{v}; c)$ identification. As we aren’t explicitly modelling the valve dynamics, the random signals were low passed filtered (30 Hz butterworth) below the valve’s critical frequency (125 Hz) before they were executed on the hardware.

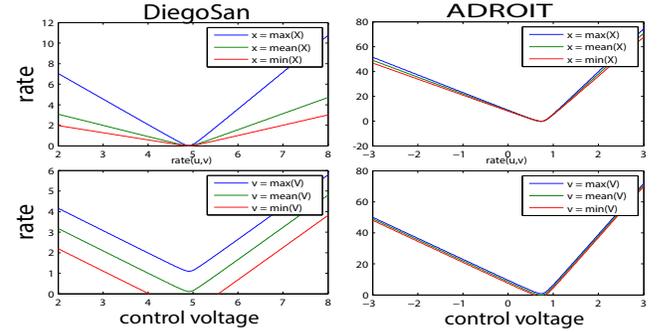


Fig. 14: Rate comparison between Adroit and DieogSan’s [8] actuators. Adroit’s actuators have much higher rates and are fairly insensitive to bore position(x) & velocity(v)

The rate (Figure 14) reveals valuable information about the dynamics, and its sensitivity, of our pneumatic system. A system with high flow rate valve and small chamber volume should have fast dynamics. Which is indeed the case – the rate of the identified model grows rapidly with control input. While fast dynamics means better throughput, it also means that the system is highly sensitive to the control input and timing delays. Small control input will be enough to drive the system (Section VI) and the wide proportional regime of our expensive valves will be hard to exploit. We also observe that rate is relatively insensitive to the volume v and volume rate \dot{v} . This can be attributed to the fact that the chamber has small volume and stroke length. As a result, the changes induced by the bore movement is insignificant. This insensitivity can be exploited to simplify the pressure model⁵.

$$\dot{p}(u, p_c, p_r)|c = (c_b + c_s \text{sigmoid}(c_3 \hat{u} + c_4 \hat{u}^3) - p)(c_7(\text{sabs}(\hat{u}) + c_8 \hat{u}) + c_9) \quad (7)$$

V. HIGH PERFORMANCE PNEUMATIC CONTROLLER

High level motion planners have seen significant advancements in recent times. We now have planners that can plan in real time for high DoFs robots [5] [6] [7]. The planners are no longer short sighted and impulsive. They reason about the movements in longer time horizon and opt for instant replanning if the execution derails from the plan. Most of these improvement primarily stem from the fact that we now have more computational resources at our disposal, that are being utilized for better understanding of the system via simulation, data analysis, optimization, learning etc.

Pneumatic systems are known to have significant latencies. Compressibility of the air results in long time constants (Figure 4) that throttles the bandwidth of the system. Local feedback controllers are unable to deal with these challenges to deliver performance required by a dynamical system. The proposed controller does that by leveraging the dynamics model of the system to unfold the system forward in time and reasoning about the actions over a longer time window (called horizon). The predictive capability from the pneumatics models (from Section IV), lets our controller

⁵We haven’t simplified the models for our use case at this point as the simplification doesn’t hold for the arm cylinders which are significantly bigger than the hand cylinders. From the software architecture standpoint, we would like to treat all the cylinder identical, if possible.

Algorithm 1 Trajectory Optimization

Input: Dynamics $f(\mathbf{x}, \mathbf{u})$, running cost $l_i(\mathbf{x}_i, \mathbf{u}_i)$, final costs $l_f(\mathbf{x}_N)$, current state \mathbf{x}_0 , warm-start sequence \mathbf{U} .

Output: Locally optimal control sequence $\hat{\mathbf{U}}$

- 1: *Rollout:* Integrate \mathbf{U} to get the initial trajectory $(\mathbf{x}_i, \mathbf{u}_i)$
 - 2: *Derivatives:* Get derivatives for l_i and l_f
 - 3: *Backward pass:* Calculate the second order approximations of $\mathbf{V}(\mathbf{x}, i)$. Obtain a search direction $\Delta\mathbf{U}$ as 2nd-order solution to the Equation (9)
 - 4: *Forward Pass:* Rollout \mathbf{x}_0 and $\mathbf{U} + \alpha\Delta\mathbf{U}$ forward with different line search parameters $0 < \alpha < 1$ to pick the winner
-

make anticipatory preparation for actions well in advance, resulting in improved performance. We outline the necessary details of our controller below

A. Controller design

Instead of having a myopic view presented by the immediate desired pressure value \mathbf{p}_{dest} , our controller design focuses on designing a policy using a macroscopic view as presented by a desired pressure trajectory over a time horizon $T = N * dt$. Given the desired pressure trajectory $\mathbf{P}_{\text{des}} = [\mathbf{p}_{\text{dest}}, \mathbf{p}_{\text{dest}+dt}, \dots, \mathbf{p}_{\text{dest}+(N-1)dt}]$, the goal of our controller is to find the appropriate policy, in the space of valve commands, that guides the system through $\hat{\mathbf{p}}_{\text{des}}$ over time. We pose the policy design problem as a finite horizon optimal control problem and deploy standard trajectory optimization techniques, iterative-LQG [13] in this case, for efficient solutions in real time.

B. Finite horizon optimal control

Given the state \mathbf{x}_i , controls \mathbf{u}_i at time i , let the discrete time dynamics of a system be described as $\hat{\mathbf{x}} = \mathbf{x}_{i+1} = f(\mathbf{x}_i, \mathbf{u}_i)$. The finite horizon optimal control problems can be posed as – starting from an initial state \mathbf{x}_0 solve for a time varying control law $\hat{\mathbf{U}}(\mathbf{x})$ that minimizes the cumulative sum of the *running cost* $l_i(\mathbf{x}_i, \mathbf{u}_i)$ and the *final cost* $l_f(\mathbf{x}_N)$ along a trajectory.

$$\hat{\mathbf{U}}(\mathbf{x}) = \underset{\mathbf{U}}{\operatorname{argmin}} \sum_{i=0}^{N-1} l_i(\mathbf{x}_i, \mathbf{u}_i) + l_f(\mathbf{x}_N) \quad (8)$$

C. Trajectory optimization

Our choice of trajectory optimizer (iterative-LQG [13]) solves the problem above using the principles of dynamic programming. The value $\mathbf{V}(\mathbf{x}, i)$ corresponding to a state \mathbf{x} at time i indicates the minimal cost incurred to optimally solve the problem for remaining $N - i$ steps. The final values function $\mathbf{V}(\mathbf{x}, N) = l_f(\mathbf{x}_N)$ is just the final cost. Dynamic programming principle reduces the minimization over a sequence of controls \mathbf{U}_i to a sequence of minimization over a single control, proceeding backwards in time

$$\mathbf{V}(\mathbf{x}, i) = \min_{\mathbf{u}} [l(\mathbf{x}, \mathbf{u}) + \mathbf{V}(f(\mathbf{x}, \mathbf{u}), i + 1)] \quad (9)$$

The algorithm is outlined in Algorithm 1. We recommend [13] for in-depth analysis and [5] [6] [7] for more applications.

D. Model Predictive Control (MPC)

The goal of this work is to abstract out the pneumatic actuators as ideal torque actuators. The controller in consideration will form the low level controller of the ADROIT actuation system. The goal is to design low level controllers that emulate ideal force controllers using non-linear pneumatic actuator. As the system is always in motion for the low level controllers, the \mathbf{P}_{des} (specified by the high level controllers) is constantly flowing with time, even when the high level controller is demanding a constant torque. If the system dynamics was linear and the cost we deploy was quadratic we will find the optimum in single iteration of the trajectory optimization Algorithm 1. The pressure dynamics is highly nonlinear and the cost we use is not quadratic either. Therefore, the algorithm needs multiple iterations to work through the linear approximation of the dynamic and quadratic approximation of the cost to converge on an optima.

As the system is always in motion, we deploy the trajectory optimization in a Model Predictive Control (MPC) fashion. Which means instead of solving for the optimum, starting from our current state estimates, we will only take few iteration of the algorithm, improve the policy for the current estimates, and then opt for an estimates update. There are multiple rational behind this choice – (1) The model used for planning will never be perfect and we will never reach the true optimum even if we have it; (2) Solving for optimum is computationally expensive. Fast policy update provides better performance by dragging the system closer to the optimum with each update; (3) The low level controller has no controls over the demands of the high level planner. The high level planner can decide to abruptly change the plan. The best choice for the low level controller is to respect its demands as soon as possible.

E. State and System Dynamics

The state $\mathbf{x} = [\mathbf{p}, \mathbf{w}]^T$ of our system is of dimensionality $2n_a$, where n_a is the number of actuators in the system (40 for ADROIT hand). It consists of *cylinder pressure* \mathbf{p} and *valve memory* \mathbf{w} , which is the controls (valve voltages) from the previous time step. System dynamics can be written as

$$\mathbf{x} = [\mathbf{p}, \mathbf{w}]^T; \mathbf{x}_{t+dt} = f(\mathbf{x}_t, \mathbf{u}_t)|_{\mathbf{v}, \dot{\mathbf{v}}} = [\mathbf{p}_{t+dt}, \mathbf{u}_t]^T \quad (10)$$

Where \mathbf{p}_{t+dt} is obtained using the pressure dynamics models of likes outlined in Section III-C (by euler integration of Eq3) and III-D (Eq4 can be exactly integrated forward). It is important to note that the volume \mathbf{v} and volume rate $\dot{\mathbf{v}}$ is not a part of the system's state \mathbf{x} . They are required as inputs to the pressure dynamics and enter the system as external sensor readings. This is possible because, given the volume \mathbf{v} and volume rate $\dot{\mathbf{v}}$, the pressure dynamics is independent of the dynamics of the piston (and the external load its driving). Similarly, given the net force on the piston, the piston dynamics is independent of the pressure dynamics. Since the sensor for measuring chamber pressures and piston stroke lengths are cheap and reliable (and available in ADROIT), we exploit this independence.

F. Cost Function

The running cost of the system is of the form

$$l_i(\mathbf{x}_i, \mathbf{u}_i) = \alpha_p \cdot \|\mathbf{p}_{des_i} - \mathbf{p}_i\| + \alpha_u \cdot \|\mathbf{u}_{des_i} - \mathbf{u}_i\| + \alpha_w \cdot \|\mathbf{u}_i - \mathbf{w}_i\| \quad (11)$$

The final cost $l_f(\mathbf{x}_N)$ is same as the running cost. The justification for including the *valve memory* \mathbf{w} in the state can be found in the last term of the cost i.e. $\alpha_w \cdot \|\mathbf{u}_i - \mathbf{w}_i\|$. This term forces the optimizers to pick similar controls for the time steps adjacent to each other. Thus ensuring smoothness in the final control sequence that executes on the hardware. Without this term, the optimizers are free to pick arbitrary sequence of chattering controls, as long it minimizes the cost objective. Rapidly chattering control sequence keeps the pneumatic valves always active and on the edge of their critical frequency, resulting in decreased performance and hardware wear. α_w can be adjusted to tune the amount of smoothing required.

G. Simulation results

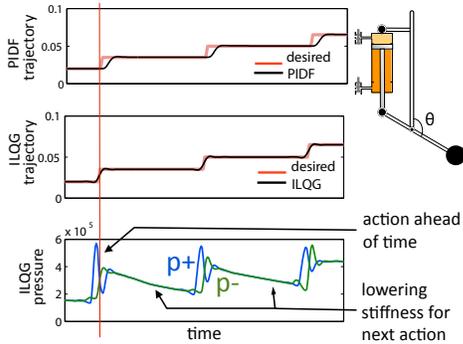


Fig. 15: Controller comparison

We use a 1 DoF platform with pneumatic actuator to compare the performance of different controllers under identical conditions. We prefer simulated platform for this comparison as it allows us to study a controller's performance independent of the unpredictability of the real world. Figure 15 outlines the performance of our trajectory optimization based controller with the PIDF controller [10]. Unlike the myopic nature of the PIDF which acts only after the change is demanded, anticipatory nature of the iLQG controllers results in lower latency and better performance. Close attention to the pressure highlights pneumatic activity ahead of time and stiffness lowering when possible.

VI. CONTROLLER PERFORMANCE

While identifying the parametric model (Figure 14) we learned that, because of the small volume of our cylinders, the pressure dynamics is quite insensitive of piston position and velocity. Leakage of the ADROIT pneumatic muscle unit further complicates the nonlinear dynamics. To circumvent these complications, we first present our controller's performance (Figure 16 & 17) on a leak free volume locked cylinder (Festo DSN104P).

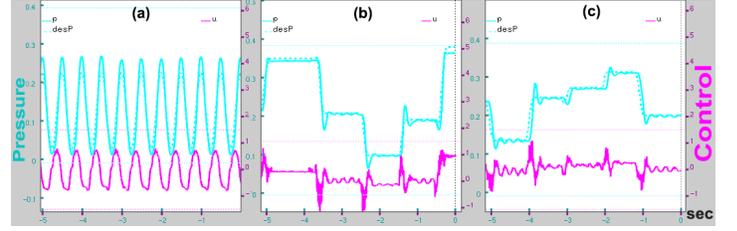


Fig. 16: pressure tracking for a volume locked FESTO cylinder+ Thin Port + short tube. (a) 2Hz (b) Random (c) Instantaneous change

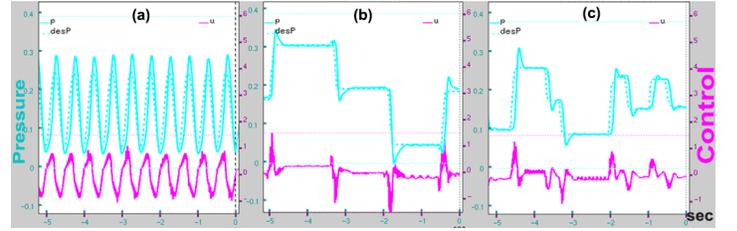


Fig. 17: pressure tracking for a volume locked FESTO cylinder+ parametric model+ short tube. (a) 2Hz (b) Random (c) Instantaneous change

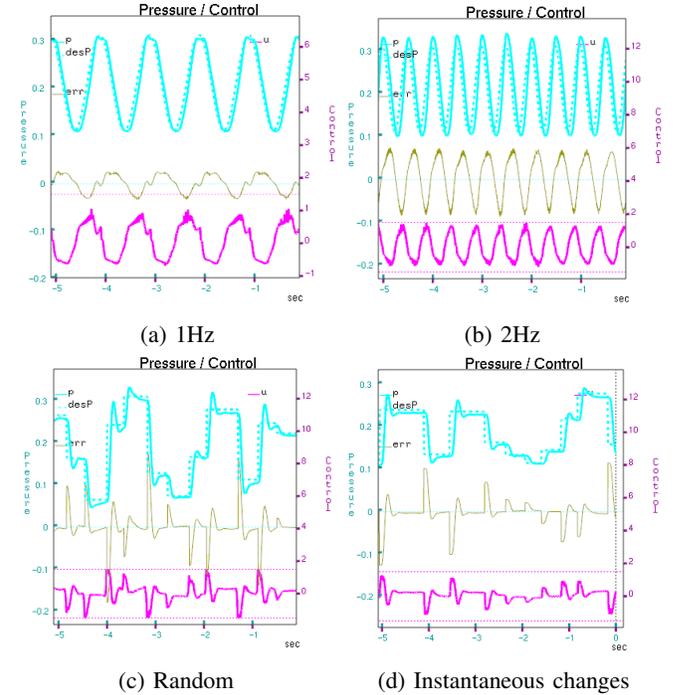


Fig. 18: Controller's performance on a free to move (Airpel M09D37) cylinder with leak, while using pressure dynamics model learned on a leak-free volume-locked (Festo DSN104) cylinder.

A. Trajectory Optimization and MPC details

To have all the optimization entities roughly in the same scale, we use pressure in Mega Pascal. The pneumatic controller and the hardware driver are two separate processes communicating using sockets. At the onset of a control loop, the controller receives the current estimates from the hardware driver and responds with the controls queried using the current

policy. The iLQG is parallelized across 8 physical cores. The planning horizon is 200 ms ($N = 100$, $dt = 0.005s$) long. The policy lag is between 7.5 to 10 ms. The cost coefficients are $\alpha_p = 1$, $\alpha_u = 10^{-4}$, $\alpha_w = 10^{-4}$.

B. Controller's robustness

In order to demonstrate (Figure 18) the effectiveness of our approach and the robustness of the resulting controller, we artificially induce modelling errors in the system, by replacing the volume locked leak free cylinder with moving cylinder with leak (ADROIT pneumatic muscle unit). In order to evaluate our controller's ability to deal with abrupt change in the demands from the high level planner (Figure 16c & 17c), we subject the controller to instantaneous random changes in entire desired future sequence starting from current time. These instantaneous changes renders the current policy useless due to the local nature of the controller's linear feedback policies. The optimizer acts on these instantaneous demands only after the next policy is pushed to the hardware.

C. Real world considerations

Results on a hardware platform is a combination of sound theory, good implementations and few practical considerations that get highlighted only in the real world. For reproducibility of the results, we outline our experiences learned while deploying the system.

- 1) *Timing is the key to performance*: The inherent pneumatic latency doesn't leave a lot of leeway for the controllers to drive the system through dynamic movements. Every effort was taken to minimize the latency of the system to the extent possible. Special attention was required for hardware-controller clock synchronization, data communication latencies, sensor data filtering latencies, estimates and policy lag, optimizer's computational needs and execution latencies (as the policy is queried over the network).
- 2) *Pneumatic nonlinearities*: The discontinuity in port areas as valve's spool moves across zero position induces severe non-linearities. The noise in the spool movement makes these nonlinearities unpredictable and hard to model. Special attention was required while modelling the pneumatic dynamics near the valve's zero point.
- 3) *Smooth operations*: As the valve chatters around its zero position making minor improvements, the pneumatic nonlinearities around zero position severely degrade the controllers performance. Introduction of valve memory \mathbf{w} and cost $\alpha_w \cdot \|\mathbf{u}_i - \mathbf{w}_i\|$ for smoothing the controls significantly improves controller's performance.
- 4) *Operational regime*: Smoothness requirements promote the controller to leverage the inherent smooth dynamics of the system over abrupt changes in the value controls resulting in plans with small and smooth controls. In practice, the controller uses only 40% of the valves input range. In light of this observation, control limits were enforced on the controller using [14] and iterative model learning methods were used to refit the pneumatic models to the restricted range of operation. Restricting the operation range, allows us to run the control loop faster (200Hz) as the critical frequency of a valves is much wider for small movements.

- 5) *feedback*: Our system tends to be more stable delivering better performance with partial feedback (around 20%).

VII. APPENDIX

The physical constants in Eq. Equation (1) are given by:

$$\alpha = C \sqrt{\frac{2M}{ZRT} \frac{\kappa}{\kappa-1}} \quad \theta = \left(\frac{\kappa+1}{2}\right)^{\frac{\kappa}{\kappa-1}}$$

$$\beta = C \sqrt{\frac{\kappa M}{ZRT} \left(\frac{2}{\kappa+1}\right)^{\frac{\kappa+1}{\kappa-1}}}$$

Gas Molecular Mass	M	0.029 for air, Kg/mol
Temperature	T	K°
Universal Gas Constant	R	$8.31 (Pa \cdot m^3)/(mol K^\circ)$
Discharge coefficient	C	0.72, dimensionless
Compressibility Factor	Z	0.99 for air, dimensionless
Specific Heat Ratio	κ	1.4 for air, dimensionless
Mass Flow	\dot{m}	Kg/s
Pressure	p	$Pascals$
Area	a	m^2

TABLE I: Parameters and units of the thin-plate port model.

REFERENCES

- [1] V. Kumar, Z. Xu, and E. Todorov, "Fast, strong and compliant pneumatic actuation for dexterous tendon-driven hands," in *IEEE International Conference on Robotics and Automation*, 2013.
- [2] Boston Dynamics, <http://www.bostondynamics.com>.
- [3] C. Semini, N. G. Tsagarakis, E. Guglielmino, M. Focchi, F. Cannella, and D. G. Caldwell, "Design of hyq—a hydraulically and electrically actuated quadruped robot," *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Eng.*, 2011.
- [4] A. Spröwitz, A. Tuleu, M. Vespignani, M. Ajallooeian, E. Badri, and A. J. Ijspeert, "Towards dynamic trot gait locomotion: Design, control, and experiments with cheetah-cub, a compliant quadruped robot," *The International Journal of Robotics Research*, 2013.
- [5] V. Kumar, Y. Tassa, T. Erez, and E. Todorov, "Real-time behaviour synthesis for dynamic hand-manipulation," in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2014.
- [6] Y. Tassa, T. Erez, and E. Todorov, "Synthesis and stabilization of complex behaviors through online trajectory optimization," in *Conference on Intelligent Robots and Systems*, 2012, pp. 4906–4913.
- [7] T. Erez, K. Lowrey, Y. Tassa, V. Kumar, S. Kolev, and E. Todorov, "An integrated system for real-time model predictive control of humanoid robots," in *International Conference on Humanoid Robots*, 2013.
- [8] T. Yuval, T. Wu, J. Movellan, and E. Todorov, "Modeling and identification of pneumatic actuators," in *International Conference on Mechatronics and Automation (ICMA)*, 2013.
- [9] E. Richer and Y. Hurmuzlu, "A high performance pneumatic force actuator system: Part I nonlinear mathematical model," *Journal of dynamic systems, measurement, and control*, vol. 122, no. 3, 2000.
- [10] E. Todorov, C. Hu, A. Simpkins, and J. Movellan, "Identification and control of a pneumatic robot," in *3rd IEEE International Conference on Biomedical Robotics & Biomechatronics (Biorob)*, '10.
- [11] Airpel, <http://www.airpot.com/antistiction-air-cylinders.html>.
- [12] C. V. V. Kumar and V. Kumar, "Pneumatic modelling for adroit manipulation platform," manuscript under review, Humanoids'16.
- [13] E. Todorov and W. Li, "A generalized iterative lqg method for locally-optimal feedback control of constrained nonlinear stochastic systems," in *Proceedings of the American Control Conference*. IEEE, 2005.
- [14] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 1168–1175.