The Dissertation Committee for Nicholas Arden Paine
certifies that this is the approved version of the following dissertation:

High-Performance Series Elastic Actuation

Committee:

Luis Sentis, Supervisor

Sriram Vishwanath, Supervisor

Jonathan Valvano

Andreas Gerstlauer

Peter Stone

Ashish Deshpande

High-Performance Series Elastic Actuation

by

Nicholas Arden Paine, B.S.E.E., M.S.E.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

August 2014

Dedicated to my family who have been a constant source of support during the many years of my education.

Acknowledgments

I must first acknowledge my advisors, Professors Luis Sentis and Sriram Vishwanath, who have guided me through the process of graduate study. Sriram, thank you for seeing potential in me and encouraging me to pursue both the Master's and Doctoral degrees. Luis, your passion for mathematics was my first true introduction to model-based control and has helped me form a strong theoretical foundation in my work. Thank you also for the many opportunities you set up for me. Your positive reinforcement of my work was certainly a primary factor in its success.

To my family, especially these past few months, thank you for giving me an opportunity to focus on my work and for always encouraging me. I attribute much of the success I have had to the values I have learned from you. Llewyn, the late nights playing Borderlands were something I always looked forward to during the long days of writing.

To the many excellent teachers I have had here at UT, especially Professors Jonathan Valvano, Ashish Deshpande, Peter Stone, Andreas Gerstlauer, and John Pearce, your passionate teaching helped me discover what it really means to understand topics like microcontrollers, mechanisms, software, and analog circuits, which are the basis for much of this work.

To Sehoon, talking about dynamics and control on the bus heading to

the Gingerman is something I will always remember. Your expert knowledge and friendship have meant a lot to me.

To Josh, James, Nic, Chris, Vienny, and the rest of the Valkyrie team, it was a once-in-a-lifetime chance to work on such a capable and energetic team. Thank you for letting us be a part of it.

To William Bard, thanks for your interest in my work and for all of the insightful conversations we have had.

To Scott Allen, Paul Landers, Danny Jares and Aaron Frost, without your help I would have gone through many more design iterations before making something that actually worked. Thank you for helping me realize my ideas.

To Liang and Drew, your guidance and knowledge turned the Pharos Lab and the Proteus robots into a great success. Thank you for teaching me the values of software engineering, documentation, and managing large, complex projects.

Thank you to Tim and Vikram for the many discussions we had on various mechanical and electrical topics.

To Robin and Marco, your work inspires and challenges me to push to new levels. Thank you for your friendship and guidance.

To Seth, thank you for your involvement in the Pharos lab and also for lending me a couch when I was stranded working in Houston. We'll have to make it back to Nobi again sometime.

To all of my peers, Matt, Pius, Somudro, Mike, Donghyun, Ye, Kenan, Kwan Suk, Alan, Travis, Dorothy, Becca, Prashant, Gray, Rachael, it has been a great opportunity for me to experience graduate school life with you all. Thank you for the many insightful discussions during the long group meetings.

To Prashant, Gray, and Rachael, the late night beers, darts, and samosas gave me something to look forward to at the end of each week while writing this. Thanks for the good times!

And finally to Chad, Karl, Ty, Greg, Paul, Joel, Letti and my many other Austin friends, you all are what have made these past few years truly enjoyable. Thank you.

Nicholas Arden Paine

The University of Texas at Austin
August 2014

High-Performance Series Elastic Actuation

Nicholas Arden Paine, Ph.D.
The University of Texas at Austin, 2014

Supervisors: Luis Sentis
Sriram Vishwanath

Mobile legged robots have the potential to restructure many aspects of our lives in the near future. Whether for applications in household care, entertainment, or disaster response, these systems depend on high-performance actuators to improve their basic capabilities. The work presented here focuses on developing new high-performance actuators, specifically series elastic actuators, to address this need. We adopt a system-wide optimization approach, dealing with factors which influence performance at the levels of mechanical design, electrical system design, and control. Using this approach and based on a set of performance metrics, we produce an actuator, the UT-SEA, which achieves leading empirical results in terms of power-to-weight, force control, size, and system efficiency. We also develop general high-performance control techniques for both force- and position-controlled actuators, some of which were adopted for use on NASA-JSC's Valkyrie Humanoid robot and were used during DARPA's DRC Trials 2013 robotics competition.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation and Objectives

Series Elastic Actuation (SEA), the practice of intentionally adding compliance to actuator drivetrains (Pratt and Williamson, 1995), is a technology which has recently experienced widespread growth in many areas associated with robotics. SEAs have been adopted for use in the fields of human-safe compliant manipulation (Edsinger-Gonzales and Weber, 2004; Diftler et al., 2011; Fitzgerald, 2013; Lens and von Stryk, 2013), exoskeletons (Veneman et al., 2005; Kwa et al., 2009; Ragonesi et al., 2011; Safavi et al., 2011), active human orthosis (Kong et al., 2009, 2012; dos Santos et al., 2013; Pott et al., 2013) and prosthesis (Au et al., 2007; Zhu et al., 2010), haptic devices (Oblak and Matjačić, 2011; Gillespie et al., 2014), laparoscopy (Basafa et al., 2009), and microrobotics (Tokatli and Patoglu, 2011), among others. Perhaps the most popular use for SEAs, though, is in the field of legged robotics. The first application of series elastic actuators was in this field, owing largely to its force sensing capabilities and its ability to absorb impacts without damage to the rest of the actuator drivetrain (Pratt et al., 1997).

Since this early work, many researchers have proposed alterations or

enhancements to SEAs, often for use in the field of legged robotics (Robinson et al., 1999; Hurst et al., 2004; Au et al., 2007; Curran et al., 2008; Hurst and Rizzi, 2008; Tsagarakis et al., 2009; Thorson and Caldwell, 2011; Hutter et al., 2011b; Parmiggiani et al., 2012; Grimes and Hurst, 2012; Slovich et al., 2012; Lee et al., 2013; Wang et al., 2013). Most of this work intends to compete directly with alternative actuation technologies, such as pneumatic (Niiyama et al., 2007, 2010), hydraulic (Raibert et al., 2008; Semini, 2010), or rigid electric actuators (Sakagami et al., 2002; Kaneko et al., 2011; Ito et al., 2014). Driven by the prodigious demands of legged robots on their actuators, a central but often implied theme unifying all of this work is the need for high-performance[1] actuation.

Given the significant interest in the performance of actuators for legged robots, it is no surprise that the performance of these legged systems has recently reached new levels. Boston Dynamics is at the center of this development, consistently impressing the world with periodic YouTube[2] videos showing Bigdog (Raibert et al., 2008), Atlas, Wildcat, and other robots performing dynamic legged locomotion. The development of many of these systems has been funded by the Defense Advanced Research Projects Agency (DARPA) for military applications. More recently though, DARPA created the DARPA Robotics Challenge (DRC) (DARPA, 2013), an international competition between teams who built and programmed their robots to complete tasks related

[1] A definition of "performance" is provided in Section 1.2

[2] `https://www.youtube.com/`

to civilian disaster response scenarios. The DRC Trials 2013 was held in December 2013 near Miami Florida and was largely considered a success (IEEE Spectrum, 2013).

To further such developments, the work in this thesis focuses on improving the state of the art in legged actuation, or more specifically, in high-performance series elastic actuation. We cannot argue that SEAs are the best actuation technology for all legged robotics applications. However, as we will show in the next section, SEAs do present an attractive collection of actuation traits, especially when overall system efficiency is a design target. Given the tight coupling between series elastic actuation and legged robotics, it is no stretch to believe that progress in one of these areas will benefit the other. At the same time, progress in the performance capability of SEAs will certainly benefit other areas of robotics as well.

1.2 State of the Art in High-Performance Actuators

Why study series elastic actuators? Hydraulic actuators are known to possess higher power output per unit weight, and rigid electromagnetic actuators can achieve better dynamic position tracking accuracy, so where do compliant actuators fit in? Each actuation technology has its particular niche where it exceeds other technologies. It is only after understanding the particular needs of legged robotics and the metrics by which this application can be characterized that the benefits of series elastic actuators become plainly apparent. We present that list of metrics here.

1. **Power-to-weight ratio.** This is the most fundamental metric because it defines a robot's ability to move its mass and thus defines dynamic performance limits. This metric is well recognized in characterizing actuator performance (Hollerbach et al., 1992).

2. **Torque-to-weight ratio.** Large torques and forces are required to support the many degrees-of-freedom found in legged systems and is often a design target (Hollerbach et al., 1992; Seok et al., 2012). This metric is necessary in our performance definition in order to exclude actuators like standalone motors which produce high power (thus satisfying the power-to-weight metric) only at high speeds.

3. **Power-to-volume ratio.** This is an important metric for versatile robots because it relates to the size or bounding volume of an actuator. Versatile robots require many actuated degrees-of-freedom (DOFs) in a small space. Even if an actuator has excellent properties otherwise, if it cannot be designed to fit within a confined volume, it will have limited use in highly articulated robots.

4. **Efficiency.** An often neglected metric, defined as the ratio of mechanical energy produced by a robotic system divided by the potential energy it consumes, this metric defines a robot's ability to operate for long periods of time on a fixed energy source. DARPA recently recognized poor actuation efficiency as a major bottleneck in many current legged robots whose specific resistance (a measure of efficiency) is two orders of

magnitude worse than legged animals (DARPA, 2012).

5. **Force control accuracy.** A more generic metric which applies more to actuators than to robots as a whole, controllable force resolution and force tracking bandwidth together determine the degree to which an actuator may be considered an ideal source of force/torque (Kong et al., 2009). Force/torque sources are widely used by many control algorithms in robotics and therefore is an important characteristic to consider (Khatib, 1987; Pratt et al., 2001; Hogan, 1984; Sentis et al., 2013; Hutter et al., 2013). Outside of legged robotics, accurate force control is an essential characteristic in any field where a robotic system must physically interact with a human (Kong et al., 2009, 2010). Additional details on this metric can be found in (Robinson, 2000).

6. **Mechanical Robustness.** As it is considered here, robustness is the sensitivity of an actuation technology to the damaging effects of impact forces. While such forces may be mitigated through means of control in some situations, in others, especially when there is environmental uncertainty, a sudden impact should not cause physical harm to the actuator.

7. **Energy source density.** For mobile robots, both the size and weight of on-board energy are important factors to consider. Therefore, actuators requiring potential energy that is not available in a dense form can be difficult to use in these mobile systems.

There are many ways of assessing the performance of a system based on these metrics. The most basic (and arguably least informative) way of assessing a system using these metrics is to score each metric based on the system design *specifications* (a.k.a. datasheet performance). This method is less informative than other methods because it only takes into consideration the mechanical system. In this work, our main emphasis is on *system level* performance, that is, the performance of the mechanical, electrical, software, and control systems operating together. While we largely focus on system performance at the actuator level, the same techniques and metrics can easily be applied to entire robots as well.

As a visual aide in understanding system level performance, we introduce the graphic in Figure 1.1. This image shows the combined effect of the various systems in a robot on its overall performance. The width of each stack represents performance as defined by the metrics above. Our objective, then, is to maximize the width of the realized performance by maximizing the performance capability of the bottom layer (i.e. the mechanical system) and then minimizing the loss of each subsequent layer. In practice, this "loss minimization" translates into designing fast and high-precision embedded systems, creating stable high-performing controllers, etc.

1.2.1 Actuation Technologies

In the following subsections, we briefly introduce various actuation technologies which have successfully demonstrated high-performance capability in

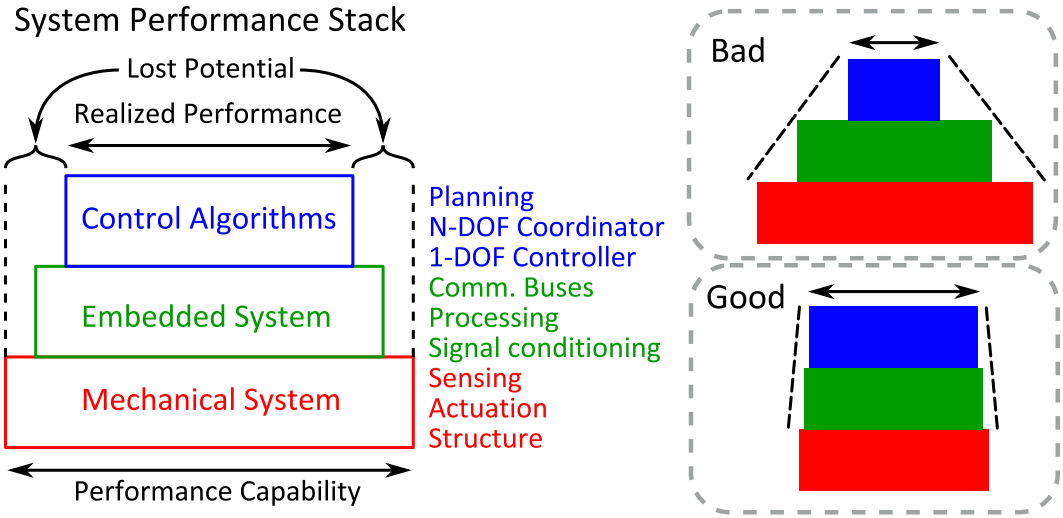

Figure 1.1: This visual representation of system performance shows how each layer in a robotic system affects overall system performance. Each layer must build upon the layer beneath it, its width representing performance capability. A perfect layer is one which preserves the performance capability of the layer beneath it. A well- designed and controlled system will maximize the mechanical system performance given some cost function (money, time, etc.) and then minimize lost performance potential of the layers above it.

legged robotic systems. Our analysis of these systems is based on the performance metrics listed above.

1.2.1.1 Pneumatic

Pneumatic actuators possess impressive power-to-weight properties, being capable of producing up to 10,000 watts of mechanical power per kilogram of mass (Hannaford and Winters, 1990). They can also produce large forces per unit weight, 8,000 N/kg and above (Festo, 2014; Villegas et al., 2012). These capabilities have been well displayed in legged robots driven by pneu-

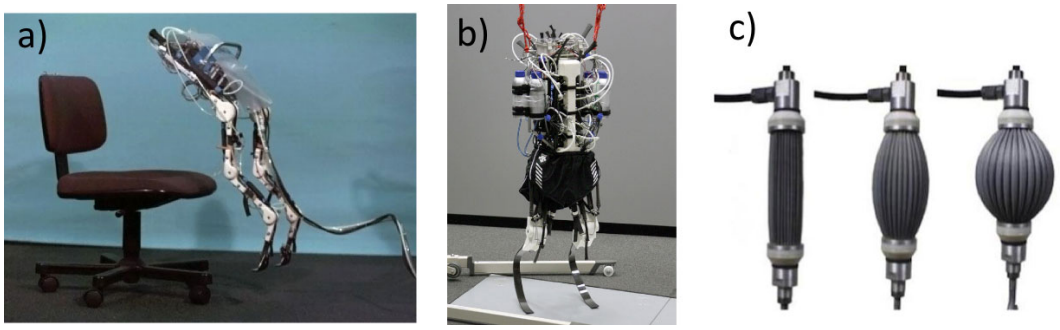

Figure 1.2: A few examples of high-performing robots with pneumatic actuators: a) Mowgli[1] (Niiyama et al., 2007) b) Athlete[2] (Niiyama et al., 2010), c) example fluidic muscle type pneumatic actuator[3].

matic actuators (see Figure 1.2) (Niiyama et al., 2007, 2010). In (Niiyama et al., 2007), Mowgli is able to jump above 50% of its body height and in (Niiyama et al., 2010) Athlete takes two strides at 1.2 m/s.

The drawbacks of this technology are most readily apparent in the energy densities of compressed air (see Table 1.1). To become feasible for mobile applications, robots based on this technology would likely require either batteries or gasoline as the primary energy source and would use pumps to compress air for use by the actuators. This added bulk and weight largely detract from the weight-saving benefits of such actuators. Additionally, pneumatic actuators can often be difficult to control due to the compressibility of air (Yeh et al., 2010) and McKibben type muscles suffer from limited contractile distance per unit length (Hannaford and Winters, 1990).

[1] http://www.isi.imi.i.u-tokyo.ac.jp/ niiyama/projects/images/Mowgli_jumping.jpg

[2] http://www.designboom.com/cms/images/andrea11/athleterobot02.jpg

[3] http://mech.vub.ac.be/multibody/topics/img/Muscles33states.jpg

Table 1.1: A Comparison of Energy Densities (Lancaster, 2002)

Technology	Volumetric (Wh/l)	Gravimetric (Wh/kg)
Gasoline	9,000	13,500
Propane	6,600	13,900
Ethanol	6,100	7,850
Lithium	250	350
Flywheel	210	120
Lead Acid	40	25
Compressed Air	17	34

1.2.1.2 Hydraulic

Excelling in power-to-weight, torque-to-weight, power-to-volume, robustness, and when powered by gasoline engines, energy source density, hydraulics have shown great promise in legged robotics applications (see Figure 1.3) (Raibert, 1986; Raibert et al., 2008; Semini, 2010; Semini et al., 2010). Because of the incompressibility of hydraulic oil, these systems are more easily controlled than their pneumatic counterparts. Additionally, these actuators naturally produce mechanical energy in a high-force low-speed form, meaning no gearboxes are required. They are centrally powered by a single powerful pump, meaning the valves (one per joint) must only modulate this energy. The routing of hydraulic fluid naturally carries heat away from each actuator allowing for a centralized cooling system as well.

Unfortunately, hydraulic systems are inefficient. For example, one of the sources of loss in robots like Bigdog is the leakage of its hydraulic valves. This leakage has been calculated to be 2000 W (Hutter, 2013), a large amount

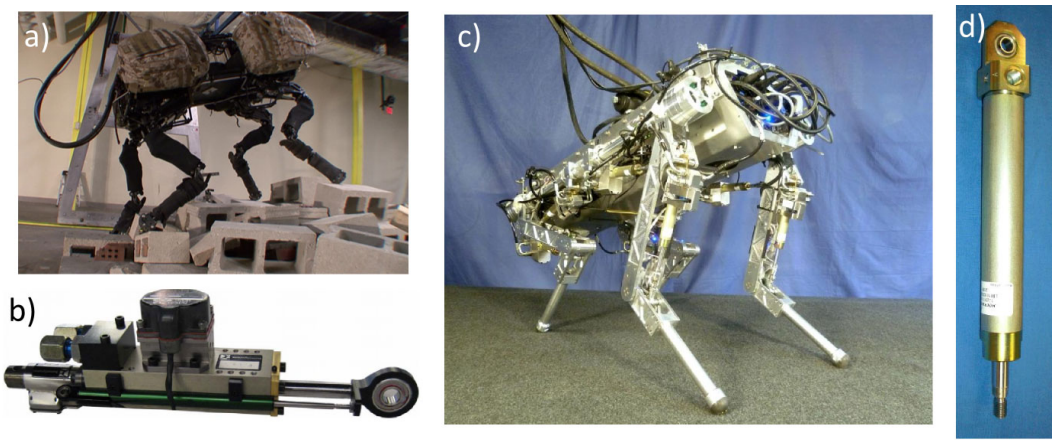

Figure 1.3: A few examples of high-performing robots with hydraulic actuators: a) Bigdog[4] (Raibert et al., 2008) b) Bigdog hydraulic actuator, c) HyQ[5] (Semini, 2010) d) HyQ hydraulic actuator.

of energy to be lost with no kinetic benefit to the robot. Additionally, to benefit from the high energy density of gasoline (see Table 1.1), internal combustion engines must be used which are also highly inefficient. The inefficiency of hydraulic machines can be calculated at the robot level in terms of specific resistance and cost of transport (Gabrielli and von Karman, 1950; Kuo, 2007). Several studies have found that hydraulic legged robots fare far worse than their natural counterparts, around two orders of magnitude, in terms of locomotion efficiency (DARPA, 2012; Seok et al., 2013).

Is it possible to do better? While not as capable in terms of power-to-weight and torque-per-weight, electric actuation typically greatly outperforms

[4] http://www.bostondynamics.com/img/BigDog_ClimbRubble.png

[5] http://spectrum.ieee.org/img/hyq%20hydraulic%20quadruped%20robot%20iit-1-1319646263602.jpg

hydraulic actuation in terms of system efficiency. There are two notable examples of this in legged robotic systems. In (Gregorio et al., 1997) the authors show an order of magnitude improvement in specific resistance for their electric version of a similar hydraulic legged robot. This improvement in efficiency comes at a slight degradation in performance; the hydraulic system can locomote at two to three times higher velocity. In (Seok et al., 2013), the authors again demonstrate an electric legged robot running at approximately 1/30th of the cost-of-transport required by an (albeit much more capable) hydraulic quadruped. Such examples give confidence that electromagnetic-based actuators offer competitive and efficient alternatives to hydraulic-based systems.

1.2.1.3 High-Gear-Ratio Electromagnetic

Compared to the power-to-weight output of a human muscle (from 50 W/kg to 200 W/kg (Hunter and Lafontaine, 1992)) modern electromagnetic motors today are capable of substantial power output (666 W/kg without forced cooling (Maxon, 2014a), 2600 W/kg with liquid cooling (Remy Motors, 2014)). However, at low speeds, electromagnetic motors produce low torque, peaking around 6 Nm/kg without forced cooling (Hollerbach et al., 1992). To make use of the high power capability of electromagnetic motors and to improve torque per unit weight, large speed reductions are often used.

High-performance robotic systems based on this technology are quite common and have achieved significant results (see Figure 1.4). Honda's Asimo robot (Sakagami et al., 2002) was one of the first examples of a self-contained

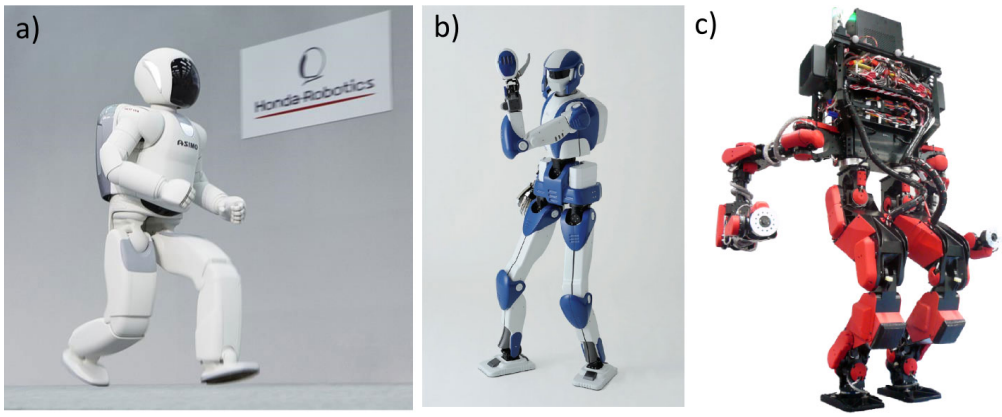

Figure 1.4: A few examples of high-performing robots with high-gear-ratio electromagnetic actuators: a) Asimo[6] b) HRP-4[7] and c) Schaft[8].

humanoid robot capable of walking in structured environments and performing household tasks. More recently, a newer version of Asimo has performed dynamic maneuvers such as hopping on one foot (Honda, 2011). With similar goals, AIST has developed the Humanoid Robot Platform 4 (HRP-4) which has demonstrated walking and basic manipulation and sensing tasks (Kaneko et al., 2011). The most recent version of this technology was revealed at the DRC Trials, where team Schaft demonstrated robust rough terrain locomotion and high joint load capability (Schaft, 2013; Ito et al., 2014). All of these actuation technologies perform well in terms of power-to-weight, most notably (Schaft, 2013; Ito et al., 2014).

These systems almost exclusively use brushless direct current (BLDC)

[6] http://spectrum.ieee.org/img/all-new-honda-asimo-1367895285823.jpeg
[7] http://www.aist.go.jp
[8] http://theroboticschallenge.org/sites/default/files/schaft.png

motors coupled with harmonic drive speed reducers. Harmonic drives are well known to be one of the most compact solutions for obtaining large speed reductions (Sensinger and Lipsey, 2012). Therefore, most of these approaches score well in terms of power-to-volume. Harmonic drives do not always operate efficiently, however, creating a drawback for their application to high-performance drives. According to the manufacturer's catalog, the efficiency of harmonic drives may be as poor as 25% and only increases above 80% when optimal combinations of input shaft speed, ambient temperature, gear ratio, and lubrication are present (Harmonic Drive, 2014). Another potential disadvantage of this type of actuation scheme is that no torque sensing feedback is available at each joint.

In contrast to the aforementioned technologies, DLR's Biped and Rollin' Justin robots (Figure 1.5) incorporate torque feedback at each actuator while maintaining the use of the harmonic drive/motor combination. This is accomplished without adding bulk by instrumenting the flexspline of the harmonic drive with strain gauges (Albu-Schäffer et al., 2007a). The flexspline deflects when a torque is applied to the joint, creating a measurable change in resistance of the strain gauges. Impressive results have been shown with this system in terms of dynamic response (power-to-weight) and force controllability metrics in ball catching (Bauml et al., 2011) and coordinated impedance control experiments (Wimbock et al., 2007).

While large speed reductions benefit the power-to-weight ratio of electromagnetic actuators, they can create significant problems in terms of actu-

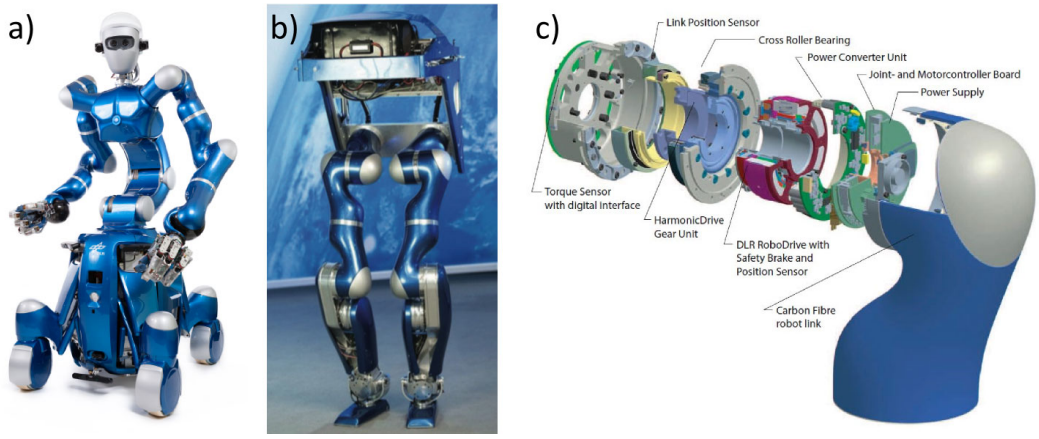

Figure 1.5: DLR's Rollin' Justin[1] (a) and the DLR Biped[2] (b) exhibiting their lightweight robot actuation technology (c) which utilizes torque sensing at each joint (Albu-Schäffer et al., 2007a).

ator robustness. Large gear ratios amplify motor inertia by the square of the reduction ratio, giving these actuators very large output inertia. From a mechanical impedance standpoint, rigid, fixed objects (the ground, trees, walls) act as admittances (Hogan, 1984). Because of their high output impedance, high-gear-ratio electromagnetic actuators also behave as admittances. According to (Hogan, 1984), stable dynamic interaction between systems are best achieved when the two systems compliment one another (i.e. an impedance interacting with an admittance or vice versa). Collisions between high-gear-ratio actuators and rigid objects violate this basic principle, often resulting in large contact forces and damage to either the actuator or the rigid object.

[1] `http://www.dlr.de/rmc/rm/en/Portaldata/52/Resources/images/institute/robotersysteme/rollin_justin/general/thumbs/Rollin_justin_feb2013_380.png`

[2] `http://spectrum.ieee.org/img/DLR%20Biped-1320692491791.png`

1.2.1.4 Low-Gear Ratio Electromagnetic

Without the use of a large speed reduction, electromagnetic actuators suffer from poor torque density. Such devices have most commonly been used in haptic devices (Massie and Salisbury, 1994) where drivetrain "transparency" (Carignan and Cleary, 2000) is a more important metric. However, recent results have shown that such actuators may effectively be used in legged robotics as well (Seok et al., 2012, 2013). While sacrifices must be made on actuator power output, low drivetrain friction enables effective end-effector force control using the motor current control techniques pioneered in the haptics community. This system was shown to be quite efficient, achieving a cost-of-transport of 0.51, a value significantly better than hydraulics-based robots and on par with land-based animals of similar mass (Seok et al., 2013).

1.2.1.5 Series Elastic

While series elastic actuation has been applied to other actuation technologies (Robinson, 2000; Pestana et al., 2010) it is most commonly found applied to high-gear ratio electromagnetic actuators. The effects of series compliance complements this technology in particular largely because of the inertia-decoupling effects of the compliance, protecting the actuator drivetrain from harm due to high impact loads. In this thesis, we therefore restrict our focus of series elastic actuation to electromagnetic technologies.

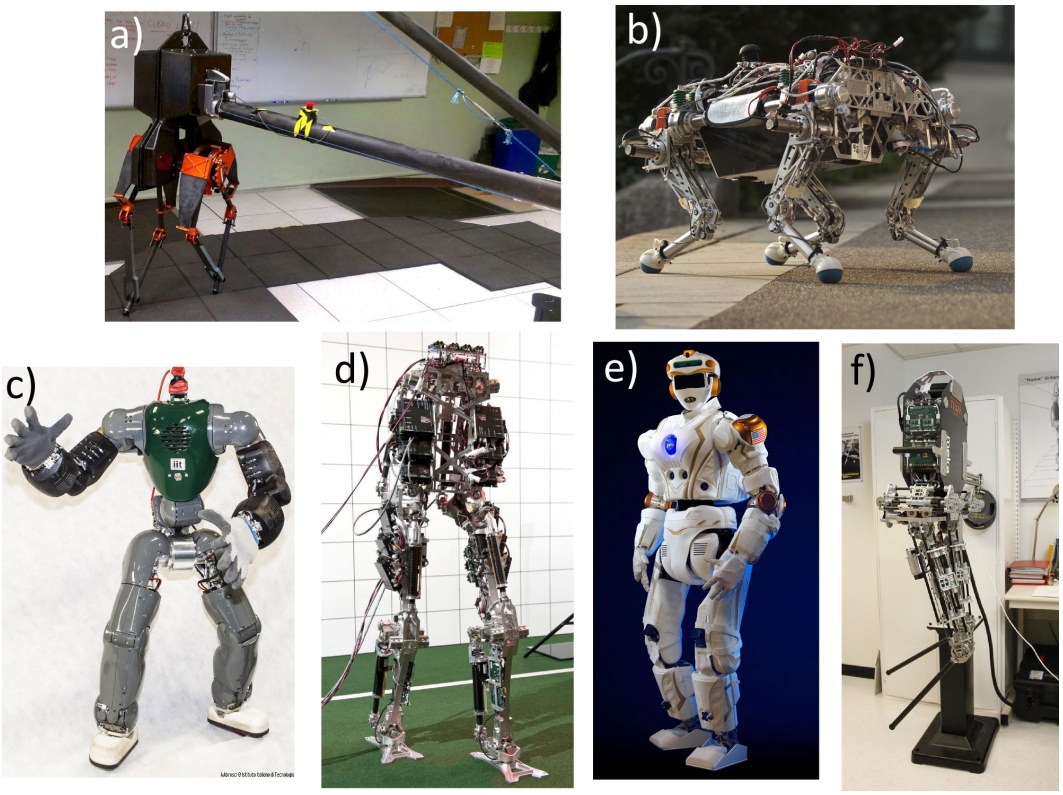

Figure 1.6: A few examples of high-performing robots with series elastic actuators: a) Atrias 2.1[9] (Hereid et al., 2014), b) StarlETH[10] (Hutter, 2013), c) Coman[11] (Li et al., 2012), d) Saffir[12] (Lee et al., 2013), e) Valkyrie[13] (NASA-JSC, 2013), and f) Hume[14] (Slovich et al., 2012).

A large number of systems have leveraged electromagnetic series elastic

[9] http://mime.oregonstate.edu/research/drl/robots/

[10] http://leggedrobotics.ethz.ch/

[11] http://www.iit.it/en/advr-labs/humanoids-a-human-centred-mechatronics/advr-humanoids-projects/compliant-humanoid-platform-coman.html

[12] http://www.romela.org/main/SAFFiR:_Shipboard_Autonomous_Fire-Fighting_Robot

[13] http://spectrum.ieee.org/automaton/robotics/military-robots/nasa-jsc-unveils-valkyrie-drc-robot

[14] http://www.me.utexas.edu/ hcrl/our_robots/index.html

actuation technology to achieve promising results (see Figure 1.6). For example, Atrias 2.1 (Hereid et al., 2014), StarlETH (Hutter, 2013) and Coman (Li et al., 2012) are all capable of walking over small obstacles and tolerating large exogenous perturbations. StarlETH is capable of several dynamic gaits, and can even be safely dropped from high distances, demonstrating the robustness of SEAs. In addition, this work has demonstrated the energy savings made possible by series compliance, where, in hopping experiments, the springs in the actuators were able to store and return 64% of the required energy per hop (Hutter, 2013). In a rigid system, all of this work would have to be performed actively, draining the system's power source more quickly.

While SEAs generally perform well in all of the metrics we have discussed, many of the actuators driving the systems in Figure 1.6 can be improved in the areas of power-to-weight, torque-to-weight, power-to-volume, force controllability, and efficiency, as will be discussed in detail in Chapter 2. Additionally, the single actuator performance of these systems is often not well characterized, making it difficult to directly compare these actuators for adoption in other applications. (Pestana et al., 2010) provides experimental data for peak power output for the actuator described in (Pratt and Krupp, 2004) and is able to achieve 64 W/kg. (Hutter et al., 2011a) provides the power exerted during a hop for one joint (close to 60 W), but because the actuator is integrated into a three degree-of-freedom leg, the actuator power-to-weight ratio is difficult to calculate. A central goal of this work, therefore, is to improve the state of the art and fully characterize the performance of a

single series elastic actuator.

1.2.2 Control

The control of actuators plays a significant role in maximizing their performance. For example, actuator power output is maximized when applying large torques at high velocities. Obtaining high velocities within a fixed range of motion requires short bursts of acceleration to and from rest. For actuators with a fixed range of motion, the boundary conditions placed on high power experiments necessitates the use of automatic control strategies to ensure the actuator operates within its permissible region. It is then the combined performance of the hardware design and the control design which determines usable power-to-weight ratio of the actuator.

The actuation control approaches used by the systems discussed in the previous subsection can largely be grouped into three different categories: position control, open-loop force control, and closed-loop force control. Position control is the predominate control method used in high-gear-ratio electromagnetic actuators (Kaneko et al., 2011; Schaft, 2013; Ito et al., 2014). Using this approach and given knowledge of the robot center of mass, inverse kinematics can be used to maintain the center of mass within the foot contact polygon to achieve poses which are statically stable. Similarly, the Zero Moment Point (ZMP) approach uses estimated or measured ground reaction forces to maintain the center-of-pressure (COP) of each foot within its supporting contact polygon (Vukobratović and Borovac, 2004).

Both hydraulics and low-gear-ratio electromagnetic actuators have successfully implemented open-loop force control strategies. In this method, the output force of the actuator is not directly sensed in terms of a force-feedback sensor such as a spring or a load cell. Instead, the force-generating element in the actuator is controlled. This can be in the form of fluid pressure in hydraulic robots (Boston Dynamics, 2013) or electric current in electromagnetic robots (Seok et al., 2012). The efficacy of this technique in generating accurate end-effector forces greatly depends on the "transparency" of the actuator drivetrain (Carignan and Cleary, 2000). Hydraulic actuators do not require a drivetrain, but when using this method the unsensed effects of seal friction can degrade force-tracking accuracy.

Finally, closed-loop force control directly senses output force with an explicit sensor. The applications of this technique are too numerous to list here, but have been widely adopted in most forms of actuation technology. In this genre of actuator control, one approach may improve over another based on the quality of the force sensor, the dynamics of the actuation control plant, and most importantly, the type of control methodology employed. We provide a detailed overview of these control methodologies in Chapter 4.

1.3 Approach and Contribution

In this thesis, we focus on improving the state of the art in series elastic actuator performance with the goal of providing new technologies for the designers of robots to overcome common actuation limitations. We adopt

a system-wide optimization approach, dealing with factors which influence performance at the levels of mechanical design, electrical system design, and control (refer back to Figure 1.1). The specific contributions of this work are listed here.

1. In the area of mechanical design, we introduce the University of Texas Series Elastic Actuator (UT-SEA). This design improves upon previous SEA designs in that it is more compact and lightweight for its power output. Because of these advantages, the UT-SEA design has been adopted by several research institutions throughout the world. Our analysis of the UT-SEA design is supplemented by a study of the theoretical limitations imposed by the dynamics of two common SEA design configurations.

2. In the area of electrical system design, we derive an equation that determines the optimal pulse-width modulation frequency to minimize loss in high-performance electromagnetic motor applications. This equation is important because it allows designers to improve the mechanical power output of electromagnetic motors, making them more competitive with their hydraulic counterparts in terms of power and torque-to-weight metrics. This is achieved without significantly reducing the excellent efficiency properties of electromagnetic motors.

3. We present several contributions in the area of control. First, we propose an improved control architecture for achieving accurate force tracking with series elastic actuators. Our disturbance observer (DOB) based

approach differs from the state of the art in that proportional-derivative (PD) feedback is used to shape the dynamic response of the actuator, eliminating the need for a feedforward filter term and improving the controller's phase margin.

4. Second, we present a new method for controlling the torque output of robots driven by series elastic actuators. This method leverages the previously mentioned DOB-based force controller to conceptually decouple actuator-level dynamics from robot-level dynamics. A condition, the Disturbance Observer Region of Convergence (DROC), is proposed which indicates when this conceptual decoupling is feasible, given properties of the system and the DOB-based controller.

5. Our third contribution in the area of control is an algorithm for optimally tuning proportional and derivative feedback gains of rigid actuators to produce a certain stability margin. Our approach is unique in that it simultaneously possesses 1) a model of real-world performance-limiting factors (i.e. filtering and delay), 2) the ability to meet performance and stability criteria, and 3) the simplicity of a single closed-form expression. This algorithm has the potential to greatly simplify and automate the tuning of feedback parameters in many existing robots.

6. Our final contributions lie in the area of real-world implementation. Using the UT-SEA and several SEAs from NASA's Valkyrie robot, we perform experiments and document how these actuators perform given

the performance metrics previously introduced. This data may be used to directly compare the performance of future SEAs with those presented here. In addition, the UT-SEA was able to achieve an empirically measured power-to-weight ratio of 423 W/kg at 54.9% total system efficiency, a 6.41x improvement in power-to-weight over previous documented work and a 2.1x improvement over the most power-dense human muscles.

1.4 Thesis Outline

This thesis is organized as follows. Chapter 2 analyzes how the mechanical design of SEAs may be leveraged to maximize performance, specifically in regards to power output, size, weight, and efficiency. Chapter 3 then studies methods by which the power output of electromagnetic motors may be increased without significantly reducing their efficiency. Shifting away from design, Chapter 4 focuses on how SEAs can be controlled to maximize their performance both in terms of power output and force controllability. It also demonstrates the high performance of the UT-SEA actuator through experimental tests. Chapter 5 studies the robustness of the controller presented in the previous chapter to model uncertainty by means of the Disturbance observer Region Of Convergence (DROC). An actuator from the Valkyrie robot is used to characterize this robustness experimentally. In Chapter 6, we apply the ideas in this thesis to NASA-JSC's Valkyrie Humanoid robot, both in terms of design and control. Finally, inspired by the few rigid actuators in Valkyrie, Chapter 7 studies how far the impedance properties of rigid actua-

tors may be increased while maintaining a certain stability margin. This work is intended as a general guide for control designers who may use the simple equations we derive to optimally tune the common rigid actuator.

Chapter 2

UT-SEA: Design of a High-Performance SEA

In this chapter we present the mechanical design of the University of Texas Series Elastic Actuator (UT-SEA), an SEA designed to optimize power-to-weight and power-to-volume ratios (see Figure 2.1). We consider the question of SEA component selection in detail and present a novel and compact arrangement of mechanical components that enables the use of a highly efficient ball screw drive mechanism. Beyond the standard benefits of SEAs, the design presented here, and those derived from it, are especially well suited for applications requiring aggressive behaviors with minimal actuation mass and minimal power consumption.

2.1 SEA Design Background

Electric SEAs contain a motor to generate mechanical power, a speed reduction to amplify motor torque, a compliant element to sense force, and a transmission mechanism to route mechanical power to the actuated joint.

This chapter contains material from the following publications: (Paine and Sentis, 2012; Paine, Oh, and Sentis, 2014b). Oh contributed to this work by aiding in the design and understanding of the DOB control structure. Sentis contributed to this work by providing insight into force control techniques.

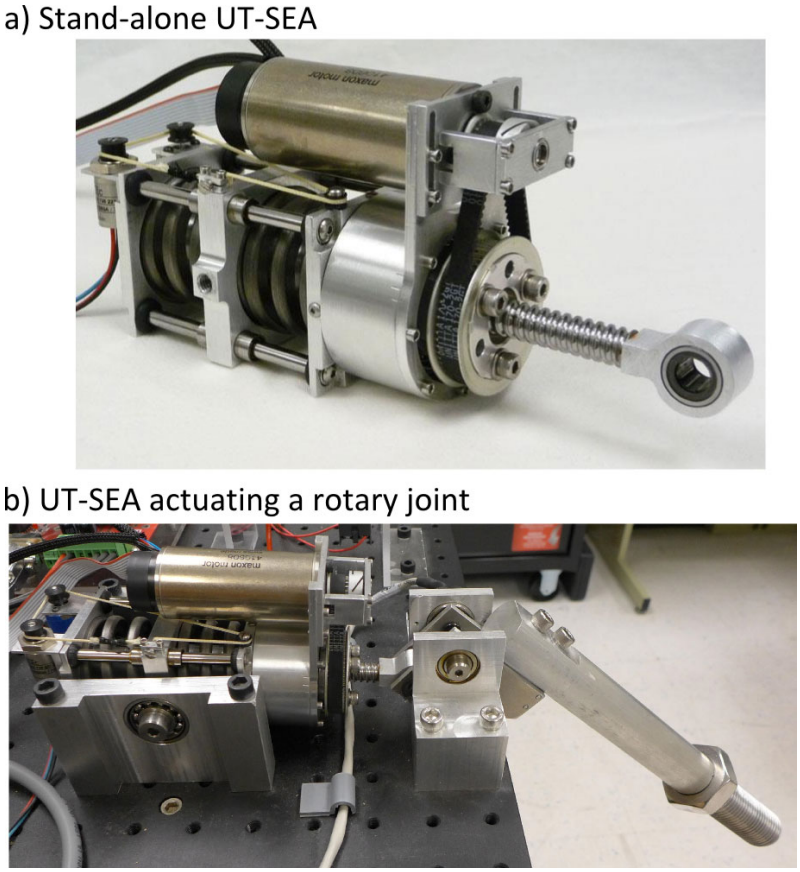

Figure 2.1: The University of Texas Series Elastic Actuator (UT-SEA). **a)** Standalone actuator. **b)** UT-SEA actuating a rotary joint.

These components can be chosen and configured in many different ways (Figure 2.2), producing designs with various trade-offs that affect power output, volumetric size, weight, efficiency, backdrivability, impact resistance, passive energy storage, backlash, and torque ripple. Existing SEA designs can be analyzed to identify such trade-offs based on their choice of speed reduction, compliant element, and transmission mechanism.

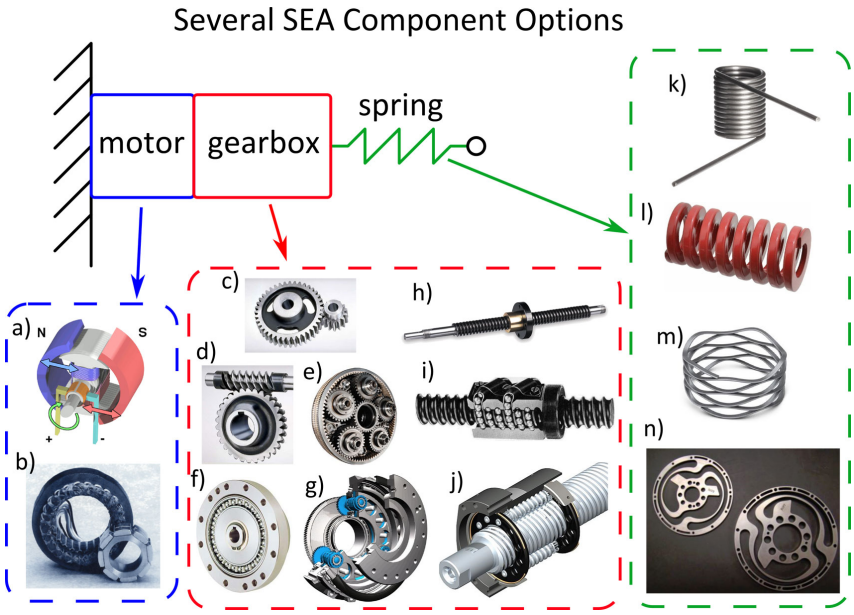

Figure 2.2: Designers of SEAs have to choose from many mechanical component options, each with advantages and disadvantages. The captions represent: a) brushed DC motor[1], b) brushless DC motor[2], c) spur gear[3], d) worm gear[4], e) planetary gear[5], f) harmonic drive[6], g) cyloid drive[7], h) lead screw[8], i) ball screw[9], j) roller screw[10], k) torsion spring[11], l) die spring[12], m) wave spring[13], n) planar spring (Diftler et al., 2011).

[1] http://upload.wikimedia.org/wikipedia/commons/0/04/Electric_motor_cycle_2.png

[2] http://www.mcfarlandassociates.net/mac_products/motors/images_motors/bps_series_brushless_direct_drive_frameless_motor_parts_set.gif

[3] http://static.ddmcdn.com/gif/gear-spur.jpg

[4] http://static.ddmcdn.com/gif/gear-worm.jpg

[5] http://www.bncgears.com/file/0_2011_06_24_194353.gif

[6] http://mfgnewsweb.com/archives/general_editorials/feb07/General-pix/REM-Harmonic-Drive.jpg

[7] http://www.nabtesco.de/en/products/principle-of-operation/

[8] http://cdn2.hubspot.net/hub/214143/file-70740309-png/images/lead-screw-for-security-automation.png?t=1394725170000

[9] http://www-mdp.eng.cam.ac.uk/web/library/enginfo/textbooks_dvd_only/DAN/threads/mechanics/recircBallBIG.jpg

[10] http://www.designworldonline.com/uploads/Imagegallery/roller-screws-2.jpg

Many researchers have proposed rotary SEA designs based primarily on commercially available off-the-shelf components, using a planetary gearbox for speed reduction, rotary or compression springs as the compliant element, and power transmission through a bevel gear or chain/cable (Curran and Orin, 2008; Hutter et al., 2009; Kong et al., 2009; Ragonesi et al., 2011; Grun et al., 2012; Pott et al., 2013; dos Santos et al., 2013; Hasankola et al., 2013). Use of off-the-shelf parts make these designs relatively low cost and easy to implement. However, multi-stage planetary gearboxes have poor efficiency (60-70% for 3-stage) and can be difficult to backdrive. Additionally, gear teeth introduce torque ripple and backlash and have limited torque capability for their weight. Bevel gears compound these effects by adding additional backlash and friction loss.

A compact rotary SEA design can be achieved using a harmonic drive and a high-stiffness planar spring (Lagoda et al., 2010; Diftler et al., 2011; Sergi et al., 2012; Parmiggiani et al., 2012; Paine et al., 2014a). Harmonic drives benefit from having little backlash and being small in size, but suffer from poor efficiency (25-80% depending on ratio, speed, and lubricant), poor backdrivability, torque ripple, and are more expensive than planetary/spur gears. The high-stiffness planar springs of (Lagoda et al., 2010) and (Diftler et al., 2011) deflect only a small amount under load and therefore store less energy than designs with softer springs. Conversely, stiff springs increase an

[11] `http://ecx.images-amazon.com/images/I/71IvLsPYqyL.jpg`
[12] `http://cdn.mscdirect.com/global/images/ProductImages/7601544-23.jpg`
[13] `http://www.smalley.com/images/renders/spring_wave.jpg`

actuator's open loop bandwidth, which may be desirable for applications with high force bandwidth requirements.

Another effective way to achieve a compact design without using expensive custom torsion springs is by using linear springs coupled to rotary shafts and placing the springs between the motor and chassis ground (Torres-Jara and Banks, 2004; Hutter et al., 2011b; Hutter, 2013). These designs are also able to reduce the spring stiffness in comparison the torsion spring method.

Other designs place the spring within the reduction phase (Kong et al., 2012; Taylor, 2011). This arrangement reduces the torque requirement on the spring compared to designs with the spring at the output. A spring's wire gauge is directly correlated to the amount of torque it can safely support, which allows the spring in these designs to be smaller. However, because the torque compressing the spring is reduced, the energy stored in the spring is reduced as well. (Kong et al., 2012) uses a novel worm-gear/rotary-spring/spur-gear design which allows an orthogonal placement of the motor relative to the joint axis at the cost of reduced efficiency and non-backdrivability due to the worm gear. (Taylor, 2011) uses two motors in parallel and has a relatively small reduction through a series of gears and a cable transmission. This approach achieves the same output torque but doubles the output speed compared to designs with a single motor. The speed increase is due to the need for only half of the speed reduction ratio in the two-motor case, given a fixed output torque requirement. This is achieved at the cost of increased weight and complexity and is difficult to implement with brushless direct current (BLDC) motors due

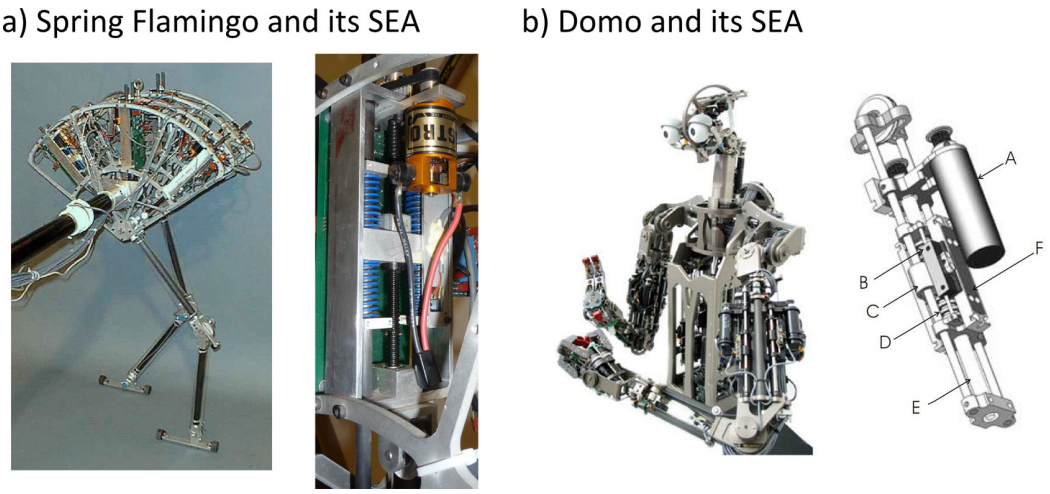

Figure 2.3: Two ball screw SEA designs from the Massachusetts Institute of Technology. a) The Spring Flamingo SEA and b) The Domo Humanoid SEA (Edsinger-Gonzales and Weber, 2004).

to commutation synchronization issues.

Prismatic-to-rotary SEA designs are possible through the use of ball screws as the primary reduction mechanism followed by a cable drive to remotely drive a revolute joint (Edsinger-Gonzales and Weber, 2004; Gregorio et al., 1997; Pratt and Pratt, 1998; Slovich et al., 2012). Ball screws are highly efficient, even for large speed reductions (85-90%), are backdrivable, are tolerant to impact loads, and do not introduce significant torque ripple. (Edsinger-Gonzales and Weber, 2004), shown in Figure 2.3, includes a belt drive between the motor and the ball screw which enables an additional speed reduction due to the pulley diameter ratio. (Pratt and Krupp, 2004; Slovich et al., 2012), shown in Figure 2.4, uses a ball screw speed reduction and removes the need for a cable transmission by directly driving the output joint

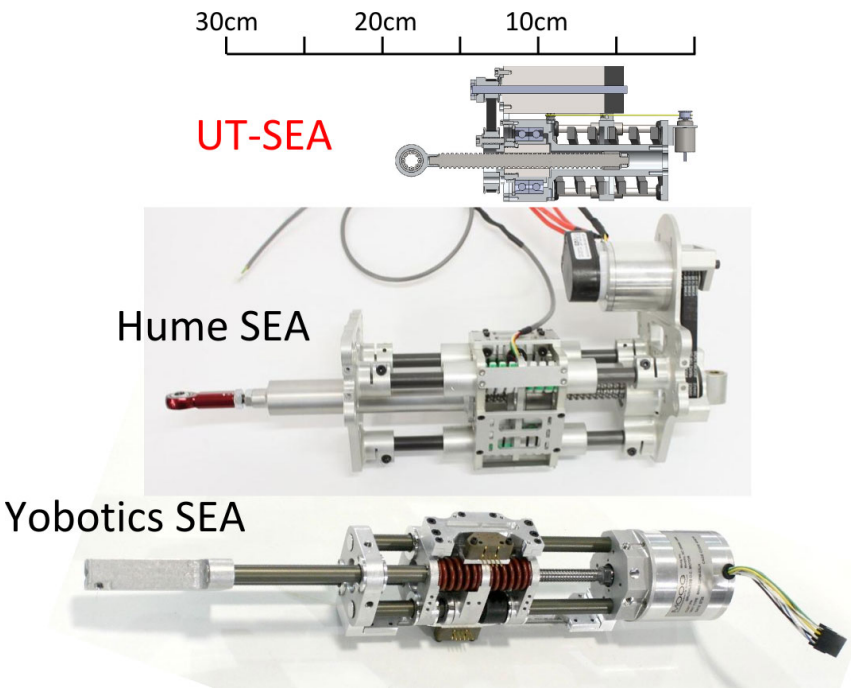

Figure 2.4: To-scale comparison of several prismatic SEAs with similar performance profiles. The UT-SEA reduces the space required to implement a high-performance SEA.

with a pushrod mechanism.

While several effective prismatic SEA designs exist, they share one notable drawback. Designs based on ball screws have not achieved the levels of compactness seen in rotary actuator design and therefore are difficult to integrate into the dense mechanical designs often found in highly complex robots (see Figure 2.5). This is an unfortunate characteristic because ball-screw-based actuators otherwise have excellent properties. The UT-SEA brings the high-performance properties of ball screw SEAs to complex, high-performance robots by shrinking the actuator package compared to previous designs as

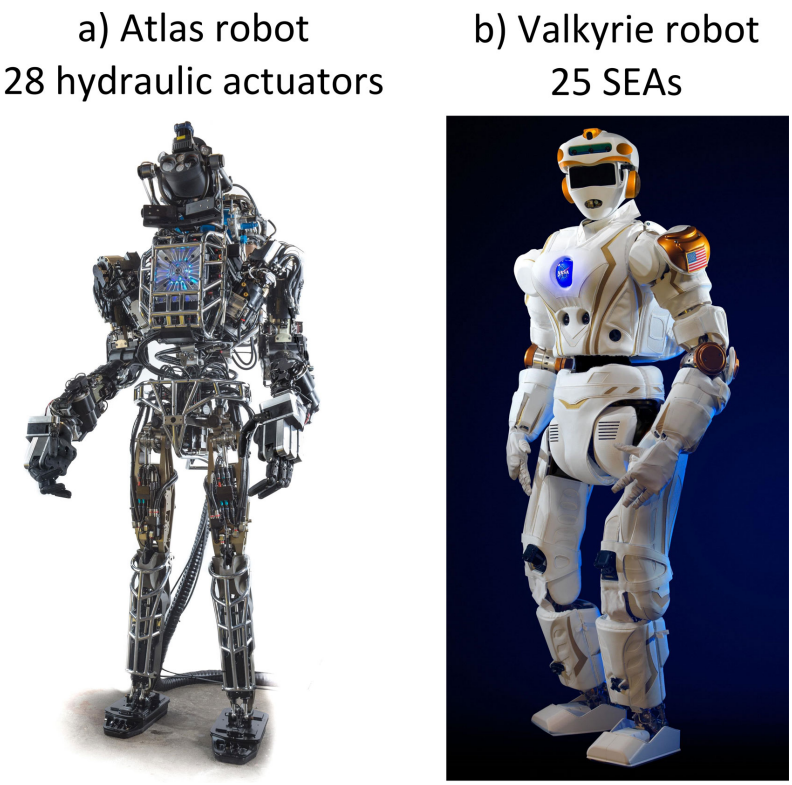

Figure 2.5: Two examples of complex high-performance robots. a) Boston Dynamic's Atlas robot has 28 hydraulic actuated degrees of freedom. b) NASA-JSC's Valkyrie Robot has 44 electric actuated degrees of freedom. To fit this many actuators in such a small space, a compact actuator design is essential.

shown in Figure 2.4. In the next section, we show how this compactness is achieved while at the same time maximizing actuator speed and torque.

2.2 Design of the UT-SEA

Nature provides many examples of well-designed actuators. An average human adult male can produce 1500 watts of mechanical power during pedaling exercises, which corresponds to a whole-body power-to-weight ratio of 19.5

watts per kilogram (Beelen and Sargeant, 1991). To achieve similar performance in man-made machines, great care must be taken during the actuator design phase to maximize mechanical power output while keeping actuator size and weight small. Excess actuator weight reduces the whole-body power-to-weight ratio while large size limits a modular actuator's applicability in dense high-degree-of-freedom robot designs. Hydraulic actuation is one approach which achieves these goals but suffers from inefficient operation as discussed in (Zoss et al., 2006).

The UT-SEA design process began with a set of loose performance specifications for a robotic knee actuator (peak joint torque of 70 Nm and maximum velocity of 15 rad/sec). These specifications were obtained from simulations of legged locomotion of a 15kg robot in rough terrain (Sentis and Fernandez, 2011). However, it is important to note that the design goal of this particular actuator was not for use in any specific robot, only for experimental study of mechanical power output.

2.2.1 Motor Interfacing and Speed Reduction Sizing

The Maxon EC-4pole 30 motor used in the UT-SEA is rated to operate at 48 volts, but is driven with 80 volts in our design. Increasing the bus voltage used to drive a motor increases the motor's maximum permissible speed. The high motor speed produced by high bus voltage enables the use of a large speed reduction which increases both intermittent and continuous torque capability compared to designs with lower voltages and lower speed reductions (assuming

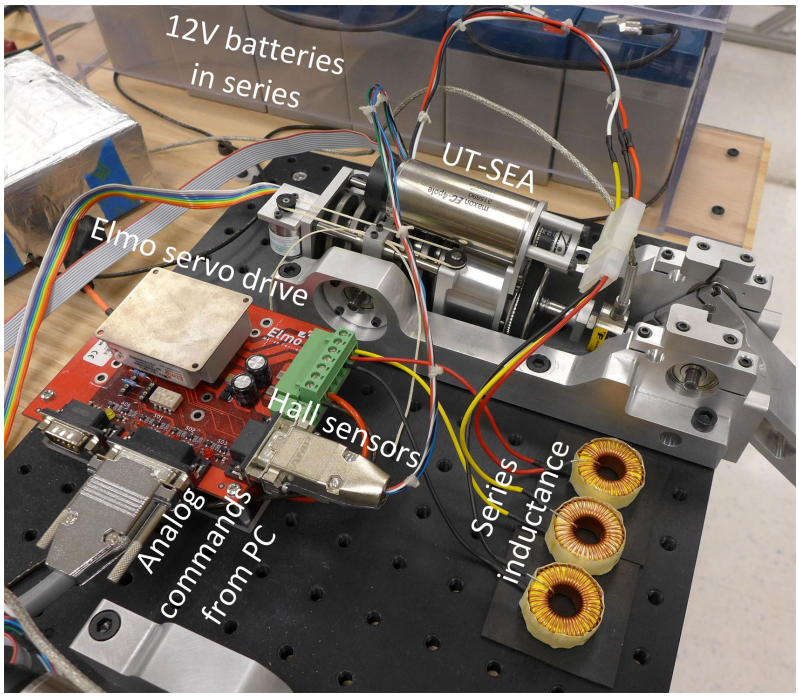

Figure 2.6: Motor interface of the UT-SEA.

a fixed actuator output speed requirement). In the case of the UT-SEA, we were able to use a speed reduction of approximately 175:1 while still achieving our target output speed of 15 rad/sec. Transient motor current is regulated by using a 32kHz PWM servo drive (Elmo Ocarina 15/100) and by adding high-current inductors in series with the motor (see Figure 2.6). Calculations provided by Elmo indicated that a series inductance of 0.082mH would keep transient current within reasonable values (this point is further discussed in Chapter 3). The small added mass of the inductors (115g, 11% of the actuator mass) is justified in that they allow the continuous force of the actuator to be increased by 66% without sacrificing output speed.

Several factors limit the degree to which motor bus voltage should be increased:

1. The motor must be able to mechanically withstand the speeds induced by the high voltage. The bearing assemblies used in motors can only tolerate certain speeds before their service life is reduced. Additionally, the large centrifugal forces present at high speed must be below the strength limits of laminates and bonding agents used to assemble the motor rotor. Maximum permissible motor speeds are typically provided by the motor manufacturer.

2. If a drivetrain is used, its maximum speed must be considered as well. Both gear assemblies and screw drives each have maximum input speeds which should not be exceeded.

3. Increasing bus voltage increases current rippled in pulse-width-modulated (PWM) drives which may lead to decreased motor efficiency and increased heating, a point discussed in detail in Chapter 3.

4. Large speed reductions are generally less efficient than small speed reductions, even if only by a small amount (different ball screw pitches, for example). This factor, combined with the inertia amplifying effects of speed reductions can contribute to bring the passive output impedance of an actuator to unacceptably high levels. For the UT-SEA, this point is largely avoided by using a series compliant element, which decouples the actuator inertia from ground, creating a soft passive output impedance.

5. If an actuator will be operated in close proximity to humans, care must be taken to protect humans from large voltages. High voltage can induce significant currents in the human body which can have life-threatening consequences.

2.2.2 Drivetrain

To preserve mechanical energy produced by the motor, energy must be transmitted from the motor to the joint with as few losses as possible. We chose a pulley/ball-screw speed reduction design similar to (Edsinger-Gonzales and Weber, 2004) for several reasons. A pulley/ball-screw reduction is efficient (typically above 90%), impact resistant, and backdrivable while the pulley ratio reduces the high motor speed to a speed more suitable for driving the ball screw.

Unlike (Edsinger-Gonzales and Weber, 2004) and other ball screw SEA designs, our design drives the ball nut instead of the ball screw ((Garrec, 2010) uses a similar ball-nut-driven design but is a non-series-elastic cable-driven actuator). Driving the ball nut enables two key features which reduce the size and weight of the UT-SEA. First, ball screw support is incorporated directly into the actuator housing using an innovative piston-style guide (see Figure 2.7). This feature replaces the long, bulky rails used to support the output carriage in conventional prismatic SEA designs. Secondly, the compliant element is placed concentrically around the piston-style ball screw support which gives series elasticity without adding to the length of the actuator. These two fea-

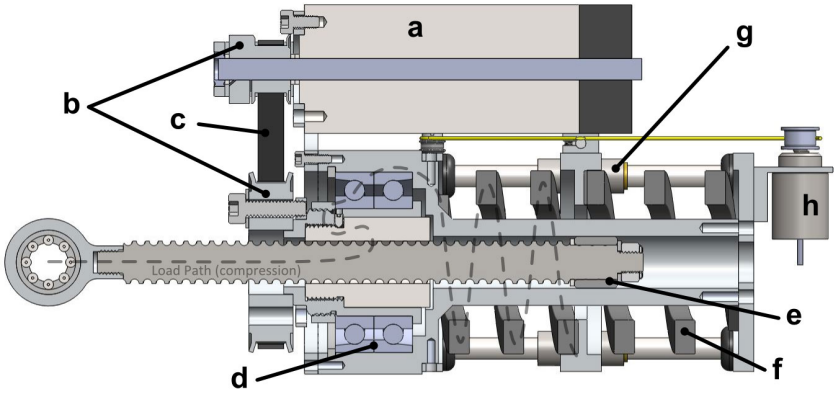

Figure 2.7: Cross section of the UT-SEA, showing drivetrain components including: a) Maxon EC-4pole 30 200W BLDC motor, b) 3:1 pulley speed reduction, c) low backlash timing belt, d) angular contact bearings, e) piston-style ball screw support, f) high compliance springs, g) miniature ball bearing guides, and h) absolute encoder. The compression load path is depicted as well.

tures combine together to define the compact form factor of the UT-SEA.

The ball nut is supported by dual angular contact bearings which allow the ball nut to rotate within the housing while transmitting axial force from the ball nut to the housing. Custom preloaded die springs (manufactured by Diamond Wire Spring Co.) transmit force from the actuator housing to the chassis ground. The die springs are supported by four miniature ball bearing guide rails (Misumi) which are mounted to the housing using grommets that allow for slight misalignment during operation. The miniature ball bearing guides offer both lower friction and higher tolerance to torsional loads than bushing style guides. Force is sensed using a 20,000 count-per-revolution incremental encoder (Avago AEDA 3300) along with an absolute sensor (Novotechnik Vert-X 1302) to remove the need for startup calibration. A low stretch, low

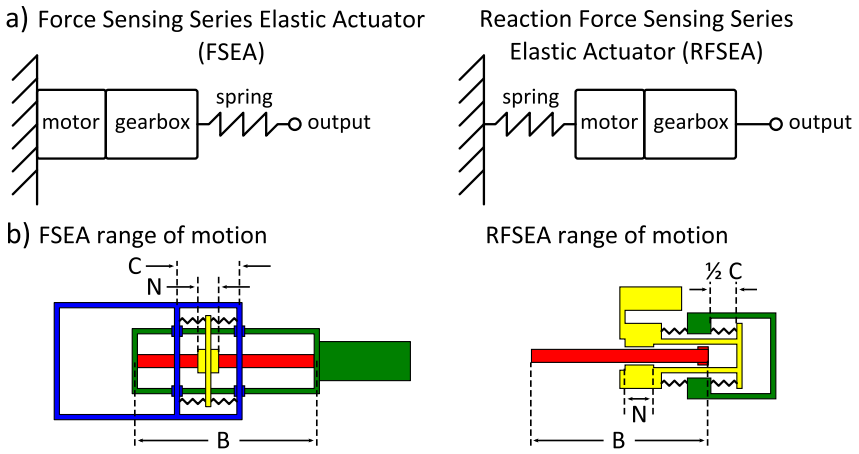

Figure 2.8: Two SEA topologies and their respective range of motion. For simplification, we assume that springs are fully compressible, spring plates have zero thickness, and the FSEA carriage travel is constrained to the length of the ball screw. The notations represent B: ball screw length, C: carriage length, N: ball nut length. Range of motion is then $B - C$ for the FSEA and $B + C - N$ for the RFSEA.

creep Vectran cable is attached to the chassis ground and is routed around the two spring deflection sensors using pulleys and an idler. Overall actuator position is measured by combining readings from the motor encoder and the spring encoder. An absolute rotary sensor on the driven joint is used to initialize actuator position.

2.2.3 Spring Placement and Stiffness

There are two common arrangements of components found in SEA designs (see Figure 2.8). The first arrangement, which we will refer to in this thesis as Force Sensing Series Elastic Actuator (FSEA), places the compliant element between the gearbox output and the load. The second arrangement,

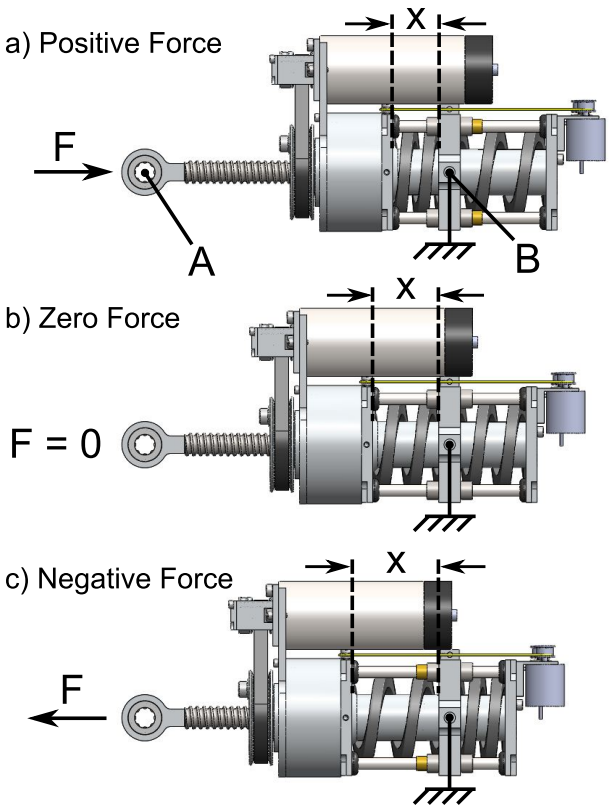

Figure 2.9: UT-SEA operation for a fixed-displacement variable-force scenario. Actuator displacement is defined as the distance between points A and B. This distance remains constant while spring deflection (x) depends on actuator force.

which we will refer to as Reaction Force Sensing Series Elastic Actuator (RF-SEA), places the spring between the motor housing and the chassis ground.

From a design standpoint there are several trade-offs between the two arrangements. RFSEA style actuators have the advantage of being more compact since the compliant element does not have to travel with the load but may be placed statically behind the actuator (Figure 2.9) or it can be remotely

located as shown in (Hutter et al., 2011b; Torres-Jara and Banks, 2004). Prismatic RFSEAs also have greater range of motion for a given ball screw travel length compared to prismatic FSEAs as shown in Figure 2.8. The primary drawbacks of RFSEAs are less direct force sensing, reduced force tracking performance, and decreased protection from impact loads as described in Section 4.2 of Chapter 4. An RFSEA style design was chosen to minimize the bounding volume of the UT-SEA. However, this design decision was heavily influenced by the selection of the pushrod/ball-screw drivetrain. The drivetrain exhibited strong radial symmetry and possessed long, narrow ball screw support structure which allowed die springs to be integrated without excess bulk. Additionally, the decreased protection from impact loads of an RFSEA design is somewhat, though not completely, mitigated by the impact-tolerance of the ball screw drivetrain in the UT-SEA.

Spring stiffness for UT-SEA was chosen to maximize energy storage. For a given force, soft springs are able to store more energy than stiff springs. Peak force, desired deflection (maximum possible deflection to minimize stiffness), and the geometric constraints of the actuator were given as design specifications to Diamond Wire Spring Co. They then designed and manufactured a spring with a stiffness rate of 138 N/mm which effectively doubles to 277 N/mm for the actuator spring constant because two springs are used with precompression.

Table 2.1: UT-SEA specifications

UT-SEA Design Specifications		
Weight	1013 g	2.23 lbs
Stroke	6 cm	2.36 in
Max Speed	32.5 cm/sec	12.79 in/sec
Continuous Force	848 N	190 lbs
Intermittent Force	2800 N	629 lbs
Spring Stiffness	278 N/mm	1587 lbs/in
Force Resolution	0.31 N	0.069 lbs
Operating Voltage	80V	

2.2.4 Summary

The end result of the design process is a pushrod RFSEA-style actuator that is compact and modular enough to be integrated into dense mechanical designs. Rotary joint designs using linear actuators can benefit from the non-linear linkage kinematics created at the joint at the cost of a fixed range of motion. Torque generated by such a linkage has an angle-dependent moment arm which can be used to provide high torque and high speed capability where they are needed. A summary of the design parameters for the actuator can be seen in Table 2.1.

The specifications listed in Table 2.1 represent the mechanical performance capability of the UT-SEA. Referring back to Figure 1.1, these specifications represent the mechanical layer of performance capability. While they do

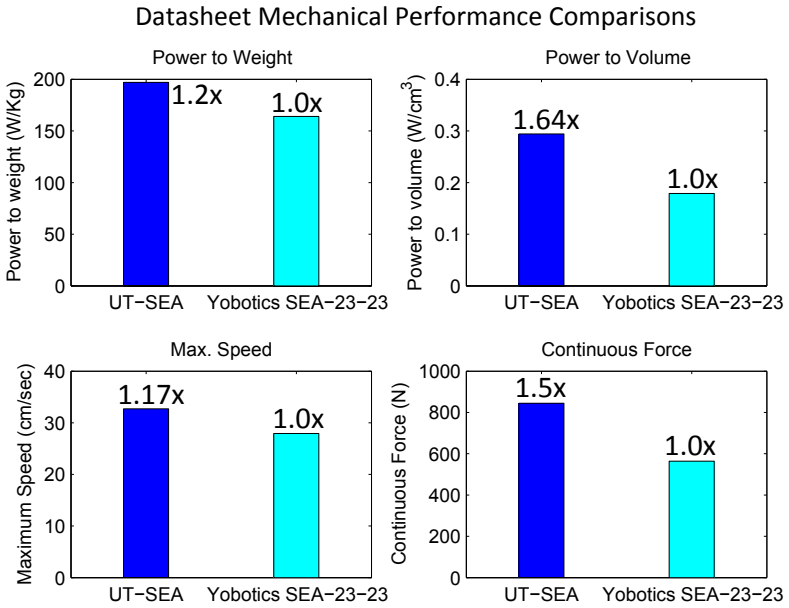

Figure 2.10: Comparison of the UT-SEA mechanical performance versus that of the Yobotics SEA-23-23 (Pratt and Krupp, 2004). This comparison only includes datasheet values (motor rated continuous power, for example) and therefore does not take into consideration control or other factors such as non-standardized meanings of "continuous power" between motor manufacturers. As such, these results may serve as a rough indication of mechanical performance. An accurate system-level performance characterization is presented in Chapter 4.

not represent the performance capability of the full robotic system including the embedded system and control layers, they define the maximum achievable performance of an idealized system. If we compare the mechanical performance capability of the UT-SEA against other prismatic SEAs, the Yobotics SEA-23-23 (Pratt and Krupp, 2004) for instance, we obtain a rough measure of the performance improvements achieved by the mechanical design aspect of UT-SEA over previous designs (see Figure 2.10). This comparison alone does

not guarantee better performance realization at the system level, but suggests an improvement in the performance foundation we are able to build upon. Maintaining this performance through the embedded system and control performance layers is the focus of much of the remainder of this thesis and will be explored in later chapters.

2.3 Design Iteration for Modularity

Soon after completing experiments with the initial UT-SEA design, we set our sights on improving the design. The initial design was difficult to assemble, required several complex high-tolerance parts, and had over-sized angular contact bearings. In the next iteration of the UT-SEA design (Figure 2.11) we addressed these issues.

In addition, an interest arose to study both rigid and compliant actuators. Any SEA may be converted into a rigid actuator if the springs are removed and replaced with a rigid material, but this approach results in an overly complex and heavy actuator. On top of the several design improvements previously listed, we wanted to create a design which was intended to act either as a compliant or as a rigid actuator.

To achieve this goal, we created an actuator mounting interface which could accept either a rigid mount or a compliant mount. Using this approach, a single part of the actuator could be exchanged to convert a compliant actuator to a rigid actuator, or vice versa.

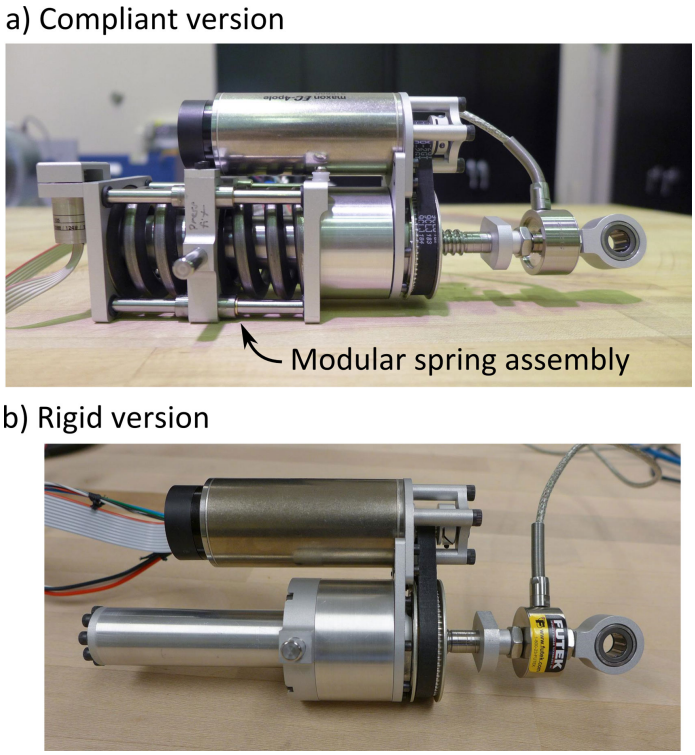

Figure 2.11: A design iteration of the UT-SEA for spring assembly modularity. In a) the UT-SEA is configured to use a compliant connection to ground while in b) a rigid connection is used. This design also includes a load cell which may be used for force feedback in the rigid actuator, and as a redundant force feedback in the compliant actuator.

2.4 Derived SEA Designs

The performance and size benefits of the UT-SEA mechanical design have found use in several external projects. This section presents a few of these applications.

The Mina Exoskeleton[1] is a project by the Florida Institute for Human

[1]http://www.ihmc.us/research/mina/Mina.php

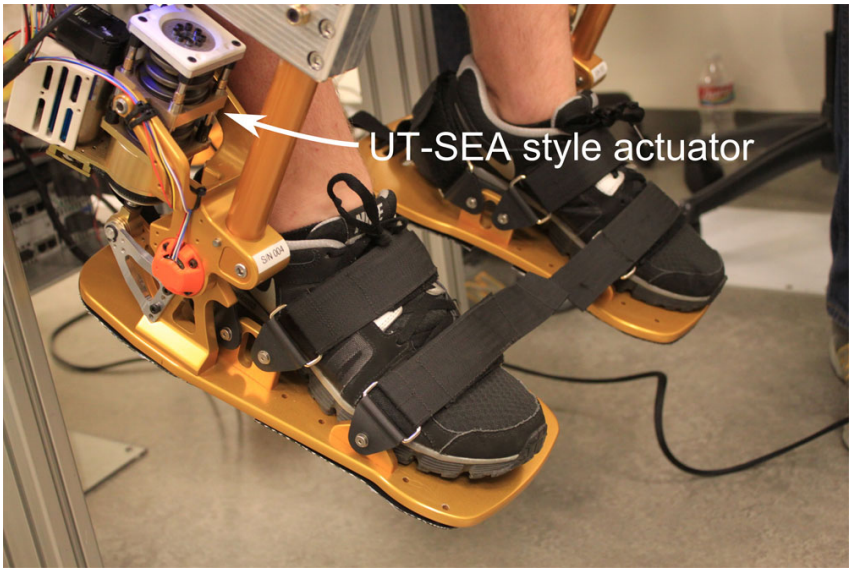

Figure 2.12: An SEA built by IHMC and NASA-JSC and used on the ankle of an exoskeleton currently under development at NASA-JSC.

and Machine Cognition (IHMC) focused on enabling persons with paraplegia to walk. Building on this work, IHMC has developed an ankle actuator based on the UT-SEA design for use in research on powered ankle orthotics (see Figure 2.12). The design of the modified IHMC actuator was then used by a collaborative effort between IHMC and NASA-JSC to design, build, and test ankle exoskeletons for exercise and dynamometry.

Another project led by NASA-JSC is the Valkyrie robot (shown in Figure 2.5). Valkyrie is an advanced humanoid robot designed to respond to disaster scenarios and to advance human spaceflight by one day assisting human explorers in extraterrestrial settings. Valkyrie leverages the compact UT-SEA topology in six joints, located in the robot's torso and at its ankles.

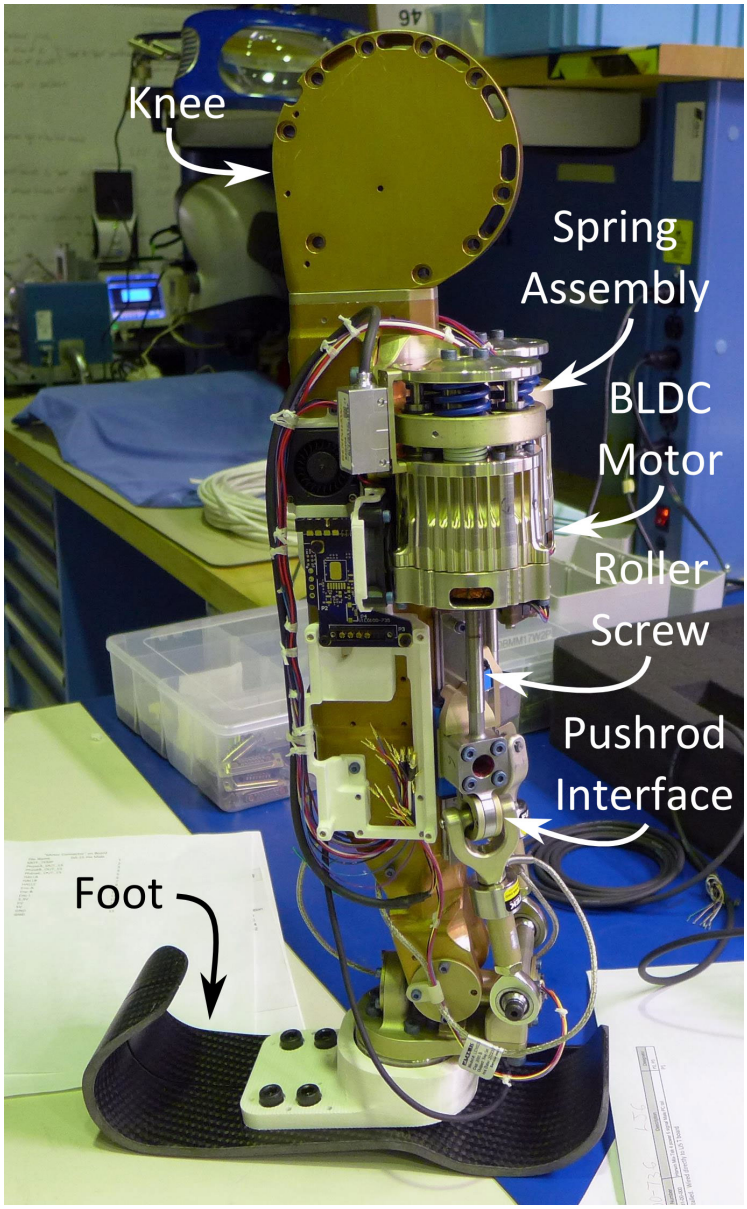

Figure 2.13: Valkyrie ankle with UT-SEA inspired actuator design.

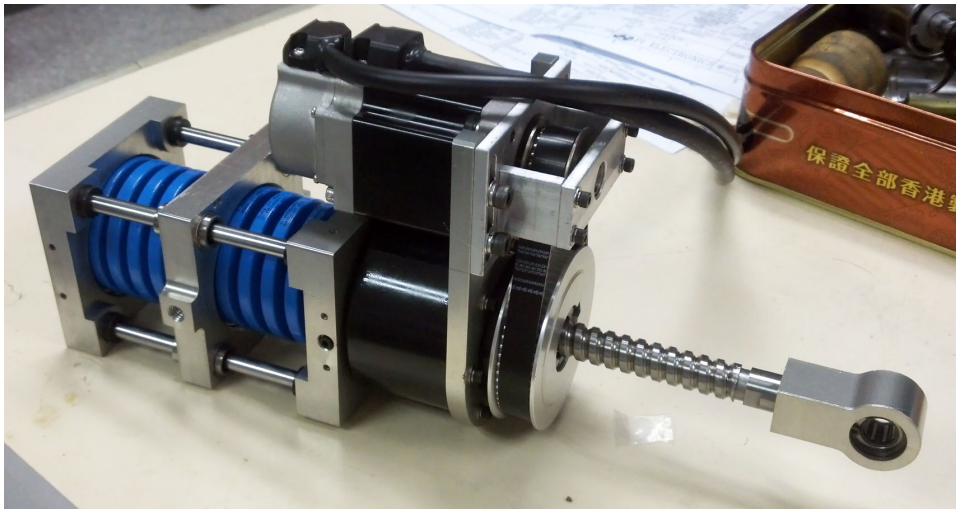

Figure 2.14: Stone Age Robotics SEA based on UT-SEA design.

However, the Valkyrie linear actuator design (Figure 2.13) differs in several ways: 1) it does not use a belt to reduce motor speed before the screw phase, 2) it uses three springs instead of one as the concentric compliant element, 3) it uses an optical spring deflection sensor instead of a string-type encoder, 4) it uses a roller screw instead of a ball screw, and 5) the pushrod output is constrained by a linear guide and transmits force to a two-dimensional linkage.

Finally, the Human Centered Robotics Lab (HCRL) has worked with a Chinese robotics company, Stone Age Robotics, to develop a low-cost industrial-grade version of the UT-SEA (see Figure 2.14). This actuator is intended to serve markets requiring a high-power actuator with high-fidelity force control capability. Such characteristics are often found in human-centric work, where robots work along side humans to achieve a common goal.

Chapter 3

Maximizing Performance of Electric Motors

In the previous chapter, we focused on methods to minimize lost performance potential (refer back to Figure 1.1) in the mechanical design of an SEA. From a performance metric standpoint, we aimed to maximize power-to-weight and power-to-volume ratios by maximizing mechanical efficiency (i.e. increasing usable mechanical power) and by minimizing both weight and volume. In this chapter, we instead focus on improving these metrics by increasing the mechanical power output of electromagnetic motors. We do this not by proposing new motor designs, but by analyzing the transfer of energy from electric potential energy through a power modulation device and to the mechanical energy output of an electromagnetic motor (see Figure 3.1). The end results of this analysis are methods and equations that may be used to optimally select the various parameters (i.e. bus voltage, PWM frequency) needed to interface an electromagnetic motor to an energy source.

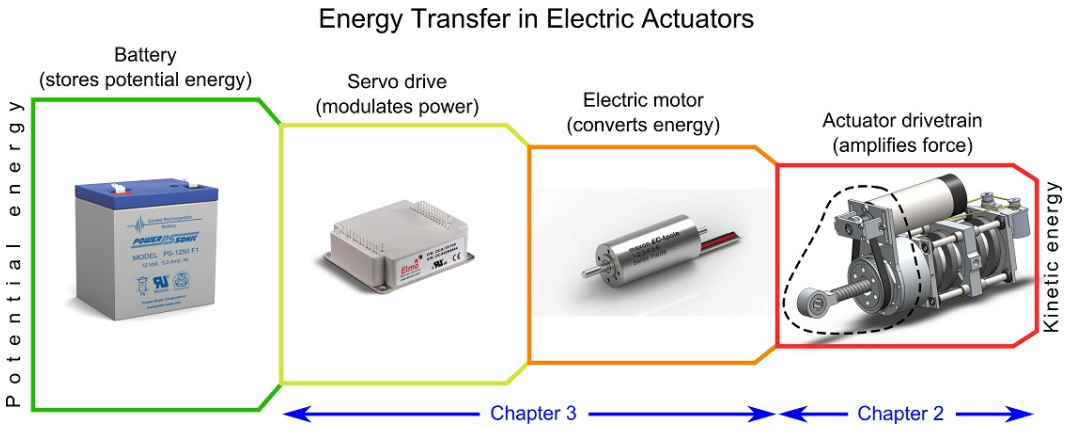

Figure 3.1: Energy transfer in electric actuators depicting loss at each stage (not to scale). A battery[14] transfers energy to a servo drive[15] which modulates the energy supplied to an electric motor[16]. The actuator drivetrain amplifies motor torque to operate the motor at high speeds where it is most efficient at generating mechanical power. Chapater 2 analyzed aspects of the actuator mechanical design while this chapter (Chapter 3) considers the combined effects of the servo drive and electric motor.

3.1 Motors as Energy Transducers

In an electric actuator, the motor is the source of mechanical energy. Therefore its selection, integration, and control significantly influences overall system performance. A DC motor may simply be thought of as an energy transducer. That is, they convert electrical voltage and current into mechanical velocity and torque. The greater the applied electrical energy, the greater the mechanical energy output becomes.

[14] http://www.power-sonic.com/images/powersonic/sla_batteries/ps_psg_series/12volt/PS-1250_11_Feb_21.pdf

[15] http://www.elmomc.com/products/ocarina-miniature-analog-amplifier-main.htm

[16] http://www.maxonmotorusa.com/maxon/view/catalog/

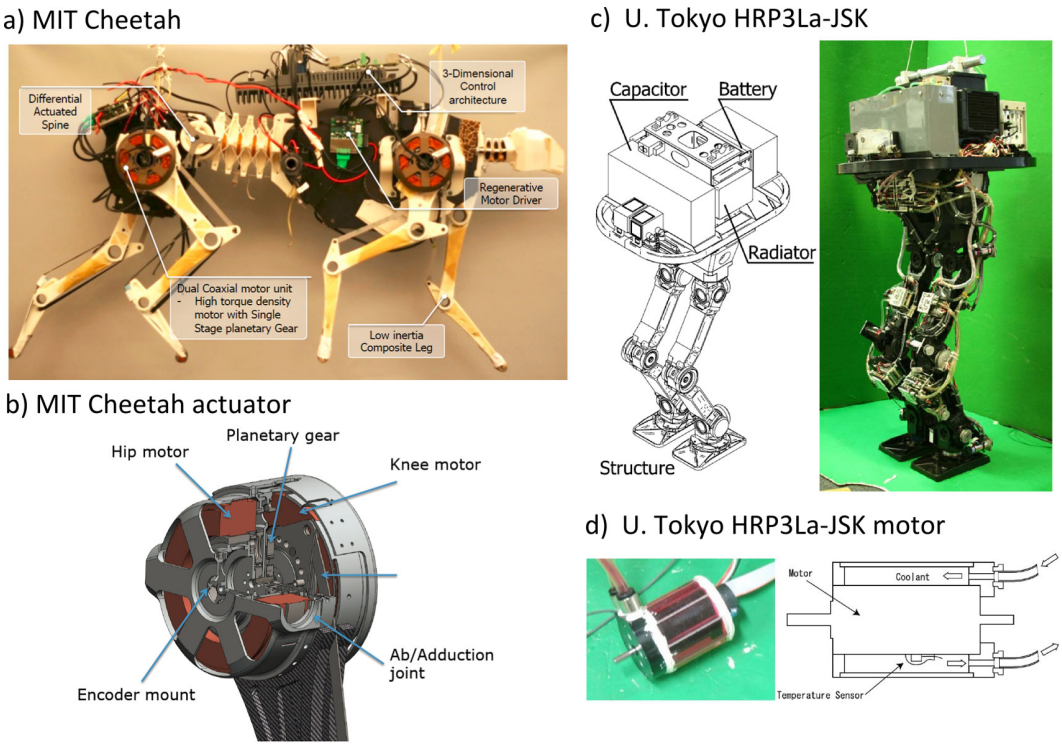

Figure 3.2: Two recent high power electric robots. a) MIT Cheetah (Seok et al., 2013), b) MIT Cheetah high torque density actuator (shoulder assembly) (Seok et al., 2012), c) University of Tokyo HRP3La-JSK (Urata et al., 2010), and d) University of Tokyo HRP3La-JSK liquid cooled motor (Urata et al., 2008, 2010).

Recently, several impressive electrically actuated robots (shown in Figure 3.2) have leveraged this idea to increase motor output power far beyond manufacturer-rated values without damaging the motor (Urata et al., 2008; Ito et al., 2010; Urata et al., 2010; Seok et al., 2012, 2013). Their approaches can be separated into two categories. The first approach increases mechanical energy by commanding motor currents beyond the manufacturer's rated value. A comprehensive description of this idea may be found in (Urata et al.,

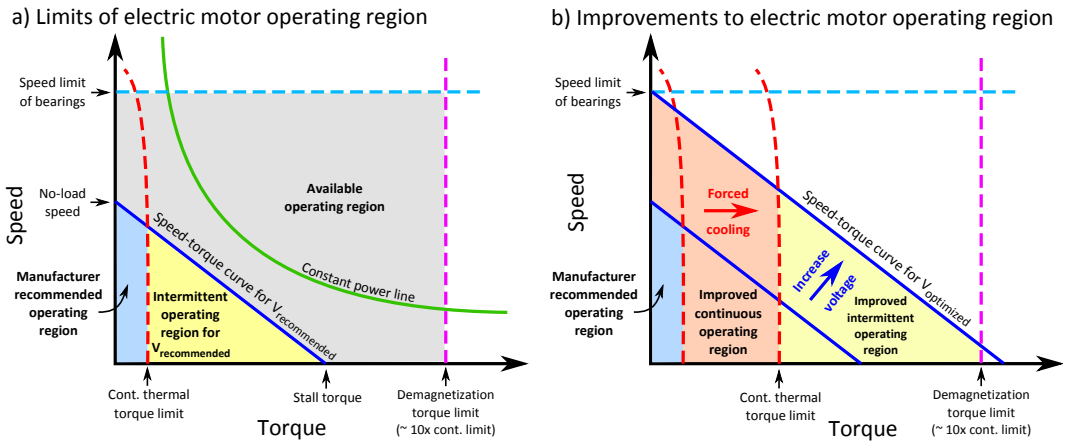

Figure 3.3: A comparison between conventional motor operation techniques (a) and techniques optimized for power output (b). Continuous torque may be increased using forced cooling techniques (Urata et al., 2008) while maximum speed may be increased by raising motor bus voltage.

2008). In brief, for short periods of time, large currents are permissible due to the relatively slow thermal time constant of the motor's core compared to the motor's electrical and mechanical time constants. By modeling internal thermal capacitance and resistance and estimating motor core temperature, large current magnitudes are permissible and can be used without damaging the motor. This model-based approach contrasts naive sensor-based methods which only measure motor surface temperature, and therefore cannot predict, but can only react to, motor overheating. The efficacy of this approach is demonstrated in (Urata et al., 2008) where peak motor output power of 7.5 times that of the continuous rated value is achieved without noticeable damage to the motor.

For longer time scales, forced liquid cooling enables the safe usage of

large currents in electric motors (Aghili et al., 2007; Urata et al., 2008; Ito et al., 2010; Urata et al., 2010). As shown in Figure 3.3, forced cooling increases the continuous torque limit of a motor by rapidly removing heat, reducing motor core temperature for a given applied current compared to passive or forced air cooling. In (Urata et al., 2008), this approach increases the continuous torque limit by a factor of four over passive air cooling.

The second approach mentioned in prior work for increasing motor power output is to raise motor bus voltage (or maximum motor voltage, V_{max}) beyond the manufacturer's rated value (refer again to Figure 3.3). However, this approach is less thoroughly discussed than the aforementioned method of increasing motor current. (Urata et al., 2008) appears to be using 80 volts to drive a 48 volt motor, but does not explicitly discuss why this is done or what effect increased voltage has on the system design process. (Seok et al., 2012) acknowledges increasing voltage as a means of achieving large motor output power, but chooses not to do so because of a desire to use a low-ratio speed reduction to maintain a "transparent" actuator drivetrain.

In our work, we choose to increase motor output power by using large motor bus voltage (as shown in Figure 3.4). While this technique in itself is not necessarily novel, the reason for its use and the accompanying analysis on efficiency optimization using large bus voltage (Section 3.3) offer new insights compared to previous work. For example, the design of the UT-SEA naturally lends itself to high-ratio speed reductions through the use of a ball screw and therefore directly benefits from higher motor speeds. Additionally, unlike

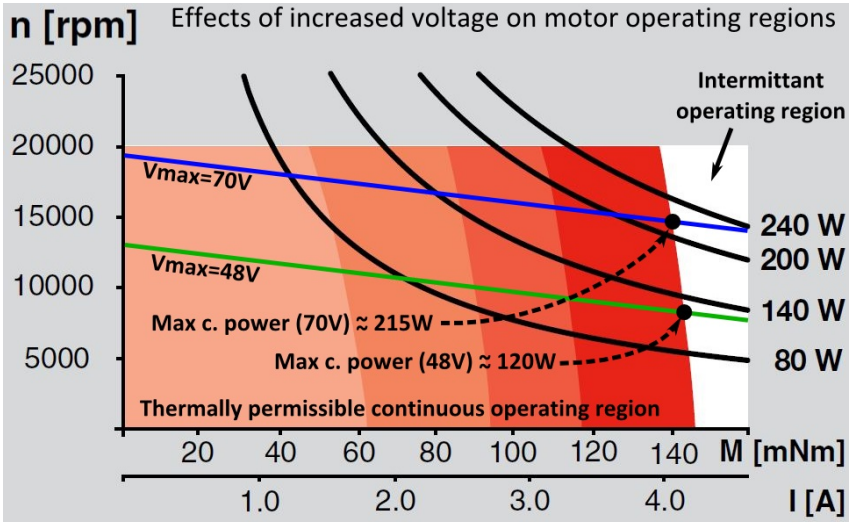

Figure 3.4: Effects of increased voltage on Maxon EC 22 operating regions (Maxon, 2014c). The lower line represents the motor speed-torque curve for $V_{max} = 48$ volts (the manufacturer specified value) and shows that maximum continuous power at this voltage is approximately 120 watts. By increasing bus voltage to 70 volts, both no-load speed and maximum continuous power are increased without violating the mechanical speed or thermal limits of the motor.

(Seok et al., 2012) which relies on a transparent drivetrain for force control accuracy, the UT-SEA incorporates mechanical force feedback in its compliant element, thus addressing the issue of reduced drivetrain transparency.

An additional benefit of increasing motor voltage instead of increasing motor current is preservation of motor efficiency. As illustrated in Figure 3.5, a motor's operating efficiency is a load-dependent quantity. At high speeds and small loads, electric motors are most efficient. Optimization of torque density in electric motors (Seok et al., 2012) ignores these effects, producing actuators with large currents and low-ratio speed reductions which do not operate at

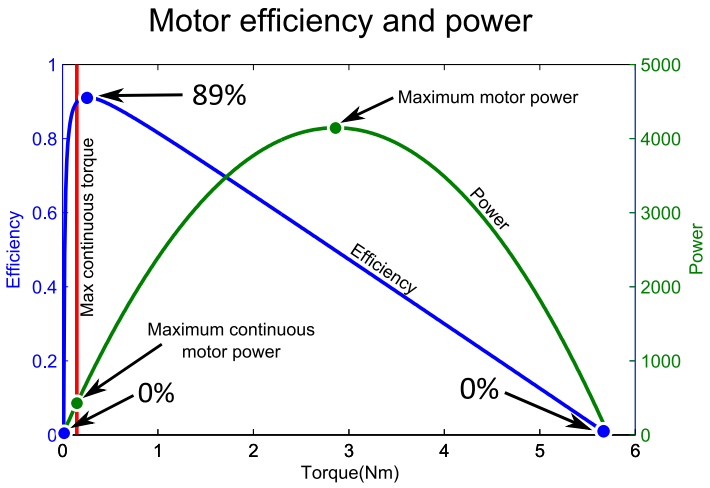

Figure 3.5: Plots of motor efficiency and motor power versus motor load (torque) (Maxon, 2014b). Note that peak efficiency and peak power do not coincide. Maximum motor power occurs at half stall torque. Maximum motor efficiency occurs at roughly 1/7th of stall torque. The linear section of the efficiency curve shows that motor efficiency generally is highest at high speed and low torque. This indicates that a large gear reduction, high-speed drivetrain design is optimal from a motor efficiency standpoint.

peak efficiency. Efficient high-ratio speed reductions preserve motor operating efficiency by allowing them to operate at high speeds.

3.2 Sources of Loss in the Energy Conversion Process

As is the case with all physical systems, many imperfections and losses occur in the process of converting energy from one domain to another. Attempts to increase electric actuator power output often neglect to consider system efficiency, deeming power output as the highest priority (Urata et al., 2008). However, as we will show in Section 3.3, designing for high power output does not preclude a design from selecting the local optima for efficiency as

well. After all, as discussed back in Chapter 1, system efficiency is a primary argument supporting electric actuation over other technologies (i.e. hydraulic), so its degradation should be minimized where possible.

In an electric actuation system (refer back to Figure 3.1) energy is first lost as heat due to internal battery resistance (Energizer, 2005). Then, energy is lost at the power modulation stage, including heating from bridge resistance, and energy loss due to high frequency state transitions of the switching elements (Rashid, 2001; Dodge, 2006; Fairchild Semiconductor, 1999). Additional losses occur when electric energy is converted into mechanical energy in the motor, such as: 1) resistive winding loss, 2) core losses which consist of eddy current loss and magnetic hysteresis loss, and 3) bearing friction loss which increases with speed and load torque (Kuria and Hwang, 2011). And finally, as discussed in Chapter 2, loss occurs during force amplification in the mechanical drivetrain of the actuator.

While many of these losses are unavoidable given existing electric motor technology, others can be reduced if certain system knowledge is available. In the following section, we leverage this system knowledge to reveal how loss can be minimized, and thus efficiency maximized, in high power electric actuation systems.

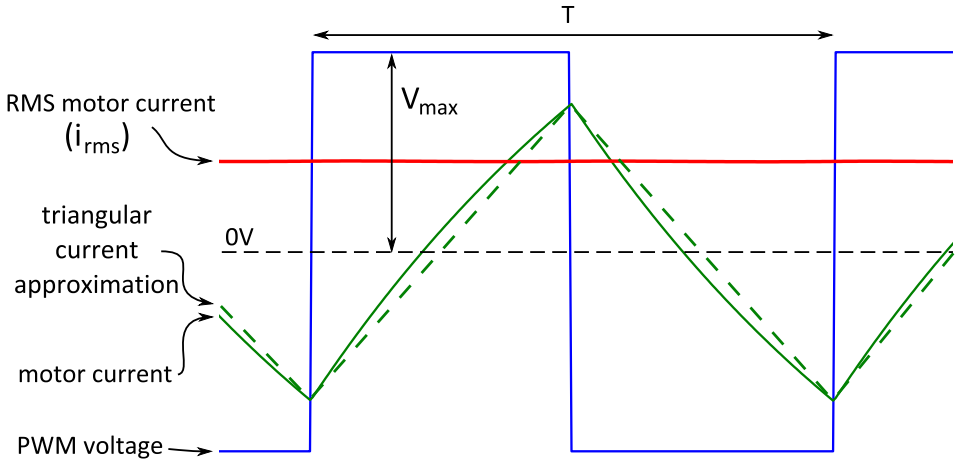

Figure 3.6: Current waveform induced by a PWM voltage applied to a motor.

3.3 Optimizing Energy Transfer in High Power Electric Actuation Systems

A Pulse-Width-Modulated (PWM) signal is typically used to efficiently modulate voltage applied to a motor (see Figure 3.6). For large bus voltages (high V_{max}), the transient current ripple created by this PWM voltage acting on the R-L motor circuit can rise to large values if the PWM switching frequency (f_{sw}) is not high enough (Figure 3.7). Even for an average current of zero (i.e. no useful mechanical work is performed), substantial RMS current can lead to motor overheating.

The steady-state RMS current (i_{rms}) induced in the motor winding with resistance R, inductance L, and time constant τ, due to PWM voltage (bus voltage V_{max} and frequency $f_{sw} = 1/T$) can be represented as follows (Williams, 1992):

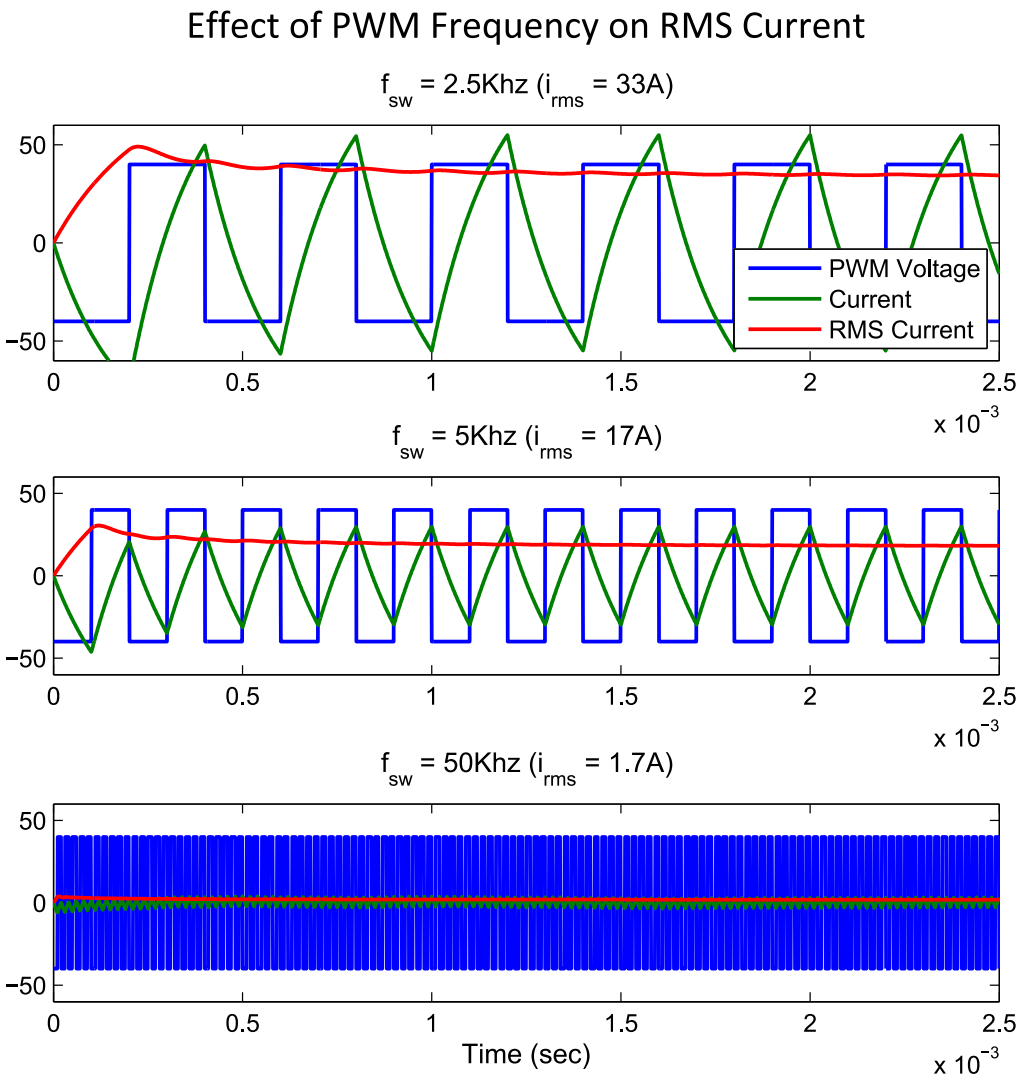

Figure 3.7: Current control through a motor is accomplished using a PWM voltage signal. For high motor voltage, transient currents may increase to large values. These series of plots shows the effect that increasing PWM frequency has on transient current. In each case, average current is zero. However, winding temperature is determined by RMS current, which is shown to decrease as PWM frequency increases.

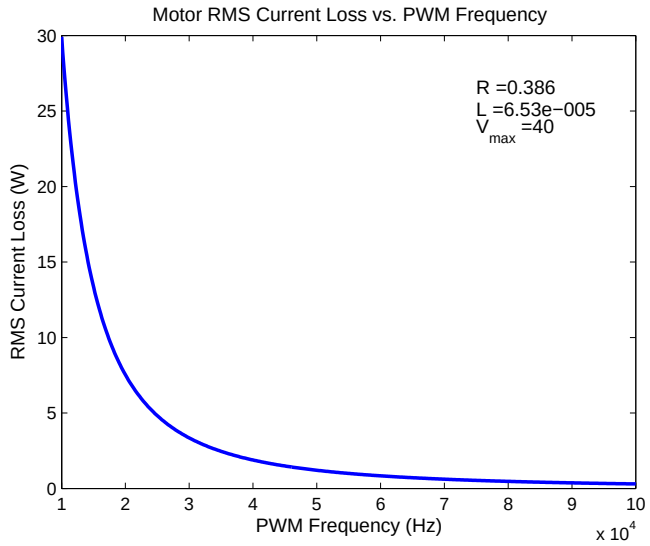

Figure 3.8: RMS current loss versus PWM frequency.

$$i_{rms} = \frac{V_{max}}{R}\sqrt{1 - 4\tau f_{sw}\tanh\left(\frac{1}{4\tau f_{sw}}\right)}. \tag{3.1}$$

RMS power loss due to heating of the motor windings can then be calculated

$$P_{rms} = {i_{rms}}^2 R = \frac{{V_{max}}^2}{R}\left[1 - 4\tau f_{sw}\tanh\left(\frac{1}{4\tau f_{sw}}\right)\right]. \tag{3.2}$$

Plotting (3.2) versus PWM frequency yields the relationship illustrated in Figure 3.8. As is shown, low PWM frequencies lead to large heating of the motor's windings. This result agrees with the three scenarios depicted in Figure 3.7. From this observation, we can see how high PWM frequencies are required to minimize loss due to RMS winding current.

Conversely, if the PWM frequency is too high, current drawn from the

capacitive charging and discharging of switching elements in the servo drive becomes large (Rashid, 2001; Dodge, 2006; Fairchild Semiconductor, 1999). For instance, to switch a MOSFET transistor to its ON state, a finite amount of charge (q_g) must flow into the gate terminal. The quantity of this charge is proportional to the MOSFET's gate capacitance. For each PWM cycle, a current (i_{sw}) flows between the low and high voltage supplies for each of n switching elements:

$$i_{sw} = q_g \, n \, f_{sw}. \tag{3.3}$$

Switching power loss is then computed given gate drive voltage (V_{dr}):

$$P_{sw} = V_{dr} \, i_{sw} = V_{dr} \, q_g \, n \, f_{sw}. \tag{3.4}$$

Equation (3.4) is plotted in Figure 3.9.

These two sources of power loss, RMS current ripple loss and switching loss, are both dependent on f_{sw}, but in opposite directions (see Figure 3.10). A trade-off therefore must be made of where loss will occur: is it better to dissipate heat in the motor by using a low PWM frequency or to dissipate heat in the servo drive bridge by using a high PWM frequency? Fortunately, an optimal answer exists and we may choose system parameters to minimize the total loss from these two sources.

To find the PWM frequency ($f_{sw_{optimal}}$) which minimizes total power loss (P_{total}), we must find the minimum value of P_{total}:

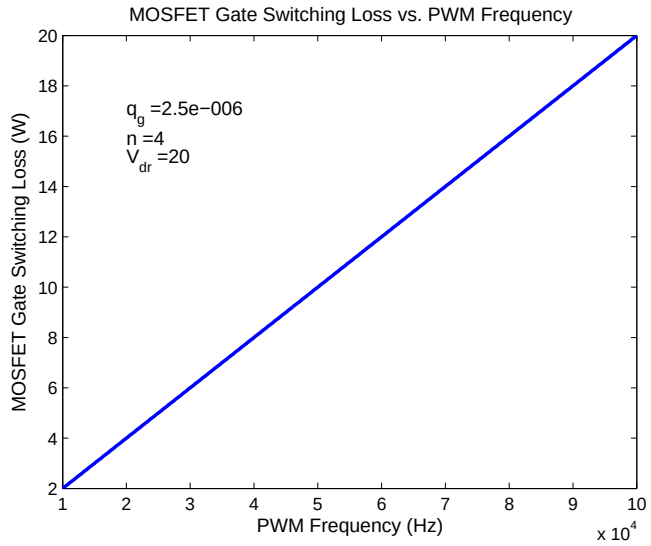

Figure 3.9: PWM switching loss versus PWM frequency.

$$P_{total} = P_{sw} + P_{rms} = V_{dr}\, q_g\, n\, f_{sw} + \frac{{V_{max}}^2}{R}\left[1 - 4\tau f_{sw} \tanh\left(\frac{1}{4\tau f_{sw}}\right)\right]. \quad (3.5)$$

However, because of the hyperbolic tangent function, we cannot minimize P_{total} analytically. Instead, we propose using a triangular waveform to approximate the RMS current (as shown in Figure 3.6). The accuracy of this approximation improves as f_{sw} increases (see Figure 3.11). Using this approximation, the triangular RMS current ($\widetilde{i}_{rms}$) is

$$\widetilde{i}_{rms} = \frac{V_{max}}{L f_{sw} 4\sqrt{3}}. \quad (3.6)$$

The approximate total power loss ($\widetilde{P}_{total}$) is then

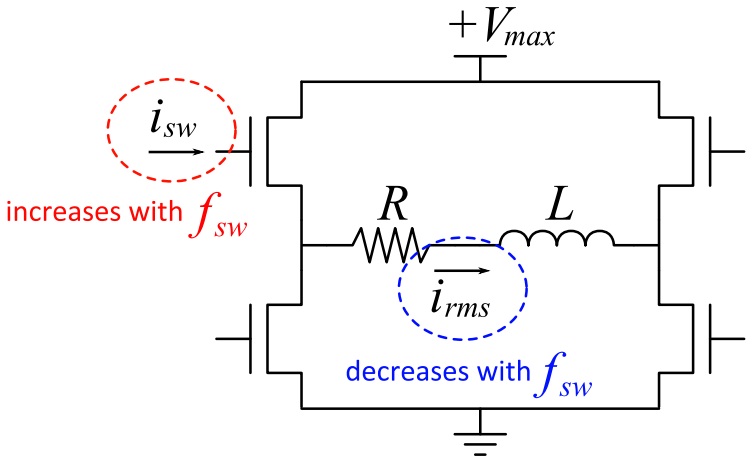

Figure 3.10: R-L circuit DC motor model including H-bridge circuitry showing capacitive gate switching current and RMS motor current. Loss due to heating from RMS current in the motor (i_{rms}) decreases with PWM switching frequency (f_{sw}) while capacitive loss due to switching elements in the servo drive increases with f_{sw}. These two sources of loss offer competing objectives.

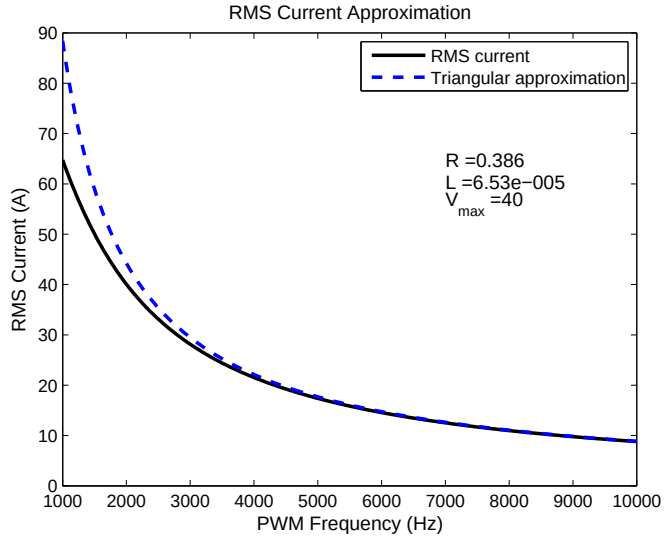

Figure 3.11: Approximation accuracy of triangular current waveform. The accuracy of the triangular approximation increases with PWM frequency.

$$\widetilde{P}_{total} = V_{dr}\, f_{sw}\, n\, q_g + \frac{R {V_{max}}^2}{48\, L^2\, {f_{sw}}^2} \tag{3.7}$$

which can be minimized analytically. Solving for the minimum power loss, we find $f_{sw_{optimal}}$ to be

$$f_{sw_{optimal}} = \sqrt[3]{\frac{R\, {V_{max}}^2}{24\, L^2\, V_{dr}\, n\, q_g}}. \tag{3.8}$$

Figure 3.12 illustrates the combined losses of RMS current ripple and capacitive MOSFET switching losses. The minimum loss point, which is produced by selecting the optimal PWM frequency using (3.8), is also shown.

3.4 Summary

Maximizing performance of electric motors is important because their performance improvements directly benefit actuator, and consequently, whole-robot capability. Several existing robots have explored exotic and effective techniques for increasing power output by increasing motor current. A less documented technique for increasing power output is to drive an electric motor with high voltage, thus enabling higher motor speeds.

Along with output power, the efficiency of power transfer from the servo drive to the motor output is a key metric of performance. This efficiency can be optimized for any given motor power maximization technique (i.e. using large currents or using large voltages). In optimizing efficiency, competing objectives are found in loss due to RMS current ripple in the motor, and loss

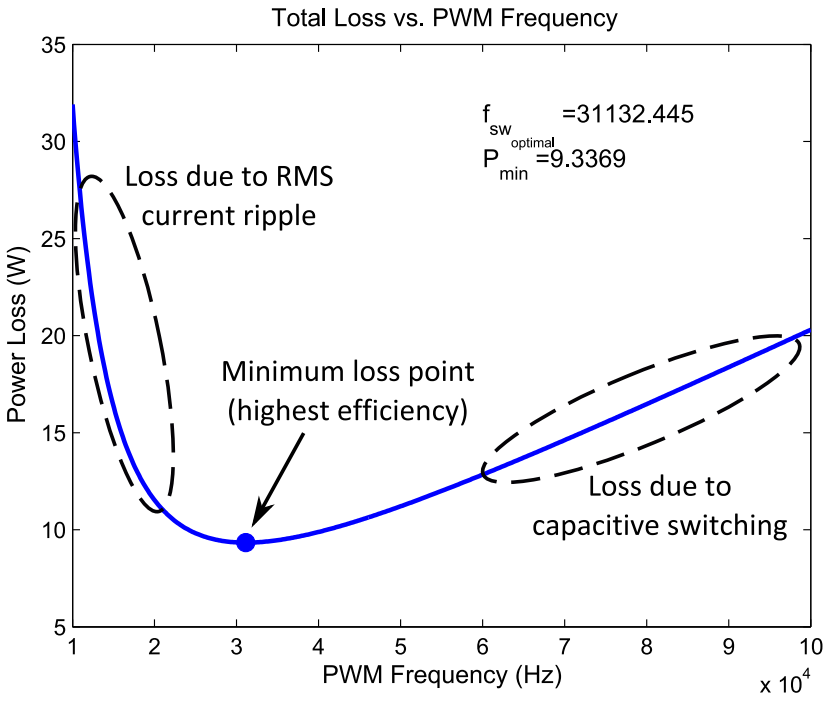

Figure 3.12: Optimization of PWM frequency for minimum loss in electric motor drives. Using (3.8) we can find the optimal PWM frequency which minimizes loss between the switching elements in the servo drive and current ripple in the electric motor.

due to capacitive switching of servo drive elements. Equation (3.8) solves this optimization problem and produces the PWM frequency which results in the most efficient power transfer from the servo drive to the motor output.

Chapter 4

High-Performance Model-Based Control of SEAs

In Chapters 2 and 3 we discussed designs and methods for increasing base actuation hardware capability and efficiency. Referring back to Figure 1.1, Chapters 2 and 3 therefore focused on the "bottom layer" hardware aspects of robot performance, defining the maximum performance potential. However, as discussed in Chapter 2, hardware alone does not create high-performing actuators and robots. In this chapter, we transition into the analysis of the "upper layers" of the system performance stack; that is, we discuss the algorithms and software which generate commands for the actuation hardware to follow. If these commands are not wisely chosen, the hardware will not be used to its full potential and performance will be lost.

The control approach demonstrated in this chapter heavily relies on force control, a control technique well studied in literature (Khatib, 1987; Pratt et al., 2001; Hogan, 1984). In the context of this chapter and Chapter 5, force

This chapter contains material from the following publications: (Paine, Oh, and Sentis, 2014b; Paine, Mehling, Holley, Radford, Johnson, Fok, and Sentis, 2014a). Mehling, Holley, and Radford contributed to this work by designing the hardware on which the proposed controllers were implemented. Johnson, Fok, and Sentis contributed by interfacing the presented work into the whole-body control framework on the Valkyrie robot.

control is used to decouple the dynamics of the actuator from the dynamics of its load. SEAs are unique in this regard as the spring itself decouples the reflected motor inertia from the load inertia. The spring in an SEA therefore creates a logical (and physical) boundary between these two systems: the actuator and its load. Chapter 5 studies this idea further, considering the effects of unmodeled loads on actuator force tracking capability.

We begin with an overview of work relating to force control of SEAs and highlight how our force control approach leverages high-performance ideas from this body of work while improving internal controller latency. We then develop generic SEA models considering both types of spring placement architectures (FSEA and RFSEA, previously discussed in Chapter 2). Using these models, an SEA force control architecture is presented along with advice on how to properly tune it for optimal performance. Finally, a model-based position controller is introduced which builds upon the force controller and is used to demonstrate high-performance behaviors using the previously discussed UT-SEA actuator (Chapter 2).

4.1 SEA Force Control Background

Many different control architectures have been proposed for force control of series elastic actuators. Some of the variation in controller design is rooted in differences imposed by hardware. For example, force can be observed either by measuring change in resistance, as is accomplished using strain gauges in (Pratt and Williamson, 1995), or by measuring spring deflection and

applying Hooke's law, as shown in (Kong et al., 2010). A control strategy for hardware designs using spring deflection sensors may treat a motor as a velocity source, and transform desired spring forces into desired spring deflections. However, for hardware designs using strain gauges, the force sensor does not output an intermediate displacement value, but maps change in resistance directly to applied force. For such a system, modeling the motor as a force source is more convenient.

Further classification of SEA force control strategies may be made based on the types and combinations of control structures used. Some approaches measure spring force and control motor force using some variant of PID control structures (P, PD, etc.) (Sensinger and Weir, 2006; Garcia et al., 2011; Ragonesi et al., 2011). If friction and backlash are significant, a single-loop PID force controller may become unstable before the desired force tracking is achieved. To remedy this issue, an inner position or velocity control loop may be used with an outer force control loop as proposed by (Pratt et al., 2004). This idea has been adopted and carried on by many others, translating force control into a position or velocity tracking problem (Thorson and Caldwell, 2011; Lagoda et al., 2010; Vallery et al., 2007; Wyeth, 2006; Hutter, 2013). Another class of controllers use PID control structures but also consider the dynamics of the mechanical system to improve the frequency response of force control (Pratt and Williamson, 1995; Hurst et al., 2010). Others show how steady state tracking and disturbance rejection can be significantly improved through the use of disturbance observers (Kong et al., 2009). Our approach

differs from this earlier work in that PD feedback is used to shape the dynamic response of the actuator, eliminating the need for a feedforward filter term and improving the controller's phase margin.

4.2 Generic Models of FSEAs and RFSEAs

As previously discussed in Chapter 2, the compliant element in an SEA may be placed between the gearbox output and the load (FSEA) or between the motor housing and the chassis ground (RFSEA). In this section we develop models which apply to generic SEAs of either type. Based on these models, we will observe properties of each style of actuator which will help us develop our model-based SEA controllers.

Figure 4.1 shows simple models for both FSEA and RFSEA style actuators. In the FSEA model (Figure 4.1c), generalized motor force (F_m) is generated between chassis ground and a lumped sprung mass (m_k) which includes rotor inertia, the gearbox reduction, and transmission inertia. If the motor is unpowered and backdriven, a viscous backdriving friction (b_b) is felt from transmission friction and motor friction. The spring is between the transmission output and output mass (m_o) and has stiffness (k) and viscous friction (b_k) generated by the spring support mechanism. In the RFSEA model the spring and force generation elements are switched. In addition, the distribution of sprung mass and output mass is different for an RFSEA. m_o in the RFSEA model includes rotor inertia, the gearbox reduction, and transmission inertia. m_k varies by design. For the UT-SEA, m_k includes the mass of the

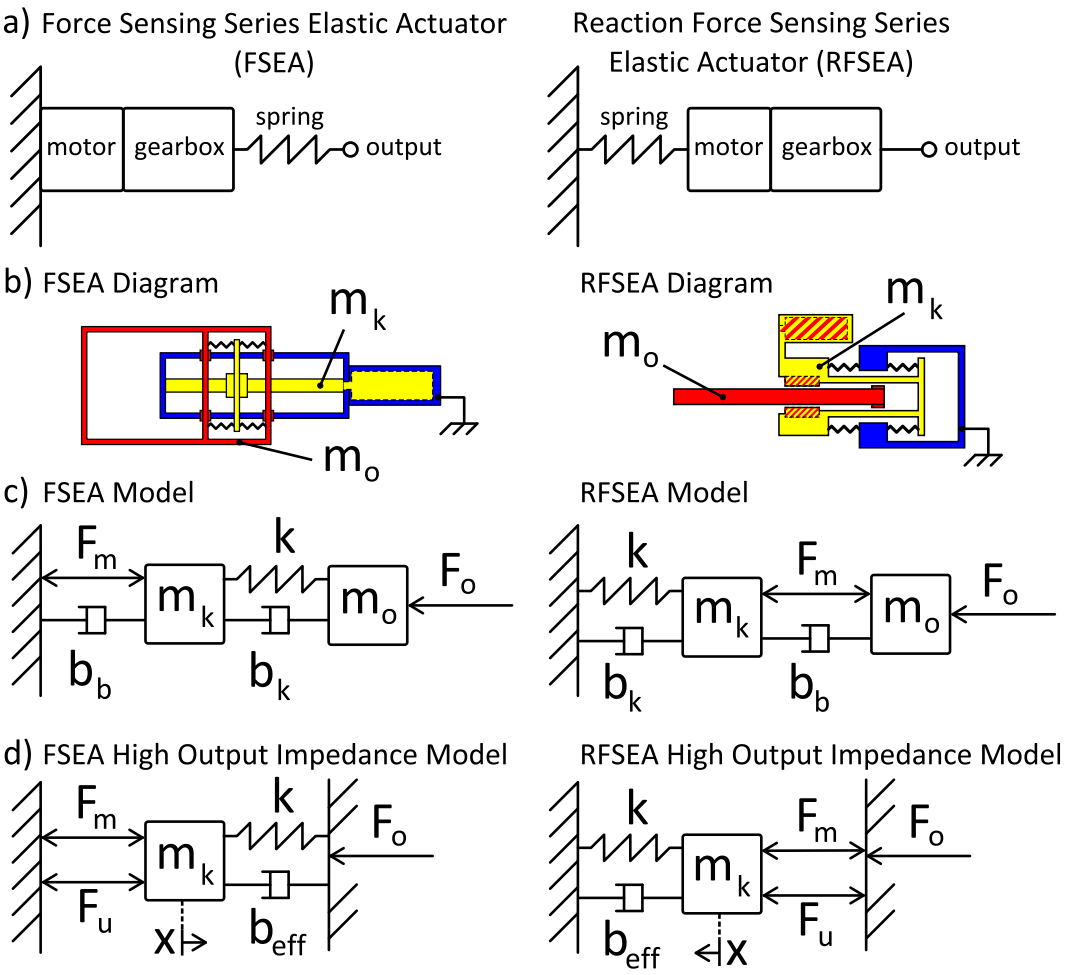

Figure 4.1: Models for FSEA and RFSEA style actuators. The notations represent F_m: motor force, F_o: output force, b_b: viscous backdriving friction, b_k: viscous spring friction, k: spring constant, x: spring deflection, m_k: lumped sprung mass, m_o: output mass, b_{eff}: lumped damping which equals $b_b + b_k$, F_u: disturbance forces and forces which are difficult to model.

actuator housing and mass of the motor, including rotor mass.

A high output impedance model is useful for simplifying the force controller design problem (Figure 4.1d). It assumes that the actuator output is rigidly connected to an infinite mass, which cannot be moved. For the high output impedance models, the sprung mass experiences a summation of forces from 1) the motor (F_m), 2) the spring (F_k), 3) lumped viscous friction ($F_{b_{eff}}$) where $b_{eff} = b_b + b_k$, and 4) from other disturbances that are difficult to model (F_u) such as torque ripple from commutation, torque ripple from gearboxes due to teeth engaging and disengaging, backlash, and various forms of friction such as stiction, and coulomb friction.

The force sensing challenge for each SEA configuration is to calculate force felt at the actuator output (F_o) given measurement of spring deflection (x). For both FSEAs and RFSEAs, the spring acts as the force sensor. The difference between the two is where the spring is located relative to the output force. This discrepancy does not affect low frequency force measurement but must be taken into account to measure high frequency forces accurately.

For FSEAs with high output impedance (Figure 4.1d), $F_o = F_k + F_{b_{eff}}$. If $F_{b_{eff}}$ is small, which can be accomplished with careful mechanical design, then F_o can be closely approximated by measuring F_k alone. This simplification removes the need to measure or calculate time derivatives of x.

$$F_{o_{FSEA}} = F_k = kx \tag{4.1}$$

For RFSEAs with high output impedance, $F_o = F_m + F_u = F_{m_k} + F_{b_{eff}} + F_k$. Here, measurement or calculation of time derivatives of x is critical in considering the large forces of F_{m_k} and $F_{b_{eff}}$. Output force can be observed as follows.

$$F_{o_{RFSEA}} = F_{m_k} + F_{b_{eff}} + F_k = m_k\ddot{x} + b_{eff}\dot{x} + kx \tag{4.2}$$

These equations tell us that force sensing for RFSEA style actuators should possess models of sprung mass and viscous damping and should be able to measure or calculate both $\dot{x}$ and $\ddot{x}$ for accurate force sensing across the frequency spectrum.

An additional drawback of RFSEAs is revealed when considering internal forces required to generate a desired output force. For an FSEA, the relationship between F_k and F_o is given by:

$$\frac{F_k(s)}{F_o(s)} = \frac{k}{sb_{eff} + k} \tag{4.3}$$

Whereas for an RFSEA the same relationship is given by:

$$\frac{F_k(s)}{F_o(s)} = \frac{k}{s^2m_k + sb_{eff} + k} \tag{4.4}$$

Similarly, the relation between F_m and F_o for FSEAs (4.5) and RFSEAs (4.6) are as follows:

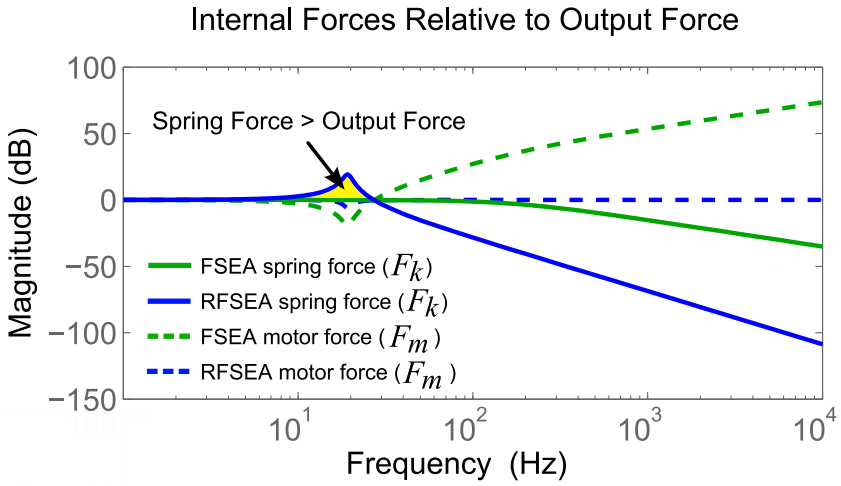

Figure 4.2: Relation of internal actuator forces to output force. The four lines represent equations (4.3), (4.4), (4.5), and (4.6). As the figure shows, RFSEA spring force exceeds output force for resonant frequencies. Parameters are selected to match the UT-SEA design.

$$\frac{F_m(s)}{F_o(s)} = \frac{s^2 m_k + s b_{eff} + k}{s b_{eff} + k} \tag{4.5}$$

$$\frac{F_m(s)}{F_o(s)} = \frac{s^2 m_k + s b_{eff} + k}{s^2 m_k + s(b_{eff} + b_b) + k} \tag{4.6}$$

Plotting the frequency response of (4.3), (4.4), (4.5), and (4.6) yields the results shown in Figure 4.2. Each line represents an internal force relative to the actuator output force across the frequency spectrum. A magnitude greater than one (0 dB) indicates that, at the represented frequency, the internal force value is greater than output force. Motor force for both the FSEA and RFSEA remain less than or equal to the output force below the resonant frequency. Motor force for the FSEA increases for frequencies greater than the resonant

frequency because the motor force must counteract the low-pass filter created by the mass-spring system. The RFSEA is able to produce high frequency forces without an increased burden to the motor force. The problem regarding the RFSEA design lies in the spring force, which is the only internal force to increase above output force for resonant frequencies and below. The resonant peak for spring force is about 15 dB (a factor of around five) for the UT-SEA design (see Chapter 2), meaning spring force is five times greater than output force. One way to address this issue is to regulate spring force rather than output force for RFSEA style actuators. In doing so, force tracking accuracy is sacrificed at the actuator's resonant frequency.

To summarize key differences between FSEAs and RFSEAs:

1. Accurate force sensing for RFSEAs requires knowledge of k, b_{eff}, m_k, x, $\dot{x}$, and $\ddot{x}$ whereas FSEAs only require k and x for a close approximation of output force.

2. FSEA output force can safely track a reference force signal up to and past resonant frequencies but will require large motor effort at high frequencies. RFSEAs cannot safely track reference force signals close to their resonant frequencies due to large resonant spring forces, but can track high frequency force signals with low motor effort.

3. FSEAs place a mechanical low-pass filter between the output and the gearbox, making them more tolerant to impact forces than RFSEAs.

Based on these observations FSEAs are better suited for force control applications. However, as discussed in Chapter 2, the excellent size and packaging characteristics of RFSEAs can provide significant advantages for integration with dense system designs.

4.3 Model-Based Control using Disturbance Observers

In this section, we apply the models obtained in Section 4.2 to develop high-performance force and position controllers for SEAs. The end goal of this section is to produce behaviors which achieve high power (high torque and high speed simultaneously) in a fixed range of motion. Obtaining high velocities safely within an actuator's fixed range of motion requires well-controlled bursts of acceleration to and from rest.

Control effort for such motions can come from two sources: feedforward (model-based) effort and feedback (reactive) effort (Bèlanger, 1995). One of the primary influences of the compliant element in an SEA is to reduce actuator output impedance. As a result, achieving stiff active impedance behaviors with SEAs is difficult due to large motor energy requirements (Pratt and Williamson, 1995) and control loop stability issues, the limits of which have not yet been formally studied. Because of these factors, high-performance behaviors are difficult to achieve with SEAs using predominately reactive control techniques.

Instead, our SEA position control approach uses a model of the load's dynamics, obtained using system identification techniques, to achieve high-

performance behaviors. As such, our method heavily relies on feedforward control effort. Using such an approach, the position bandwidth of an SEA is only limited by its force bandwidth through Newton's Second Law:

$$\begin{aligned} x_d &= A\,sin(\omega t) \\ \ddot{x}_d &= A\,\omega^2\,sin(\omega t) \\ F &= m\,\ddot{x}_d = m\,A\,\omega^2\,sin(\omega t) \end{aligned} \tag{4.7}$$

where x_d is a desired position profile, A and ω are the signal amplitude and frequency, respectively, m is the load mass, and F is the necessary tracking force. As is evident in (4.7), both position and force signals have the same signal frequency.

While in this chapter we use our proposed force controller to act as an inner loop for our model-based position controller, this force controller is generic and can be applied to any other application where actuator forces are the control target (see Chapter 6 for another example).

4.3.1 Force Control

Our proposed force controller uses a structure similar to (Kong et al., 2009) but differs in that the closed-loop dynamics are shaped using PD feedback alone instead of using a feedforward filter. By excluding a feedforward filter, a source of delay in the signal tracking transfer function is removed. In this section we give an explanation of the force control structure and characterize its performance through a series of experimental tests.

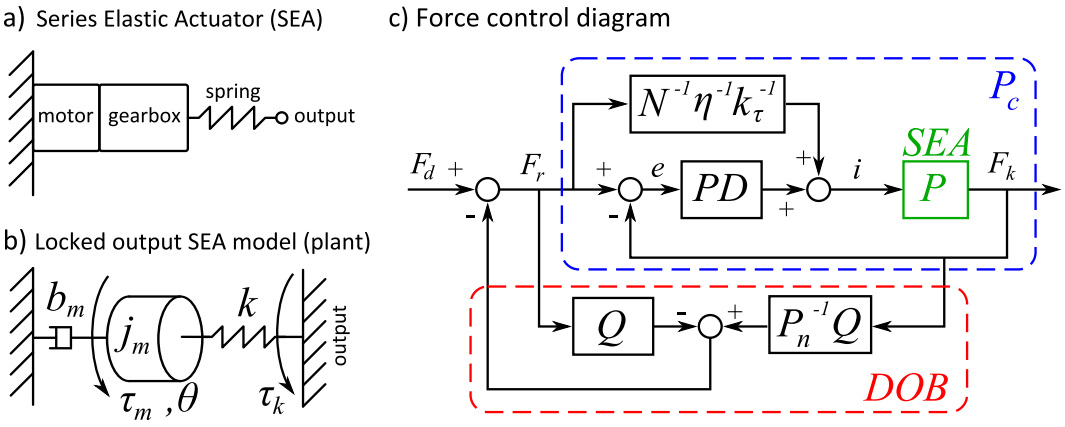

Figure 4.3: **a)** Schematic representation of a series elastic actuator. **b)** Plant model (P) used for control design. Note that a locked-output assumption is used. **c)** Diagram of the proposed force feedback controller for series elastic actuators. The PD compensator is used to shape the dynamics of the force response while the DOB is used to improve disturbance rejection, especially at low frequencies. The Q filter is defined by (4.17).

As discussed in Section 4.2, both FSEAs and RFSEAs should be controlled using their spring force. Therefore, the controller proposed here equally applies to both FSEA and RFSEA style actuators.

Figure 4.3 shows our proposed force control structure. The control plant is an SEA with a locked output, as shown in Figure 4.3b. The inner PD compensator is tuned to produce the desired frequency response based on this locked-output assumption. A disturbance observer (DOB) is then used to reject deviations from this nominal locked-output model and maintain force tracking accuracy. Further discussion on DOB disturbance rejection for non-locked-output scenarios will be held in Chapter 5.

The control plant transfer function (P) from motor current (i) to mea-

sured spring force (F_k) is a function of the effective motor inertia felt by the spring (m_k), the effective motor-side damping felt by the spring (b_{eff}), spring stiffness (k), speed reduction (N), motor torque constant (k_τ), and drivetrain efficiency (η):

$$P(s) = \frac{F_k(s)}{i(s)} = \frac{Nk_\tau \eta k}{m_k s^2 + b_{eff} s + k} = \frac{\beta k}{m_k s^2 + b_{eff} s + k}. \tag{4.8}$$

The scaling factor between motor current and motor-side output force is represented by β where $\beta = Nk_\tau \eta$. Assuming k is measured beforehand, all of the parameters in (4.8) can be found using system identification techniques (see Appendix A.1) with the actuator output locked.

The closed-loop transfer function from reference force (F_r) to F_k for the locked-output control plant is represented by P_c:

$$P_c(s) = \frac{F_k(s)}{F_r(s)} = \frac{(k\beta k_d)s + k(1 + \beta k_p)}{m_k s^2 + (b_{eff} + k\beta k_d)s + k(1 + \beta k_p)}. \tag{4.9}$$

As depicted in Figure 4.3c, P_c is composed of a feedforward term ($N^{-1}\eta^{-1}{k_\tau}^{-1}$) and a feedback term ($PD$). The feedforward term is used to scale desired actuator forces into approximate actuator output forces to minimize control effort from the feedback term. The feedback term is represented by the following transfer function:

$$PD(s) = \frac{i(s)}{e(s)} = k_d s + k_p. \tag{4.10}$$

For simplicity, we do not model the effects of a low-pass filter that is applied to the derivative term. In practice, the cutoff frequency for the $k_d s$ term (f_{kd}) in (4.10) is chosen to be sufficiently higher than the actuator's maximum closed-loop system bandwidth. The effects of f_{kd} on the force tracking response will be shown later in Figure 4.4.

4.3.1.1 Simplification of Feedback Gain Selection

When faced with highly parameterized feedback controllers, it is often unclear how each parameter should be chosen. Clearly, a trade-off between poor performance (low gains) and poor stability (high gains) must be found. In practice, gain tuning is often performed manually to quickly locate parameters that balance these two trade-offs. From (4.10) we see that two parameters must be found: k_p and k_d. Instead of choosing k_p and k_d directly, it can be more insightful to consider an equivalent pair of feedback parameters: bandwidth and damping ratio. This is because a desired "shape" of the frequency response can be chosen by selecting a desired damping ratio, and then bandwidth may be chosen independently to satisfy the trade-off between performance and stability. Here, we derive the equations mapping damping ratio to k_p and k_d.

Notice that the characteristic polynomial of (4.9) is a second order system that can be represented in terms of an effective mass ($\hat{M}$), spring, ($\hat{K}$) and damper ($\hat{B}$):

$$\hat{M} = m_k \tag{4.11}$$

$$\hat{B} = (b_{eff} + k\beta k_d) \tag{4.12}$$

$$\hat{K} = k(1 + \beta k_p). \tag{4.13}$$

The damping ratio (ζ) for such a second order system is

$$\zeta = \frac{\hat{B}}{2\sqrt{\hat{M}\hat{K}}} \tag{4.14}$$

Combining (4.11), (4.12), (4.13) and (4.14) forms a relation between desired closed-loop damping ratio (ζ_d) and k_p, k_d:

$$\zeta_d = \frac{(b_{eff} + k\beta k_d)}{2\sqrt{m_k k(1 + \beta k_p)}}. \tag{4.15}$$

Solving (4.15) for k_d yields

$$k_d = \frac{2\zeta_d\sqrt{m_k k(1 + \beta k_p)} - b_{eff}}{k\beta}. \tag{4.16}$$

Using (4.16), a desired damping ratio can be chosen and k_d can then be automatically calculated for a given k_p (see Figure 4.4). In tuning this force controller, k_p is used to represent bandwidth, therefore simplifying the tuning of the PD compensator to a single degree of freedom. Using this one parameter, the trade-offs between performance and stability can easily be changed on-the-fly while ensuring a dynamic response with the desired damping ratio.

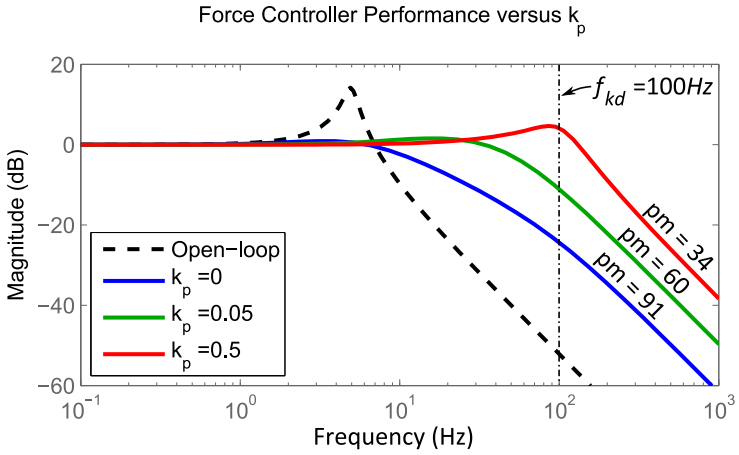

Figure 4.4: The effect of varying k_p on force controller performance (F_k/F_r). k_d is recalculated based on (4.16) for each value of k_p. As is shown, force control bandwidth increases as k_p increases while stability (phase margin, pm in the figure) decreases. While each response has the same desired damping ratio, a resonant peak begins to develop for larger values of k_p due to the limitations of the k_d cutoff frequency (f_{kd}).

For the UT-SEA, we chose a desired damping ratio of 0.9 to produce a flat force response with little overshoot ($\zeta_d = 1.0$ is critically damped, $\zeta_d = 0.7$ is underdamped with minimum settling time).

Using these tuning techniques, the primary factors determining maximum force control bandwidth are 1) closed-loop stability, which largely depends on the value of f_{kd} as demonstrated in Figure 4.4, and 2) actuator power consumption. Extending the bandwidth of any passive system inevitably results in increased power consumption at frequencies above the passive system bandwidth (Bèlanger, 1995). Large power consumption can be a problem due to loss in the actuation system, which manifests itself as heat as discussed in Chapter 3. Figure 4.5 demonstrates this behavior for two different values of

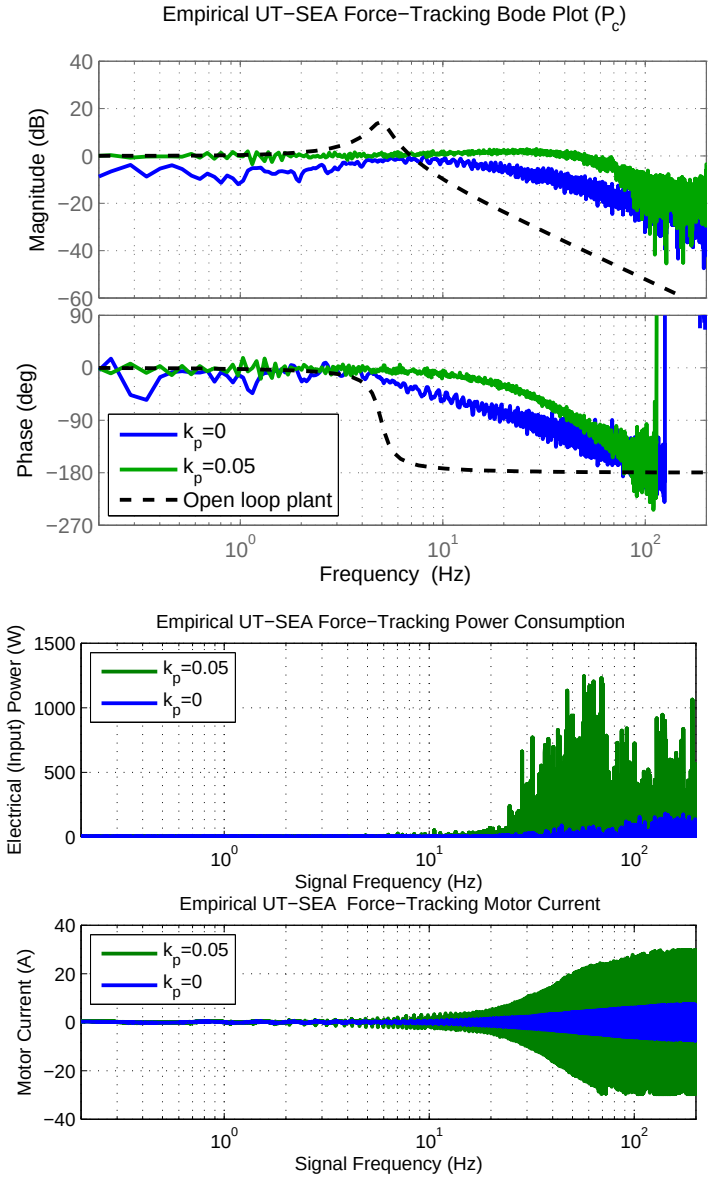

Figure 4.5: These plots demonstrate how power consumption plays an important role in feedback gain selection. The top two plots show bode plots for two different values of k_p, demonstrating better tracking and larger bandwidth for the larger k_p value. Correspondingly, the lower plots show the power consumption and motor currents for each experiment. Here, the larger feedback gain shows a significant increase in energy consumption at high frequencies.

k_p using the UT-SEA.

4.3.1.2 Disturbance Observer

The use of a disturbance observer (DOB) applied to an inner PID/PD control loop has been shown to significantly enhance the force tracking capability of SEAs (Kong et al., 2009, 2012; Paine et al., 2014b). DOBs have several useful properties for our specific application. First, they preserve and enforce a dynamic plant model through the use of a nominal model. This means that the shape of the closed-loop frequency response of P_c will not be altered by adding a DOB. The DOB will try to maintain this characteristic response in the presence of either 1) external disturbances or 2) plant model variations. The latter characteristic is of central importance to our approach and is discussed in detail in Chapter 5. Secondly, DOBs excel at removing steady state error, and therefore effectively serve as an integral feedback term. This characteristic is useful in minimizing controlled force resolution as discussed in Chapter 6.

A DOB applied to P_c is shown in Figure 4.3c. ${P_n}^{-1}$ represents the inverse of the nominal closed-loop model (P_c, using the locked-output constraint). Q is a low-pass filter which is used both to make ${P_n}^{-1}$ proper and to tune the frequency (f_q) up to which disturbances are rejected. In the tuning of the UT-SEA, f_q values in the range of 20 Hz to 50 Hz were found to adequately reject unmodeled disturbances while maintaining high control loop stability. Q takes the form of a low-pass Butterworth filter and has the following transfer

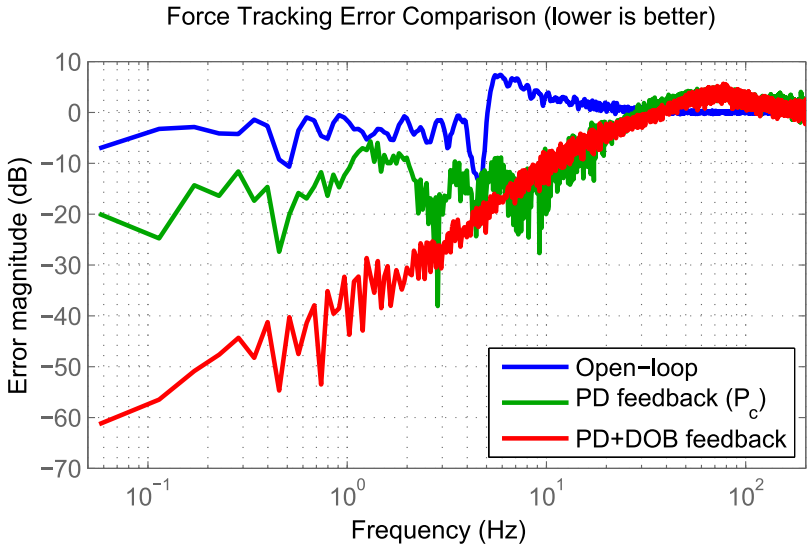

Figure 4.6: Measured force tracking error versus frequency for various controller configurations. PD feedback (P_c in Figure 4.3c) improves force tracking accuracy at most frequencies compared to open loop approaches. The addition of DOB feedback (full controller in Figure 4.3c) drastically improves low frequency tracking while preserving high frequency system dynamics.

function:

$$Q(s) = \frac{1}{(s/2\pi f_q)^2 + 1.4142(s/2\pi f_q) + 1}. \tag{4.17}$$

The complete force control parameters used for the UT-SEA are shown in Table 4.1.

4.3.2 Position Control

Our proposed position controller builds upon the force controller discussed in the previous section, treating it as a low-pass-filtered ideal torque

Table 4.1: UT-SEA force control parameters

Parameter	Value	Units
k_p	0.05	A/N
f_{kd}	100	Hz
ζ_d	0.9	n/a
f_q	40	Hz
β	219	N/A
m_k	360	kg
b_{eff}	2200	Ns/m
k	350000	N/m

source. This force source generates a force through a mechanical linkage with a moment arm (L) as depicted in Figure 4.7. Actuator force (F) generates arm torque (τ_a) depending on arm angle (θ_a) according to the following equation.

$$\tau_a = FL(\theta_a) = F\frac{cb\sin\theta_a}{\sqrt{b^2+c^2-2bc\cos\theta_a}} \tag{4.18}$$

The dynamics relating τ_a to θ_a with arm inertia (J_a) and joint friction (B_a) are:

$$\tau_a = J_a\ddot{\theta}_a + B_a\dot{\theta}_a + \tau_g(\theta_a) \tag{4.19}$$

where τ_g is the torque due to gravity and is parameterized by the mass of the output link (m_a), the distance from the point of rotation to the center of mass (l_m) and an angle (ϕ) to correct for c in Figure 4.7 not being orthogonal to the gravity vector.

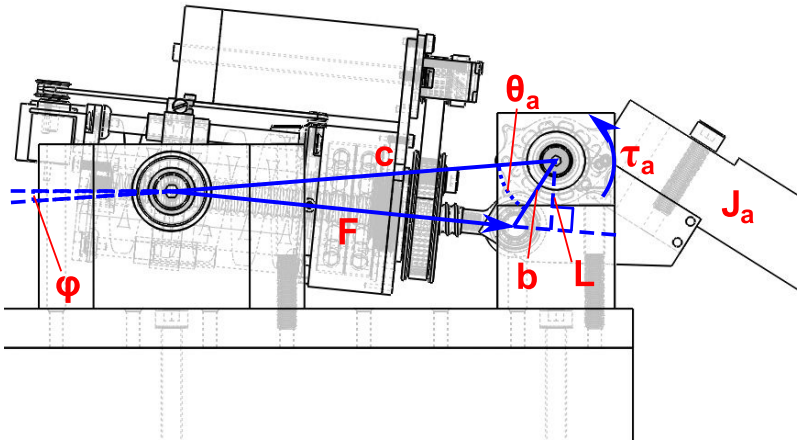

Figure 4.7: UT-SEA mounted on a test bench with the prismatic linkage geometry shown. The notations represent: L: linkage moment arm, c: distance between the actuator pivot and the arm pivot, b: distance between the arm pivot and the pushrod pivot, F: actuator force, τ_a: torque exerted on the output arm, θ_a: output arm angle, J_a: inertia of the output arm, ϕ: offset angle. Values used during testing of the actuator are: b = 0.025 [m], c = 0.125 [m]. Speed reduction from motor output to arm output is 175:1 at $\theta_a = 90°$.

$$\tau_g(\theta_a) = -m_a g l_m \cos(\theta_a + \phi) \tag{4.20}$$

Combining (4.18) (4.19) and (4.20) the full dynamics from F to θ_a are then represented by the following nonlinear differential equation.

$$F = \frac{\sqrt{b^2 + c^2 - 2bc\cos\theta_a}}{cb\sin\theta_a}\left[J_a\ddot{\theta}_a + B_a\dot{\theta}_a - m_a g l_m \cos(\theta_a + \phi)\right] \tag{4.21}$$

Our position control approach first considers the problem of controlling θ_a given τ_a, assuming no gravity is present. The relation between τ_a and θ_a in this case is given as:

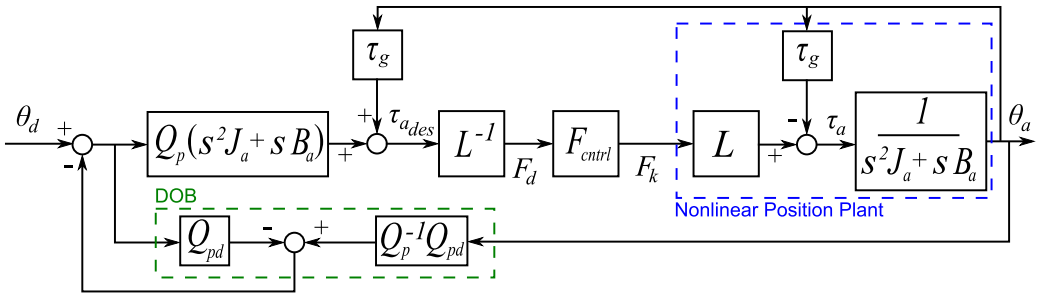

Figure 4.8: Block diagram of the control structure used for position control. The notations represent: τ_g: gravity compensation torque, L: nonlinear linkage kinematics, F_{cntrl}: the force control block shown in Figure 4.3. The Q functions are low-pass filters defined by (4.17).

$$\frac{\theta_a(s)}{\tau_a(s)} = \frac{1}{s^2 J_a + sB_a} \tag{4.22}$$

Inverting (4.22) provides a desired arm torque ($\tau_{a_{des}}$) given a desired arm angle (θ_d) and is used as the initial block in the position controller (Figure 4.8). Because (4.22) does not consider gravity, the desired arm torque signal must be summed with a gravity compensation torque (4.20) to produce the expected motion. The resulting torque value is then converted into desired actuator force by multiplying by the inverse of the nonlinear kinematics (L^{-1} from (4.18)). This desired force is then passed to the force controller.

Without some form of feedback the position controller would not be able to track a desired position due to modeling error and external disturbances. A DOB is placed in an outer loop around the model-based position controller to resolve these issues. The DOB treats modeling error and exogenous input as a disturbance and counteracts this disturbance with input to the model-based

position controller. Q_p (f_c = 10Hz) in Figure 4.8 is a feedforward low-pass filter to smooth position response, thus reducing required torques. Q_{pd} (f_c = 35Hz) is a low-pass filter that attenuates high frequency disturbance signals of the DOB.

4.4 Performance Experiments

Our goal was to design hardware and controllers that would maximize performance, but how do we know if we have been successful? One way of measuring the success of the control design is to attempt to reach the mechanical limits of actuator components in a safe and controlled manner. To this end we performed an experiment to push actuator speed to the limits of mechanical and control capabilities. A 5th order spline was used to generate a smooth position reference signal for high-speed transitions between a large angle displacement (60 degrees, see Figure 4.9). Figure 4.10 shows the experimental results. The arm is able to track the reference position closely and achieves a velocity of 15 rad/sec which is the mechanical limit of the ball screw. In this test, the motor reached a speed of 22,600 rpm which is 3,000 rpm below the maximum possible motor speed. Acceleration from rest, to maximum speed, and back to rest occurs within less than 0.2 seconds.

A critical metric for performance is power output. To maximize achievable power we designed an experiment which would require high speed and high torque simultaneously (see Figure 4.11). The design of the output arm to which the actuator is attached allows for additional weight to be added.

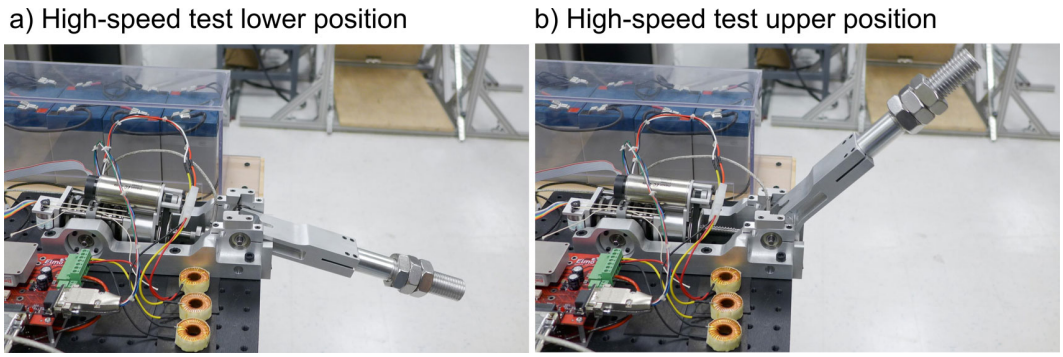

Figure 4.9: High-speed position tracking test setup.

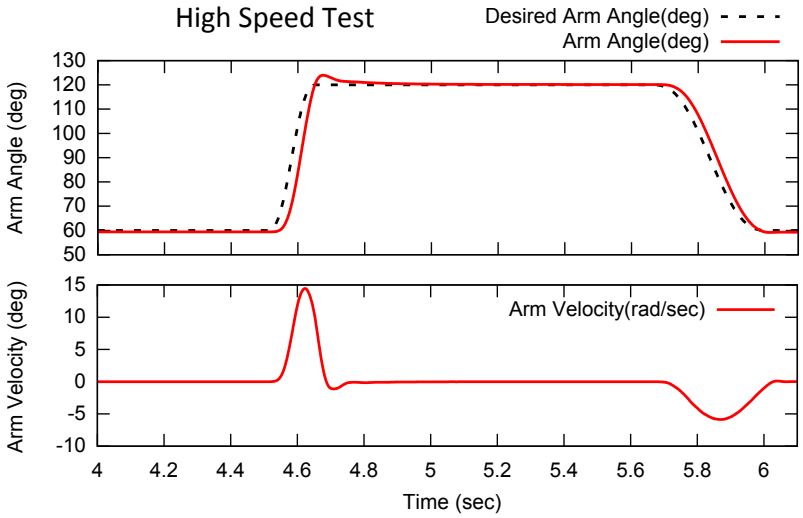

Figure 4.10: **High-speed position tracking test.** The actuator output follows a reference signal that changes 60 degrees in less than 0.2 seconds. The actuator is able to track the reference signal closely and reaches the maximum mechanical speed of the ball screw of 15 rad/sec. The torque moment arm (L) rotates through its largest length at $\theta_a = 90°$.

For the experiment, we fixed a 4.5 kg weight to the arm with a 0.23 meter moment arm. The experiment requires the arm to track a reference position which is again generated using splines. The motion is not symmetrical. When

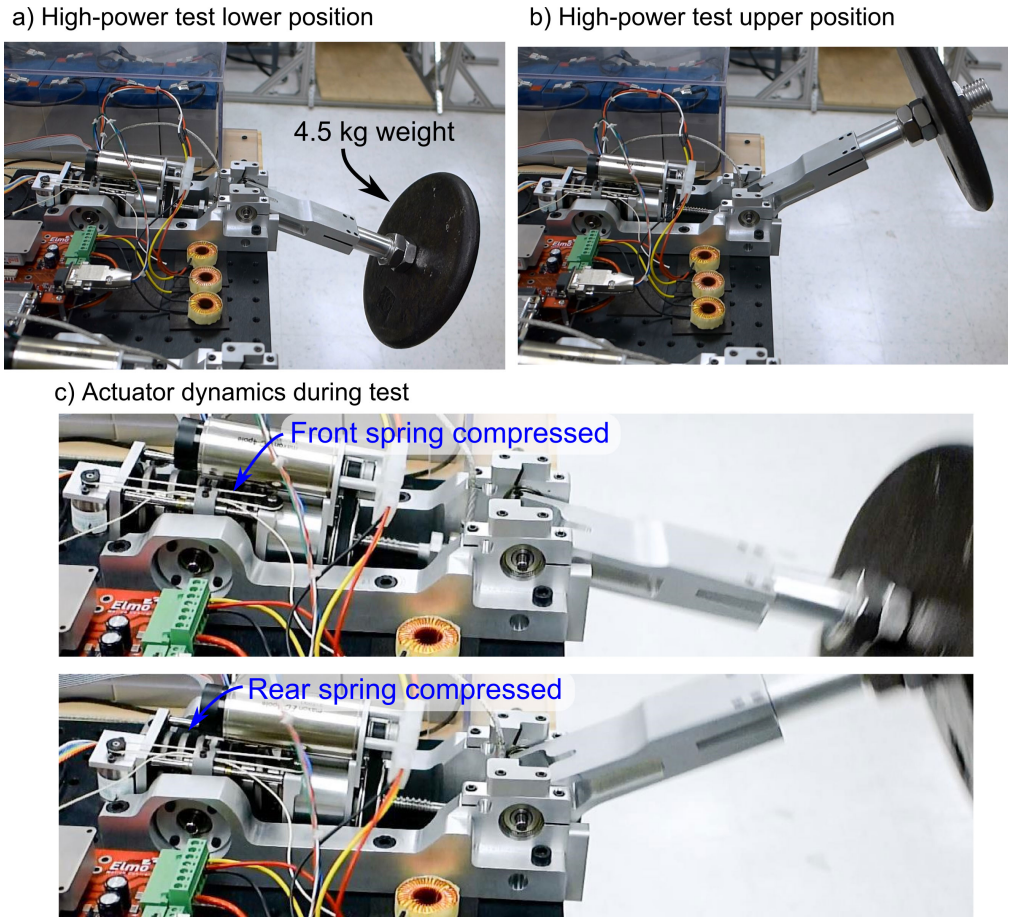

Figure 4.11: High-power position tracking.

the weight is being lowered the reference signal changes slowly and when the weight is being raised the reference signal changes quickly. The combination of fast motion and the fact that the motion is directed upwards against gravity makes this test require very high actuator power output. Figure 4.12 shows the experimental results. On the bottom graph power at the motor and power at the arm can be seen. Output power is measured as output torque times output velocity. The actuator generates peak mechanical output power of ap-

proximately 429 watts, which corresponds to a power-to-weight ratio of 423 watts per kilogram. Comparing with (Pestana et al., 2010; Paine et al., 2014b) this represents an improvement over previous attempts of 6.41x and 4.5x, respectively. Additionally, this power-to-weight ratio is 2.1x higher than the most power-dense human muscles (Hunter and Lafontaine, 1992). It should be noted that these results were obtained with the new, lighter version of the UT-SEA shown previously in Figure 2.11.

We also measured the end-to-end efficiency of the high-power behavior, including all losses from the electrical power system (battery) to the mechanical system output. These losses include the power modulation in the servo drive, the energy transduction of the motor, and the mechanical friction in the actuator drivetrain. The input power was measured using power resistors in series with the battery while the output power was measured using forces from the spring and velocity of the actuator output. Efficiency was calculated using the highest power phase of the experiment, where the arm is in the lower position and rapidly moves the 4.5 kg weight upwards. By integrating both the electrical power and the mechanical power (Figure 4.13) we were able to calculate the overall system efficiency during this stage of the maneuver. As is shown, 36.4 joules of energy are fed into the system while 20 joules of energy are felt at the output, resulting in a total efficiency of 54.9%. These are excellent results, considering that this end-to-end efficiency is comparable to the single-stage mechanical efficiency of many other actuator drivetrains (Harmonic Drives, planetary gearboxes, etc.).

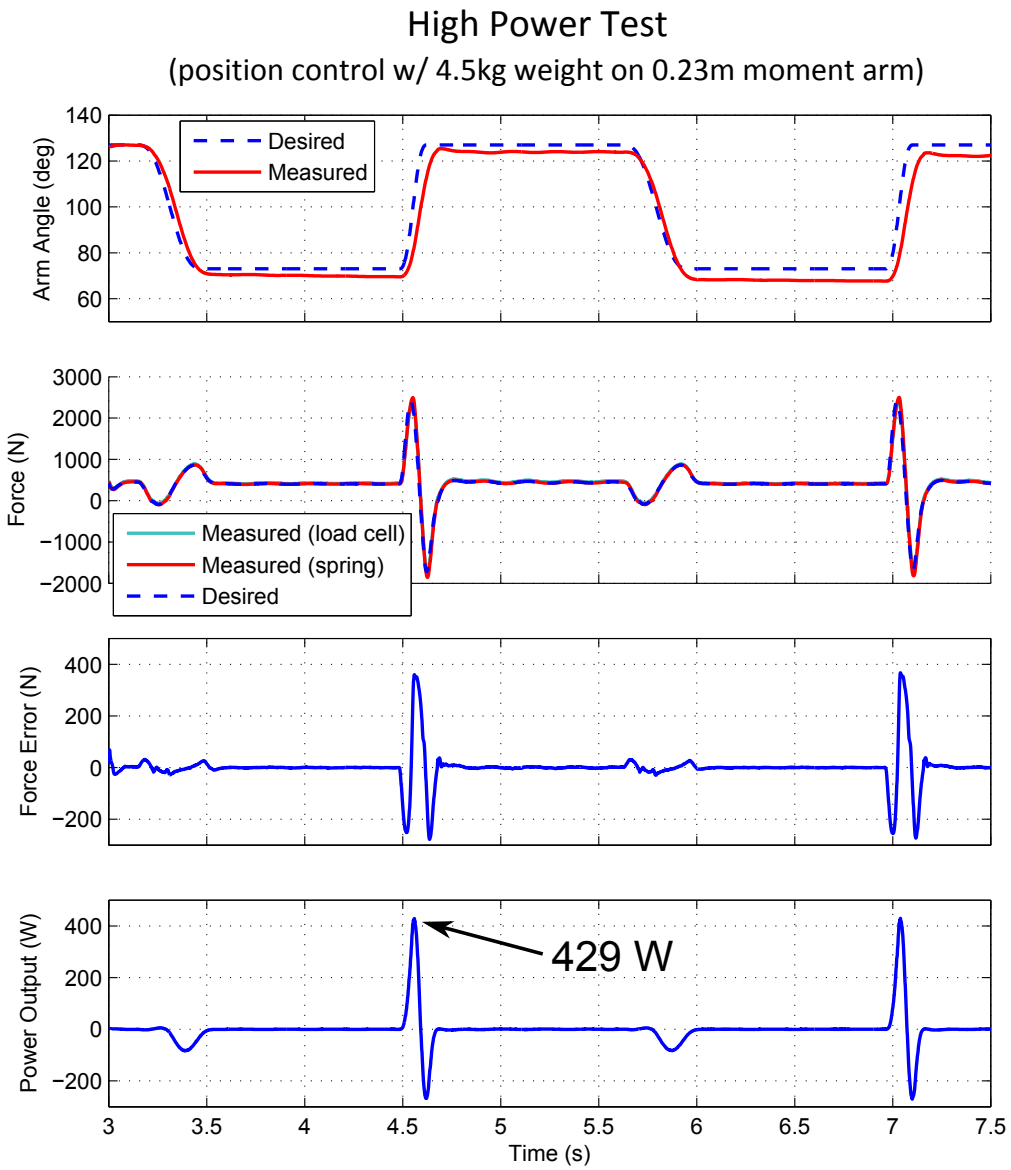

Figure 4.12: **High-power test.** Here, fast motions are tracked with heavy weights. Position and force tracking accuracy is shown in the top two graphs. The third graph shows force error. The bottom graph shows power measured at the output (measured torque times measured velocity). The following variable mappings are used: Desired Arm Angle (θ_d), Arm Angle (θ_a), Desired Force (F_d), Spring Force (F_k). This data was taken with the second generation UT-SEA shown in Figure 2.11.

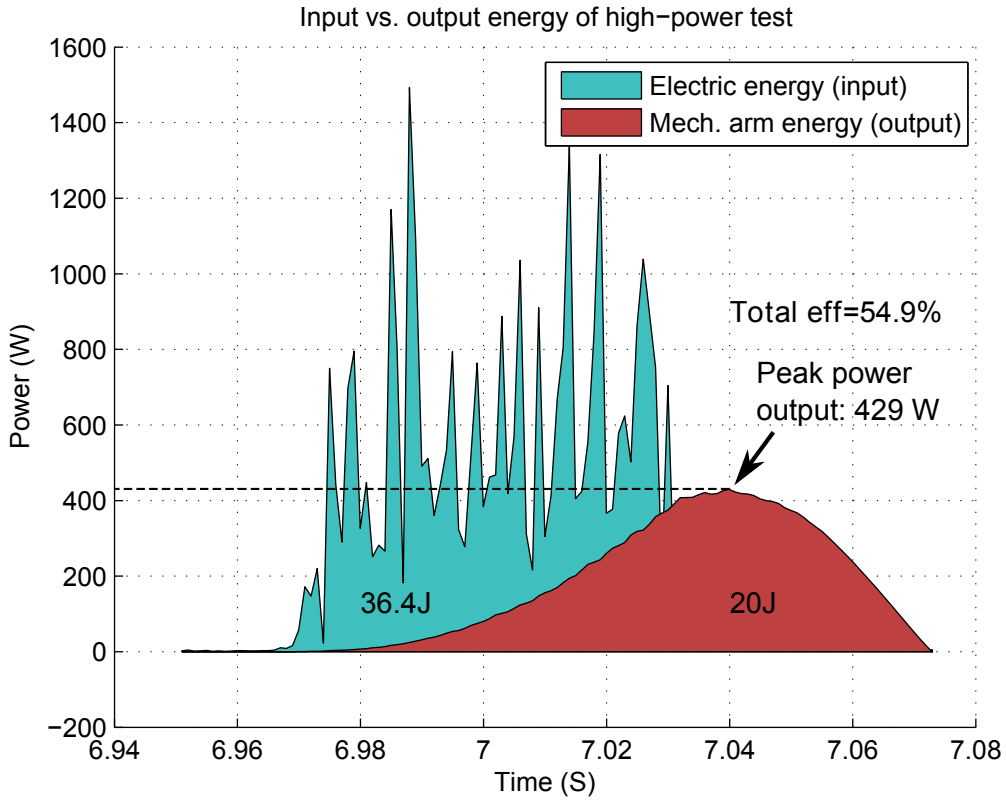

Figure 4.13: Input energy (electrical) versus output energy (mechanical) of the UT-SEA during the high-power tests shown in Figure 4.12. A total of 36.4 joules of energy are provided from the batteries while 20 joules are produced at the actuator output, resulting in a total efficiency of 54.9%.

These experiments were performed on a PC-104 form factor computer from VersaLogic (VL-EPMs-21) running Ubuntu Linux with an RTAI patched kernel to enable real time computation. Data was passed to and from the actuator using analog and quadrature signals which pass through a custom signal conditioning board. Both force and position control were performed at a servo frequency of 1kHz. All continuous time control structures and signal time

derivatives were converted to discrete time using a bilinear (Tustin) transform and were implemented in C.

Chapter 5

Robust Torque Control of SEAs with Unknown Loads

5.1 Introduction

In Chapter 4 we proposed an SEA force control structure and applied it to a model-based position controller to achieve high-performance behaviors in a single actuator testbed. One critical assumption is required for such a position control approach: that the inner force loop can be abstracted as an ideal force source, or more accurately, a low-pass-filtered ideal force source. In this chapter we study this assumption in detail, paying special attention to the effects of unmodeled control plant variations on force tracking accuracy. The results of this chapter have important implications for the application of SEAs in complex multi-DOF robots, which is the focus of Chapter 6.

A definition for an ideal force/torque source (referred to as ideal torque sources hereafter) is provided in (Kong et al., 2009) and states that an ideal torque source must 1) have zero mechanical output impedance and 2) be able to

This chapter contains material from the following publication: (Paine, Mehling, Holley, Radford, Johnson, Fok, and Sentis, 2014a). Mehling, Holley, and Radford contributed to this work by designing the hardware on which the proposed controllers were implemented. Johnson, Fok, and Sentis contributed by interfacing the presented work into the whole-body control framework on the Valkyrie robot.

track control inputs (desired torques) exactly. Implied by the latter condition is the ability to track control inputs up to infinitely large bandwidths. In practice, ideal torque sources are clearly not possible to obtain due both to the limited bandwidth and limited sensing resolution of all physical systems.

However, actuators providing a *near*-ideal torque source provide at least two primary benefits. First, an ideal torque source is the common building block for several multi-body control strategies, including but not limited to: Operational Space Control (Khatib, 1987), Whole-Body Control (Sentis et al., 2013), Virtual Model Control (Pratt et al., 2001), and impedance control (Hogan, 1984). If an ideal torque source is obtained, direct application of these types of multi-body control strategies becomes possible.

Second, abstracting SEAs as torque sources allows higher control layers (i.e. multi-joint controllers) to avoid modeling internal actuator dynamics, and instead only model the effects of the rigid body system. Prior work in this area has demonstrated that effective decoupling of the fast actuator-level dynamics from the slower multi-body dynamics is indeed possible (Ott et al., 2003). However, in this work torque errors remain large (around 50Nm) and thus do not adequately abstract an actuator as an ideal torque source. Decentralized joint-level control approaches have been shown to work well for position controlled robots (Nakao et al., 1987; Godler et al., 1999a). In these approaches subordinate joint-level controllers are co-located at each joint and use feedback to compensate for multi-body dynamics. Interestingly, the authors found that the largest disturbances to affect their joint controllers were caused by varia-

tion in the apparent load inertia due to changing robot pose. (Nakao et al., 1987) found that a Disturbance Observer (DOB) could fully reject the undesirable behavior resulting from these model variations that were unknown to their low level controllers. Decentralized DOBs were first applied to SEAs in (Kong et al., 2009, 2012) and demonstrated accurate torque tracking with no knowledge of actuator load inertia.

Several control methods have previously been proposed for controlling the torque output of an SEA with no knowledge of the load inertia (Vallery et al., 2007; Kong et al., 2009, 2012; Schepelmann et al., 2012). While it is known that a minimum load inertia is required for force control stability (Pratt and Williamson, 1995), little work has been done to identify what this minimum load *is*.

In this chapter we present a method for understanding this minimum load inertia, given the DOB-based controller presented previously in Chapter 4. We define the Disturbance Observer Region of Convergence (DROC), an analysis tool which may be used to study the effect of a DOB applied to a plant with bounded uncertainty. Finally, we show the ability of DOB-based controllers to maintain torque tracking accuracy with varying load inertia in a comparative experimental study.

We begin the study of a joint-level DOB's ability to reject the effects of multi-body dynamics by modeling the effects of variable actuator load inertia on the controller presented in Chapter 4. As previously discussed, the controller in Chapter 4 is tuned for the case where the load inertia is locked

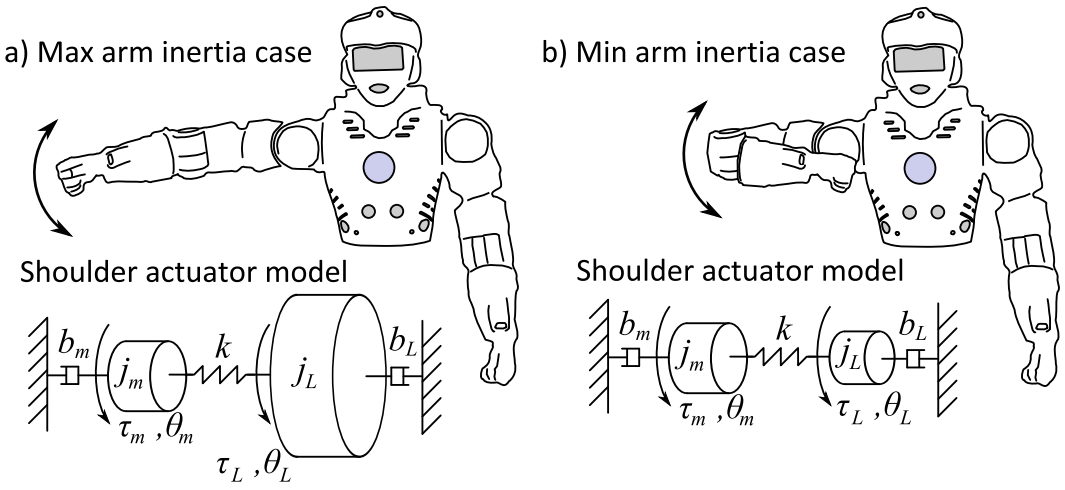

Figure 5.1: An example of a variable load inertia scenario using the abduction/adduction shoulder joint of NASA-JSC's Valkyrie robot. Approximate actuator-level models are shown for each case, a large load inertia corresponding to the outstretched case and a smaller load inertia corresponding to the bent-elbow case.

and is not altered thereafter. That is, no adaptive control techniques or gain scheduling based on robot pose is used. Instead, a DOB is used to attenuate disturbances due to control plant variation and maintain the desired closed-loop dynamics. This feature greatly benefits the tuning procedure for robots with many actuators (see Chapter 6). Each actuator can be separately tuned on a bench with minimal parameter tuning required once the actuator is assembled into a multi-joint system.

5.2 Augmented Control Plant Model

To carry out our analysis, we must augment the locked-output plant model (4.8) by including the effects of load inertia. The modeling approach

we take follows work presented in (Kong et al., 2009) and is extended to map motor torques to spring torques. Unlike Chapter 4, our model here is in the rotary domain, conforming to the rotary actuation hardware on which this work was tested.

A multi-input model for an SEA is shown in Figure 5.1, with τ_m representing motor torque, θ_m motor angle, j_m motor inertia, b_m motor damping, k spring stiffness, τ_L external torque applied at the load, θ_L load angle, j_L load inertia, and b_L load damping. The state equations for the system are

$$j_m\ddot{\theta}_m + b_m\dot{\theta}_m + k(\theta_m - \theta_L) = \tau_m \tag{5.1}$$

$$j_L\ddot{\theta}_L + b_L\dot{\theta}_L + k(\theta_L - \theta_m) = \tau_L. \tag{5.2}$$

Assuming no external torque input, combining (5.1) and (5.2), and representing spring deflection as $\theta_d = \theta_m - \theta_L$, yields a relation ($\alpha$) between motor angle and spring deflection

$$\alpha(s) = \frac{\theta_d(s)}{\theta_m(s)} = \frac{j_L s^2 + b_L s}{j_L s^2 + b_L s + k}. \tag{5.3}$$

Using Hooke's law and (5.3), spring torque (τ_k) is then represented by

$$\tau_k(s) = k\theta_d(s) = k\alpha(s)\theta_m(s). \tag{5.4}$$

Combining (5.1) and (5.4) yields a transfer function from motor torque to motor angle

$$\frac{\theta_m(s)}{\tau_m(s)} = \frac{1}{j_m s^2 + b_m s + \alpha(s)k} \tag{5.5}$$

which, combined with (5.3), provides the transfer function from motor current to spring torque for any load inertia

$$P(s) = \frac{\tau_k(s)}{i(s)} = \frac{\tau_m(s)}{i(s)} \frac{\theta_m(s)}{\tau_m(s)} \frac{\theta_d(s)}{\theta_m(s)} k = \frac{\beta\alpha(s)k}{j_m s^2 + b_m s + \alpha(s)k}. \tag{5.6}$$

Equation (5.6) represents a version of the fixed-output control plant model (4.8) augmented with the effects of finite load inertia.

5.3 Sensitivity Analysis to Variations of Load Inertia

Plotting (5.6) for a wide range of load inertias provides a visualization and intuition towards understanding how an SEA will respond to different load inertias (see Figure 5.2a). When j_L is infinite (bold black line), the system exhibits the familiar second order underdamped response, assuming underdamped system parameters. As j_L decreases, low frequency behavior changes noticeably and the resonant peak moves to higher frequencies. In the extreme case, when $j_L = 0$, the load inertia no longer carries energy and the system becomes the first order mass-damper system defined by j_m and b_m.

Similarly, we may simulate the dynamics of the closed-loop response (P_c from Figure 4.3c) for varying plant load inertias (see Figure 5.2b). In this case, we see that despite significant variation in load inertia, the high frequency

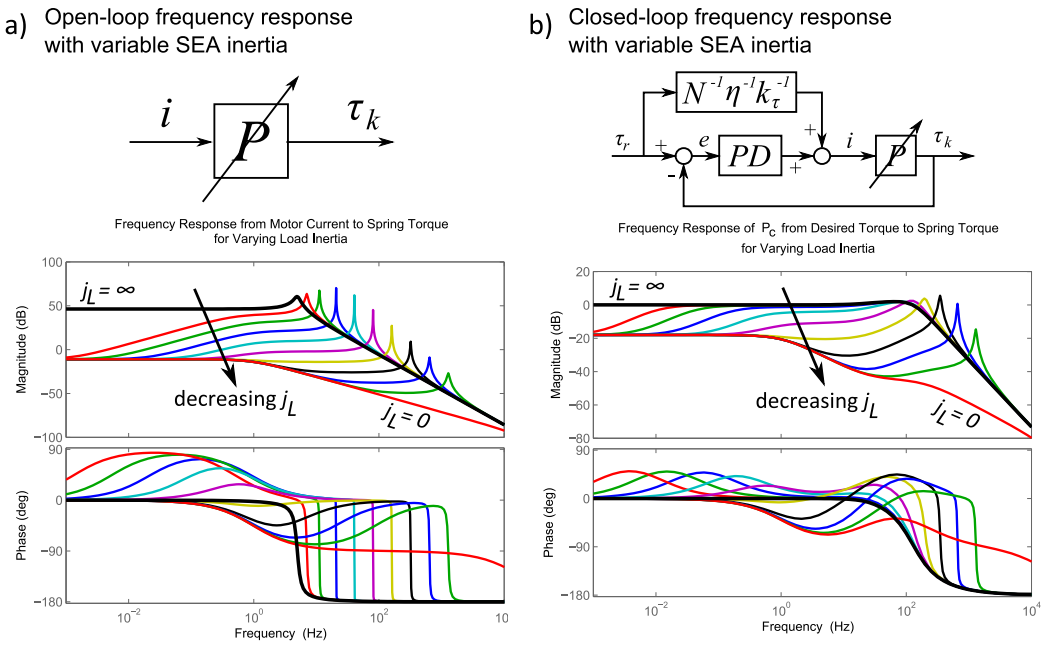

Figure 5.2: **a)** Series of SEA control plants (5.6) for varying load inertia. Reducing load inertia is shown to increase resonant frequency and alter effective level of damping in the response. **b)** Closed-loop (P_c in Figure 4.3c) torque response for a series of SEA plants with decreasing j_L. For larger values of j_L, the bandwidth of the system remains fairly constant while only low frequency behavior changes.

dynamics remain relatively constant. The main variation in the closed-loop response occurs at low frequencies, except for very small values of j_L.

5.4 DOB Disturbance Rejection

Our objective is to leverage the disturbance rejection capability of a DOB to create a controller that is insensitive to variation in load inertia. The result shown in Figure 5.2b demonstrates the qualitative behavior of finite load inertia on the inner PD compensator, P_c. In this section we seek to understand

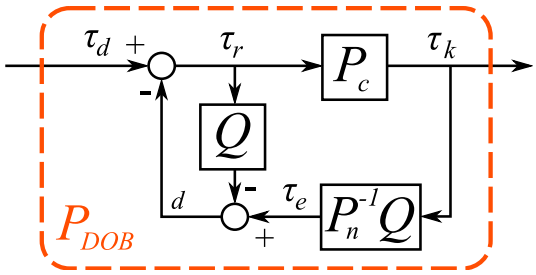

Figure 5.3: Control structure of a disturbance observer. The notations represent: P_c: control plant (a closed-loop plant in this case), P_n: nominal plant model, Q: a filter (typically low-pass) for tuning disturbance rejection bandwidth.

a DOB's disturbance rejection capability and compare this capability with the disturbances shown in Figure 5.2b. If the DOB can successfully remove the disturbances of Figure 5.2b, we will have achieved our goals of creating an inertia-independent torque controller and thus decentralize torque control of Valkyrie's SEAs.

Figure 5.3 shows the structure of a DOB. Given a nominal model (P_n) of a control plant (P_c), the measured plant output (τ_k) is passed through the nominal model inverse to produce an estimate (τ_e) of the control plant input (τ_r). A lumped disturbance value (d) is calculated by subtracting τ_r from τ_e and is compensated for by subtracting d from the desired value (τ_d). A low pass filter (Q) is used both for tuning purposes and to make the inverse nominal plant model proper.

The transfer function for the DOB shown in Figure 5.3 is

$$P_{DOB} = \frac{\tau_k}{\tau_d} = \frac{P_c}{1 + Q(P_c {P_n}^{-1} - 1)}. \tag{5.7}$$

The Q filter in (5.7) affects the range of disturbances that are rejected by the DOB. Consider the case where $Q = 1$. In such a case, deviations from P_n are rejected at all frequencies, resulting in $P_{DOB} = P_n$. In practice, actuator limitations and sensor noise place a limit on the upper bound of the cutoff frequency for Q. Nonetheless, this basic example demonstrates that large disturbances may be removed with sufficiently high Q filter cutoff frequencies.

In the remainder of this section we analyze the relation between the disturbance rejection capability of a DOB and the variations in closed-loop plant behavior of P_c caused by varying load inertia (Figure 5.2b). The result of this analysis is a criteria that guarantees bounded tracking error of the DOB. (Kong and Tomizuka, 2013) provides a similar criteria for determining stability bounds of DOB-based systems with multiplicative uncertainties, but does not give the same error bound guarantee presented here. A difference between these two approaches is that plant variation is unknown in (Kong and Tomizuka, 2013), whereas here we can model variation due to changes in inertial load.

The underlying idea behind our method of characterizing DOB disturbance rejection is as follows (see Figures 5.3 and 5.4). A DOB attenuates deviations of the control plant, P_c, from some nominal model, P_n. The discrepancy between P_n and the DOB-compensated system (P_{DOB}) is the tracking

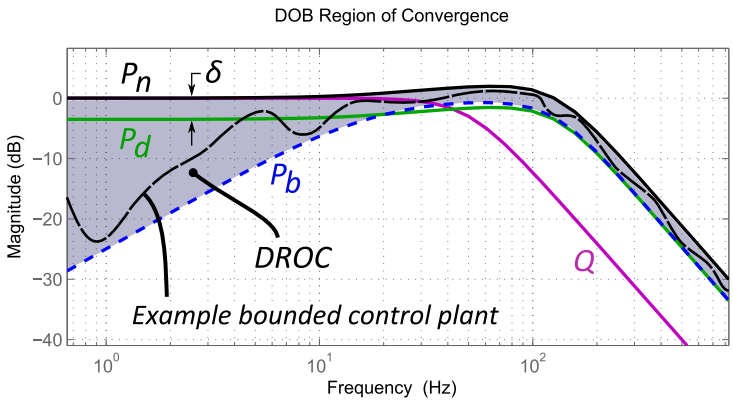

Figure 5.4: Disturbance observer Region Of Convergence (DROC) for a DOB with a given Q filter. A control plant P_c will converge to an error less than or equal to δ if P_c is contained within the shaded area between P_n and P_b.

error of the full closed-loop system. Because the DOB only *attenuates* deviations, tracking error can never be reduced to zero if finite deviations of P_c versus P_n exist. However, if a minimum acceptable tracking error is specified (δ), the worst-case deviation of the control plant from the nominal plant that satisfies the tracking error can be found (P_b). Therefore, any P_c contained within the area bounded by the best case scenario (P_n, zero deviation) and the worst case scenario (P_b, δ deviation) satisfies the minimum tracking error requirement. We refer to the area between P_n and P_b as the Disturbance observer Region Of Convergence (DROC).

A more formal description of this idea is presented here. Let us define an error tolerance (δ) relative to the nominal plant model

$$P_d = P_n(1 - \delta) \tag{5.8}$$

Table 5.1: Variable Inertia Simulation Parameters

Parameter	Value	Units
k_p	1.5	A/Nm
ζ_d	0.9	
f_q	20	Hz
δ	20	%

where P_d is the desired maximum allowable deviation from the nominal plant. Setting $P_{DOB} = P_d$ and solving (5.7) for P_c yields

$$P_b = P_c = \frac{P_d(1-Q)}{1 - P_d Q {P_n}^{-1}}. \tag{5.9}$$

P_b represents the boundary plant transfer function that satisfies the error tolerance, δ. In other words, P_b determines the maximum deviation from the nominal plant model for which the DOB is able to compensate.

5.5 DOB Disturbance Rejection Due to Variable Load Inertia

Analyzing the results shown in Figure 5.2b using our DROC method gauges how a DOB will reject disturbances due to load inertia variation and thus maintain the desired ideal torque source abstraction. In this analysis, we model a Valkyrie shoulder actuator for the two load inertia cases shown in Figure 5.1 using the control parameters shown in Table 5.1. We also simulate a smaller load inertia ($0.16\ Kg \cdot m^2$), which is included to demonstrate a case

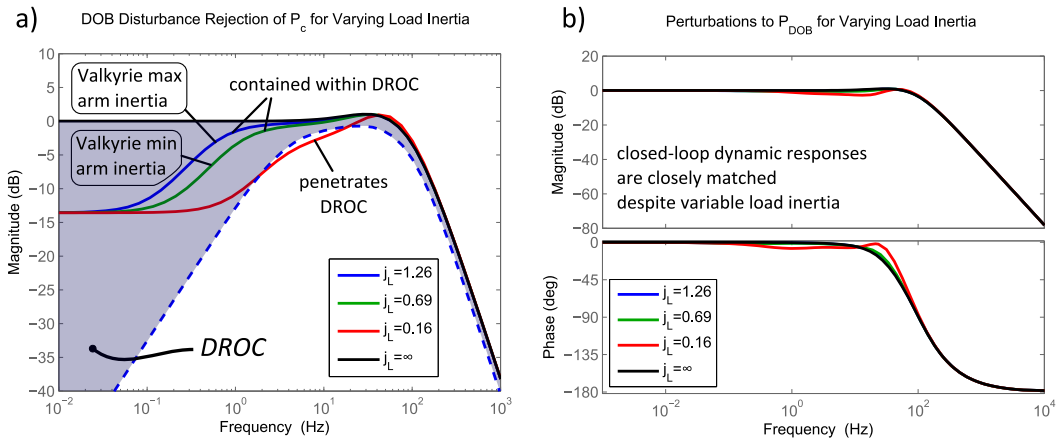

Figure 5.5: These figures demonstrate how a disturbance observer can eliminate the effects of plant deviations caused by changes to actuator load inertia. **a)** Closed-loop frequency response of P_c for different load inertias. The dashed line represents 20% error ($<2\ dB$) compared to the nominal high load inertia model while the shaded area represents the DOB region of convergence. Both the maximum and minimum arm inertia lines are contained within the DROC, meaning a DOB can be used to attenuate the disturbance due to these effects. **b)** Closed-loop frequency response from desired torque (τ_d) to measured spring torque (τ_k) for the full torque controller as shown in Figure 4.3c with a finite actuator load inertia. A series of responses to plants with varying inertias is shown representing the full inertial operating range of the Valkyrie shoulder abduction/adduction actuator ($j_L = 1.26 \rightarrow 0.69\ Kg \cdot m^2$). Note the insensitivity of the system to load inertia variation.

which violates the DROC. Figures 5.5 and 5.6 show these results.

Figure 5.5a illustrates how the closed-loop responses for both values of arm inertia remain within the DROC. Because both responses remain within the DROC, a DOB applied to either closed-loop response is guaranteed to bring the system response to within δ ($<2\ dB$ in this case) of the nominal fixed-output plant model. We can visualize this result in Figure 5.5b where the perturbations to the full system transfer functions due to load inertia

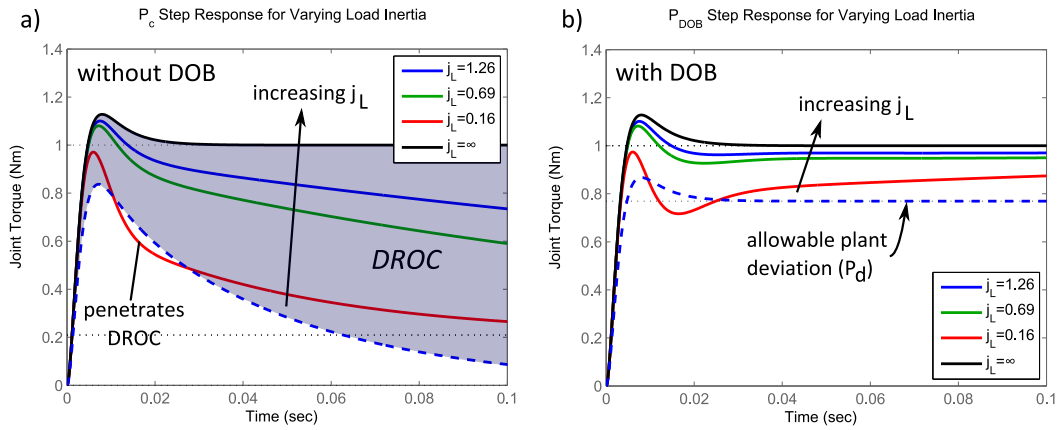

Figure 5.6: Time domain representations of Figures 5.5a and 5.5b. In **a)** the step response of only the inner portion of the proposed force controller (P_c) to various load inertias is shown. Again, for these control parameters, all values of load inertia except $j_L = 0.16 Kg \cdot m^2$ remain within the DROC. In **b)** step responses are shown for the same SEA plant parameters as in **a)**, but a DOB is now applied. The DOB is shown to dramatically improve low-frequency performance, maintaining both Valkyrie inertia values within the allowable performance envelope.

variation are imperceptible for the two arm inertia cases and only become perceptible in the scenario where $j_L = 0.16\ Kg \cdot m^2$.

Figure 5.6 shows the data of Figures 5.5a and 5.5b in the form of time domain step responses. Here, it is again shown how the DOB is able to compensate for large deviations of the control plant, provided these deviations primarily occur at frequencies below the Q filter cutoff frequency.

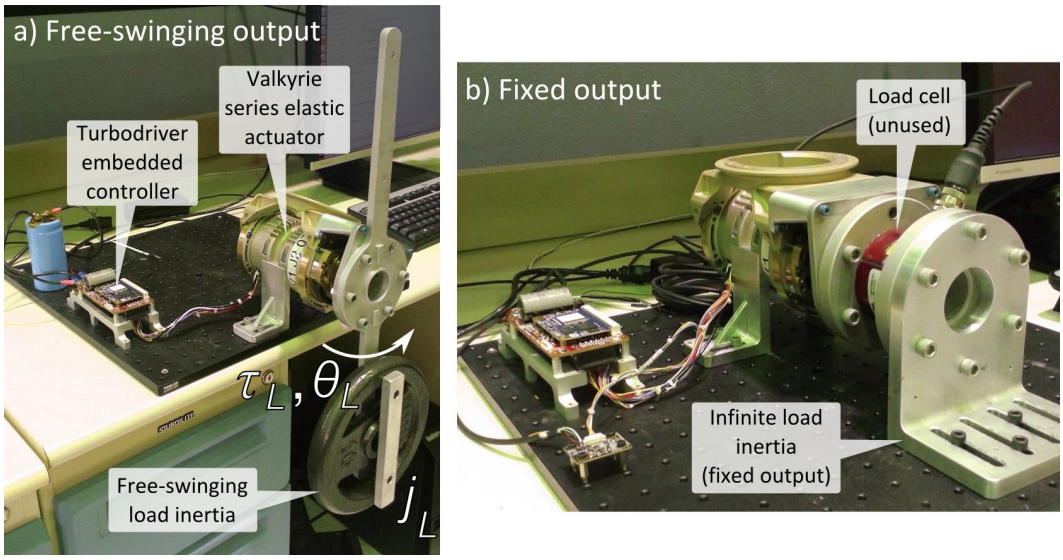

Figure 5.7: **a) Free-Swinging-Output Experimental Setup.** Actuator torque causes load inertia motion. j_L in the free-output test is 0.16 kg·m^2. **b) Fixed-Output Experimental Setup.** The output of the actuator is bolted to the table and therefore actuator torque does not cause load motion.

5.6 Experimental Results

In this section we use Valkyrie's SEAs to demonstrate our DOB-based controller's ability to accurately track torques with no knowledge of load inertia, thus overcoming the primary difficulty of implementing a decentralized torque control architecture. We contrast our DOB-based system with a controller using only the inner PD control loop (P_c). For each controller we consider torque tracking for two different actuator load scenarios (see Figure 5.7). In the first scenario the actuator output is fixed to the table, creating an effective load inertia of $j_L = \infty$. In the second scenario a weight mounted on a moment arm ($j_L = 0.16kg \cdot m^2$) is attached to the actuator output and

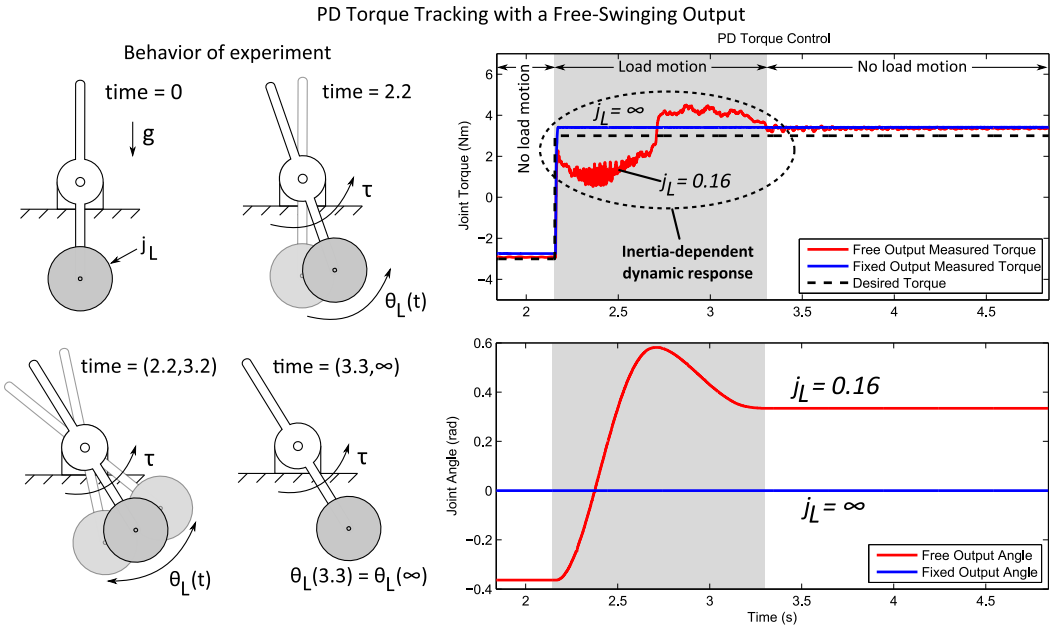

Figure 5.8: Step response of the PD torque controller (P_c in Figure 4.3c) to a torque command for two different load inertias. The dynamic response of the torque controller varies depending on load inertia. In the free-moving case, significant low frequency tracking error occurs and corresponds to motion of the load as indicated by the shaded region of the plot.

is allowed to freely rotate in the presence of gravity. The SEA used in this experiment is Valkyrie's shoulder abduction/adduction actuator.

Figure 5.8 shows the response of a stand-alone PD compensator (P_c in Figure 4.3c) to a 6 Nm step in desired torque. The dynamic response of the PD compensator heavily depends on the amount of load inertia present. In the fixed-output case, the controller successfully suppresses actuator dynamics and tracks the desired signal well. In the free-output case, motion of the load causes significant delay (over one second) before the measured torque signal tracks the desired torque signal closely. Assumptions of an ideal torque source

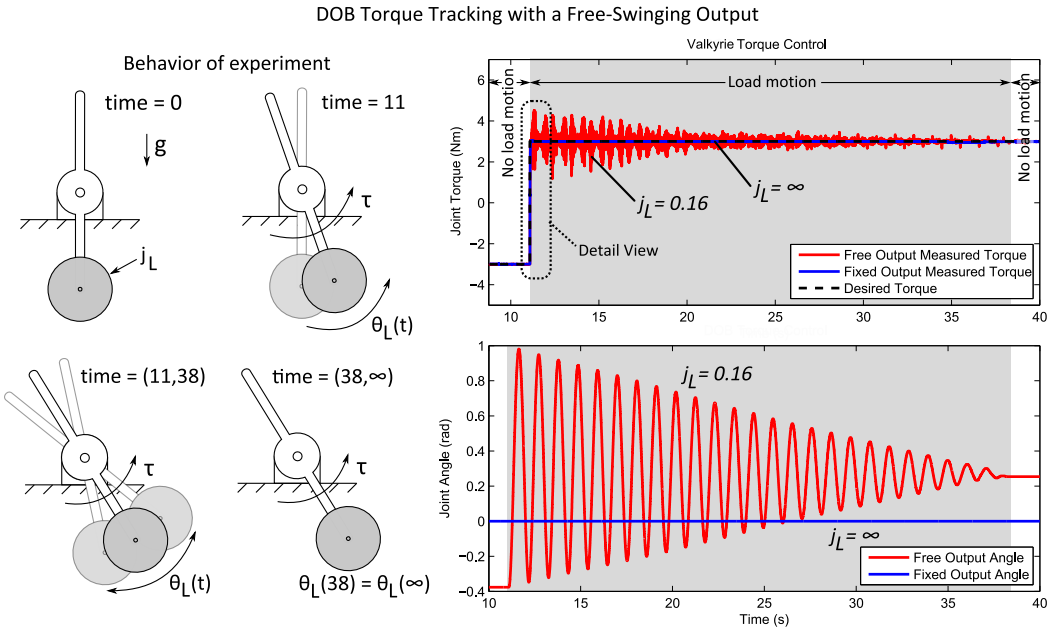

Figure 5.9: Step response of the DOB torque controller used on Valkyrie (Figure 4.3c) to a torque command for two different load inertias. Due to the low friction of the DOB-based controller, the torque step induces high velocity pendular swinging at the actuator output that persists for almost 30 seconds. See Figure 5.10 for a detailed view of the torque response.

using this control method are thus impractical in the presence of low load inertia due to the inertia-dependent dynamic response.

Figure 5.9 shows the response of the DOB-based torque controller (Figure 4.3c, including the DOB) to a 6 Nm step in desired torque. In the free-output case, the torque step causes a low-friction pendular motion that persists for almost 30 seconds before load motion ceases. Note that this is the expected behavior for a pendulum with little friction excited by a torque step. A side effect of this motion is a high frequency disturbance signal caused by torque

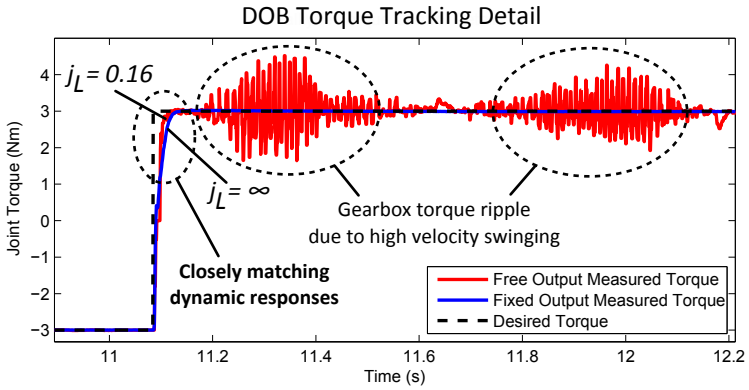

Figure 5.10: Detailed view of the torque response in Figure 5.9. showing dynamic response. The dynamic response of the torque controller is closely matched in the fixed-output versus the free-output case indicating that the DOB controller is robust to output inertia variation and that a nearly ideal torque source is achieved. Because the torque signal is tracked closely, the output inertia swings at high velocity, creating a disturbance signal from harmonic drive torque ripple (a well-known issue with torque signal quality of harmonic drives (Lu et al., 2013; Godler et al., 1999b)). The DOB does not reject this disturbance because its signal frequency content is above the Q filter cutoff frequency.

ripple in the harmonic drive gearbox which the DOB does not fully reject due to its signal frequency content being above the Q filter cutoff frequency. The dynamic response of the controller can be seen more clearly in Figure 5.10. In contrast to the PD compensator, the dynamic response of the DOB-based controller in the free-output case closely matches the response in the fixed-output case. This test indicates that the DOB-based controller produces similar dynamic responses independent of the load inertia configuration.

The results shown in Figure 5.9 confirm that the decentralized joint-level torque controller remains effective in producing an ideal torque response,

or more specifically a low-pass-filtered ideal response, for the load inertia values we have tested. Looking back to the DROC analysis shown in Section 5.5, we note that the load inertia value experimentally tested here is far lower (and thus harder to attenuate with a DOB) than those typically seen by Valkyrie's series elastic joints. As such, the positive results presented in this section give confidence to the torque control abstraction used to decouple Valkyrie's system dynamics.

Chapter 6

High-Performance SEA Implementation on NASA-JSC's Valkyrie Humanoid Robot

In this chapter, the methods and ideas presented in Chapters 2 through 5 are applied to Valkyrie, a 44-DOF humanoid robot. This work was performed in preparation for the DARPA Robotics Challenge (DRC) Trials 2013 as part of a collaboration between the Human Centered Robotics Laboratory at UT-Austin and NASA-JSC. This chapter describes the work performed during this project by the author, namely: 1) assisting with the design of Valkyrie's linear SEAs, 2) designing, implementing and characterizing Valkyrie's joint-level SEA controllers, and 3) system integration and tuning of controllers at the whole-robot level.

6.1 Background

As an entry into the 2013 DRC Trials, NASA-JSC formed a team with several external partners and led the development of Valkyrie (see Figure 6.1).

This chapter contains material from the following publication: (Paine, Mehling, Holley, Radford, Johnson, Fok, and Sentis, 2014a). Mehling, Holley, and Radford contributed to this work by designing the hardware on which the proposed controllers were implemented. Johnson, Fok, and Sentis contributed by interfacing the presented work into the whole-body control framework on the Valkyrie robot.

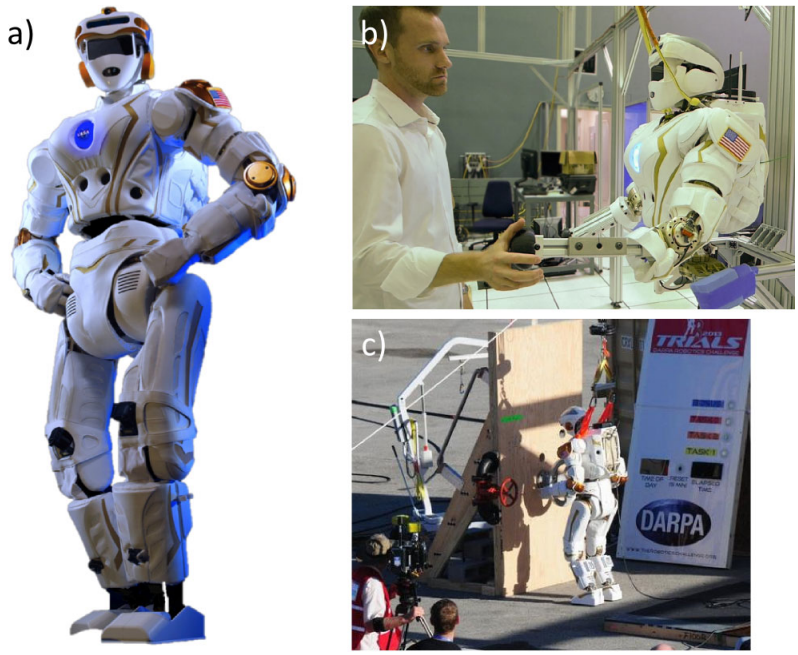

Figure 6.1: **a) NASA-JSC's Valkyrie Humanoid Robot.** Valkyrie has 44 actuated degrees of freedom and a suite of sensors including stereo vision, laser range fingers, sonar depth perception, and tactile feedback. **b) Safe Human Interaction.** Series elastic actuators and compliant torque control enable safe interaction with humans and with unexpected environmental collisions. **c) DRC Competition.** Valkyrie was one of 16 entries in the 2013 DARPA Robotics Challenge Trials. In this image Valkyrie is performing a valve turn task.

Valkyrie was designed to perform tasks required both for responding to disasters, such as the Fukushima Daiichi nuclear disaster in Japan (Nagatani et al., 2013), and for advancing human spaceflight by one day assisting human explorers in extraterrestrial settings such as Mars. While many areas of expertise were required to produce a new humanoid robot in less than 12 months, this chapter focuses on one area, namely the methods used to control Valkyrie's series elastic actuators. In the following discussion, we present our

overall control approach, an outline of our holistic robot control architecture, the actuator controller design and its implementation on Valkyrie.

NASA-JSC has a long history of developing robotic actuators and their controllers, beginning with the Robonaut project and extending to a number of other systems (Ambrose et al., 2000; Bluethmann et al., 2003; Bridgwater et al., 2012; Reiland et al., 2013). Robonaut 2, for example, is the first humanoid robot in space and is currently aboard the International Space Station (Diftler et al., 2011).

For Valkyrie, the actuation control requirements were driven by the needs of the holistic robot control approach and by the mechanical design of its actuators. Because the DRC tasks required movement and manipulation in uncertain environments, compliant control approaches were favored over high-gain rigid control approaches. Consequently, series elastic actuators were chosen as a means of achieving compliant control and protecting both Valkyrie and external objects from unexpected collisions. The naturally low output impedance provided by integrated passive compliance makes series elastic actuators particularly effective at stable interactions with the environment. This stability as well as improvements in shock tolerance, energy storage capabilities, power output, and force sensing are among the many benefits of the SEA architecture widely cited in the literature (Pratt and Williamson, 1995; Pratt et al., 2002; Paluska and Herr, 2006).

6.1.1 Multi-Joint Series Elastic Control Architectures

How to best address compliance within the context of a full body coordinated control architecture is an open question. The primary issue is whether to control the entire compliant robot with one centralized controller or to delegate individual actuator control in a more decentralized scheme (see Figure 6.2). Strengths and weaknesses can certainly be found in both approaches. The decision of how control authority is partitioned between high and low level controllers, if at all, can greatly affect the system and must be considered in any performance analysis.

Early investigations into elastic joint robots treated each motor as an ideal torque source and controlled both actuator and rigid body dynamics with a centralized algorithm (Spong, 1987). Signal latencies and the complexities of modeling the higher order dynamics inherent in a multi-Degree-Of-Freedom (DOF), series elastic robot are difficult to overcome. Nevertheless, this approach is still widely used today and proves quite effective (Ott et al., 2008, 2011; Li et al., 2012). It has been suggested, however, that robots, particularly those designed to mimic the broad versatility of humans and animals, could benefit from a more embedded, collocated control of their actuators (Pratt et al., 2004). Physically, or at least conceptually, decentralizing joint-level control allows for the assumption of a rigid body model at the high level, actuated by ideal joint torque sources at the low level. This approach is taken in (Kim et al., 2012; Sentis et al., 2013; Hutter et al., 2013). The success of such decentralized control architectures directly depends on the performance capa-

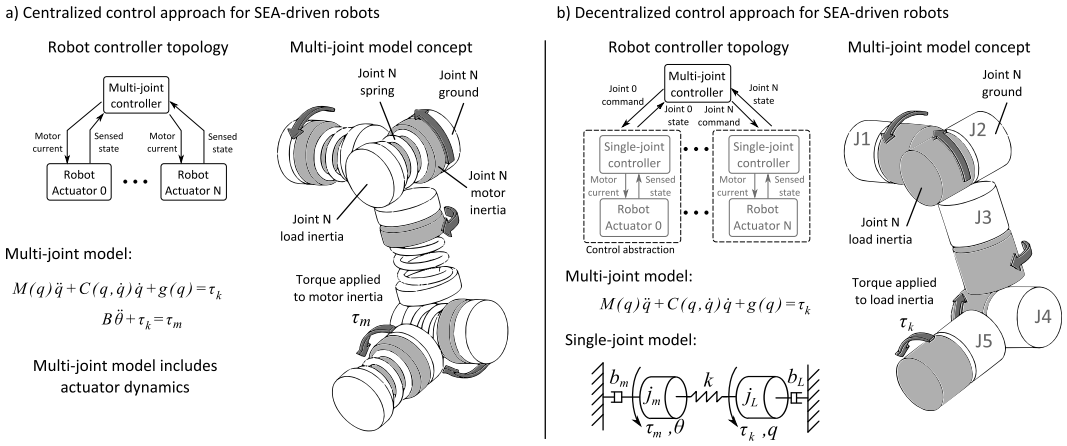

Figure 6.2: Two different approaches for controlling SEA-driven robots. **a) Centralized Approach.** In this approach, a multi-joint controller models both rigid-body dynamics and actuator dynamics. The multi-joint controller takes joint states as an input and outputs motor currents. **b) Decentralized Approach (used on Valkyrie).** In this approach, actuator-level dynamics are abstracted from the multi-joint controller. The multi-joint controller models the robot as rigid bodies actuated by joint torques. The multi-joint controller generates desired joint torques which are then passed to a subordinate set of single-joint controllers. The single-joint controllers model the actuator dynamics and enforce the received torque commands.

bilities of joint-level torque controllers and their effectiveness at suppressing actuator dynamics.

While effective series elastic robots can be designed using either a centralized or decentralized control architecture, one less obvious advantage perhaps tips the scale in favor of the decentralized approach. Hierarchical controller tuning and validation, that is, the ability to test each joint in a multi-DOF system individually, before integration into the whole robot, could speed development time and ease the process of testing higher level functionality

once the full robot is assembled. To achieve this result, torque control of each actuator must truly be decoupled from the rest of the system.

6.2 The Role of Actuator-Level Control in Valkyrie

The selection and design of actuator-level control in a series elastic robot depends on the holistic robot control strategy. In this section, we give a brief overview of our holistic decentralized control approach and provide examples of the broad utility provided by joint-level torque controllers.

6.2.1 A Decentralized Control Approach

Valkyrie uses a decentralized control approach as shown in Figure 6.2b. This approach was chosen over a centralized approach (Figure 6.2a) for several reasons. First, using hierarchical control abstractions reduces the complexity of the dynamic model at the multi-joint level, reducing the computational burden on the central control computer. The benefit of this reduced model complexity manifests itself in the achievement of a 400Hz loop rate for the multi-joint controller on Valkyrie, one which must control 44 degrees of freedom using off-the-shelf computation hardware. Second, a powerful, distributed embedded control element, using NASA-JSCs proprietary 'Robonet' high speed serial bus interface and 'Turbodriver' motor controllers, is co-located at each joint on Valkyrie. Because of this embedded processing capability, our single-joint controllers are able to utilize highly effective dynamic-model-based control schemes. Third, reducing the coupling between central and peripheral sys-

tems decreases overall communication latency due to the reduced number of signals required by the multi-joint controller, meaning they may be updated at a faster rate. Fourth, a decentralized control approach naturally lends itself to an incremental testing methodology, which aides development and debugging.

In the decentralized control approach used on Valkyrie, actuator-level dynamics are abstracted away from the central multi-joint controller. The multi-joint controller models the robot as rigid bodies actuated by joint torques. The outputs of the multi-joint controller are desired joint torques which are then passed to a subordinate set of single-joint controllers. The single-joint controllers model the actuator dynamics and enforce the received torque commands. This approach differs from centralized control approaches, which do not abstract actuator dynamics from the multi-joint model (Spong, 1987; Albu-Schäffer et al., 2007b; Ott et al., 2008).

6.2.2 Primary Actuator Control Modes

During the early stages of the project, the control system design of Valkyrie closely matched the decentralized model shown in Figure 6.2b. The desired joint torque signal created a natural interface between multi-joint and single-joint control domains. Figure 6.2b represents the robot control system from the multi-joint controller's perspective. If we instead consider a single-joint controller's perspective, the same control system could be represented by Figure 6.3a. The single-joint controller receives a desired joint torque (τ_d) and uses a torque feedback controller (discussed in Section 6.3) to enforce this

command.

Later in the project, a higher importance was placed on rejecting disturbances caused by multi-joint model uncertainty. As a result, high joint impedance became a new control target. When high impedance controllers were implemented at the multi-joint level, latencies from single- to multi-joint communication would limit the degree to which impedance could be increased before incurring control loop instability. As a solution to this issue and following work by (Pratt et al., 2004), we created another control paradigm at the single-joint level that we dubbed "impedance control mode", which allowed the position and velocity feedback to occur locally on the single-joint controller (see Figure 6.3b). This control paradigm increases achievable joint impedance due to a significant reduction of control loop latency compared to those incurred in "torque control mode."

In both control modes, a control loop that provides accurate torque control is required. In the next section, we present the design of the torque feedback controller and demonstrate its performance on Valkyrie's actuation hardware.

6.3 Implementation of Torque Feedback Control Using Series Elastic Actuators

As discussed in Section 6.2, torque control plays an important role in the control of Valkyrie. In this section, we introduce Valkyrie's series elastic actuators and quantify their performance using the force controller discussed

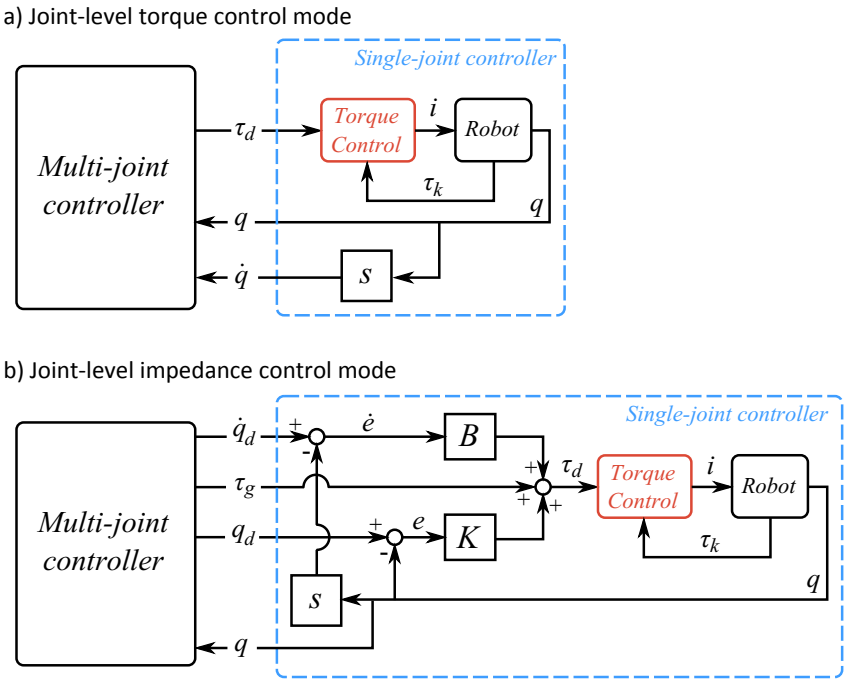

Figure 6.3: Two actuator control modes supported by Valkyrie. **a) Torque Control Mode.** This joint-level control mode matches the decentralized torque abstraction model shown in Figure 6.2b where the multi-joint controller sends desired torques to the single-joint controller. **b) Impedance Control Mode.** Desired joint position (q_d), velocity ($\dot{q}_d$), stiffness (K), damping (B) and gravity compensation torque (τ_g) are sent to the single-joint controller. An impedance control law is enforced at the single-joint level to minimize latency in the control loop, resulting in higher possible stiffness and damping gains. **Note:** Both control modes rely on an accurate torque feedback controller.

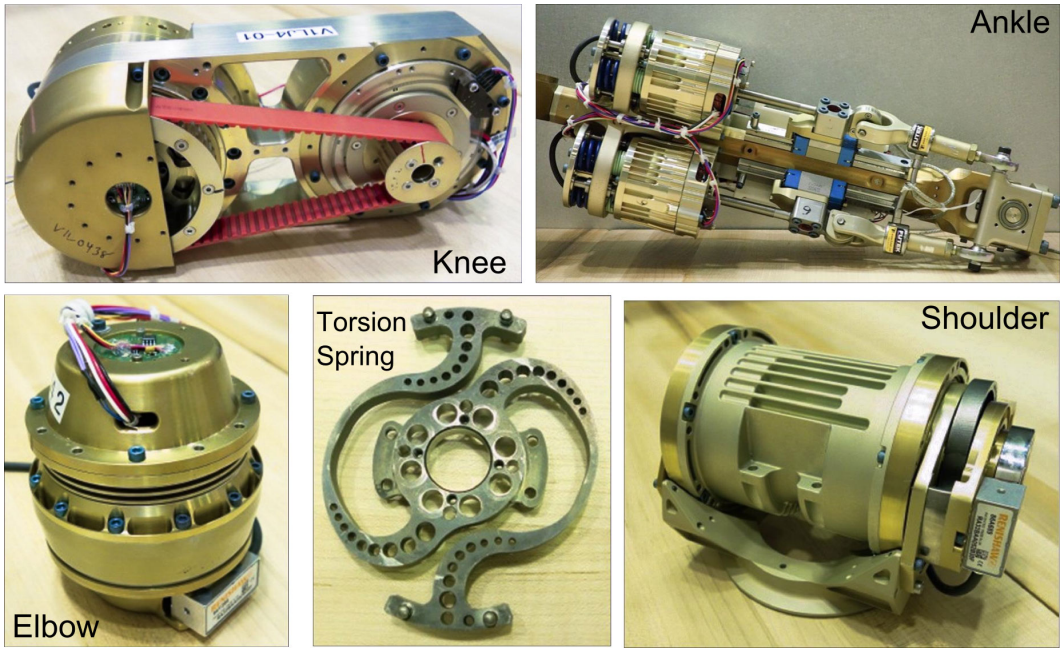

Figure 6.4: Valkyrie's series elastic actuators. For the rotary actuators, spring deflection corresponding to joint torque is sensed using a Renishaw optical sensor. For the linear actuators, load cells are included in addition to the spring deflection sensor to provide redundant force feedback. The loadcells are placed closer to the joint output giving them better dynamic sensing performance but suffer from a higher noise floor than the spring deflection sensor.

in Chapter 4.

Valkyrie's torso, legs, and arms house a number of different SEAs (see Figure 6.4). Valkyrie's torso, arms, and legs nominally use a combination of rotary and linear SEAs driven by brushless DC motors. The rotary SEAs use harmonic drive speed reducers, while custom designed torsion springs act as the compliant element and are used to sense joint torques. Inspired by the compact linear RFSEA topology presented in Chapter 2, Valkyrie's linear SEAs use roller screw drivetrains and commercial-off-the-shelf die springs as

the compliant element. The linear actuators have redundant force feedback in the form of sensed deflection of the die springs and load cells in the actuator output linkage. All of these actuators are driven by NASA's 'Turbodriver' motor controllers which perform the joint-level control discussed below.

Torque control of Valkyrie's SEAs was accomplished using the controller discussed in Chapter 4. Model parameters were identified using the techniques described in Appendix A.1.

6.3.1 Quantifying Torque Control Performance

In the development of Valkyrie, it was essential to establish useful performance metrics so that the effects of controller modifications or topological changes could be quantified. In this section we discuss the performance metrics used to measure torque control performance.

Torque resolution determines an actuator's minimum controllable torque magnitude. An important distinction must be made between *sensed* resolution and *controlled* resolution. Sensed torque resolution is the minimum torque magnitude an actuator is able to *observe* and is purely determined by the type of sensor used. Controlled torque resolution is the minimum torque magnitude that is both observed and *acted upon*. Controlled torque resolution depends on sensed torque resolution as well as an actuator's mechanical properties (friction, for example) and the properties of the torque controller. A controller with small steady state error is required to minimize controlled torque resolution.

The controlled torque resolution of Valkyrie was measured by placing

an actuator in torque control mode and placing a series of loads on the actuator output. The torque from the minimum load that caused motion was determined to be the actuator's controlled torque resolution. For the elbow actuator of Valkyrie, the controlled torque resolution was measured to be 0.002 Nm.

As a second measure of torque tracking performance, we performed a test where the actuator was commanded to track a sinusoidal reference torque. We then created disturbances by manually applying loads to the actuator output. By measuring the maximum torque tracking error, we obtained a rough estimate of nominal torque tracking performance. Figure 6.5 shows these results. As can be seen, torque tracking error remains relatively constant despite the added disturbances. Maximum torque tracking error both with and without disturbances was found to be approximately 1.0 Nm.

Torque bandwidth is an important metric that establishes the maximum signal frequency an actuator is able to accurately track. Figure 6.6a shows the torque tracking bandwidth of Valkyrie's elbow actuator using a fixed-output constraint. Using PD feedback, we extend the bandwidth of the SEA by a factor of 5.3 compared to the passive bandwidth of the SEA (70 Hz compared to 13 Hz). The torque error plot (Figure 6.6b) illustrates the effect of two different controllers on the torque tracking accuracy. The "PD feedback" line demonstrates the maximum performance a proportional controller can achieve before becoming unstable. The "PD+DOB" line, representing the full torque controller implemented in Valkyrie (that of Chapter 4), clearly illustrates torque tracking improvements in the low frequency range. This im-

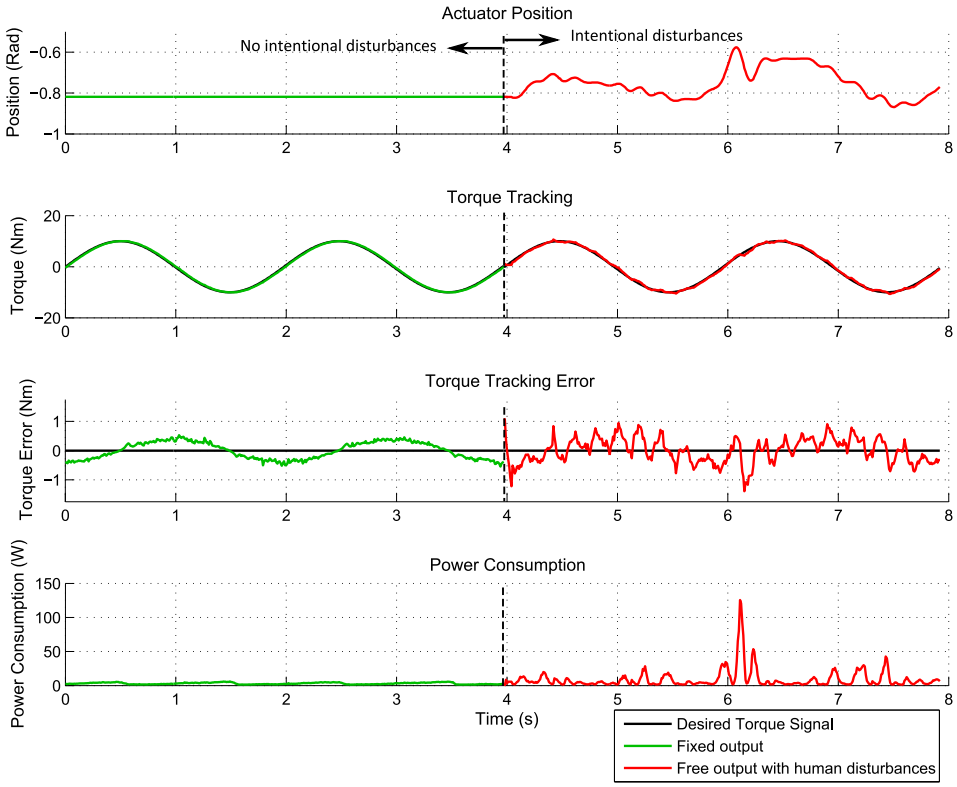

Figure 6.5: Torque tracking and disturbance rejection of Valkyrie's torque control implementation. A 10 Nm sinusoidal reference torque is tracked with intentional disturbances (right half) and without intentional disturbances (left half). Tracking error peaks around 1 Nm.

proved low-frequency torque tracking benefits the controlled torque resolution metric.

As a final metric, we consider the power consumption of a Valkyrie SEA in torque control mode. Based on the previous discussion of feedback gain selection in Section 4.3.1.1, at a first glance, the primary factors in determining feedback gains appear to be performance and stability. A study of actuator

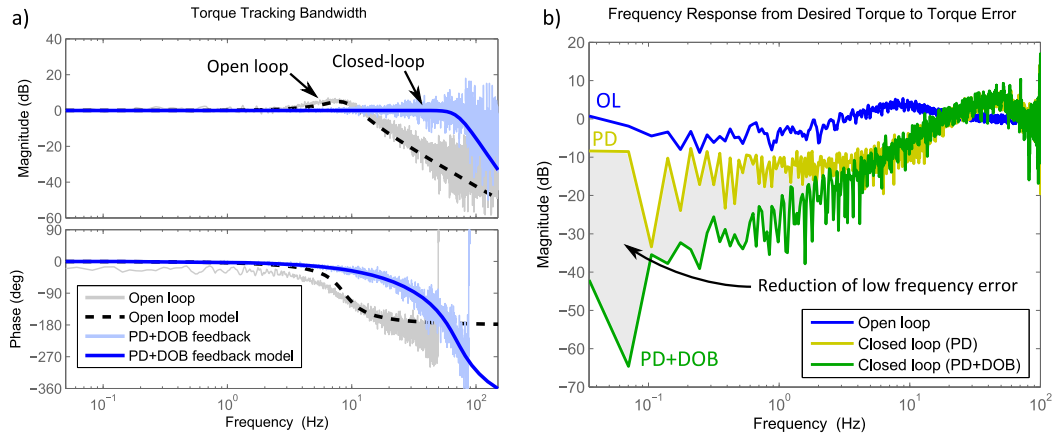

Figure 6.6: **a)** Frequency responses of a Valkyrie elbow actuator with a fixed output. The dashed line represents the open-loop plant response while the solid blue line represents torque tracking with the closed-loop controller from Chapter 4. The closed-loop response significantly increases torque tracking bandwidth and removes the resonant peak seen in the open-loop actuator response. **b)** Torque tracking error versus signal frequency. Adding PD feedback reduces tracking error at fequencies within the actuator bandwidth. The full controller (PD+DOB) further reduces error in low frequencies compared to the PD controller.

energetics quickly demonstrates that a third factor, power consumption, is of critical importance as well. Efficiency is especially important for Valkyrie due to its reliance on battery power.

Figure 6.7 shows the power consumption of Valkyrie's elbow actuator for three different scenarios. In each scenario, a chirp signal is generated as a torque reference. The three scenarios are differentiated by the selection of k_p in the torque feedback loop. As is shown in Figure 6.7, torque tracking bandwidth plays a significant role in determining the efficiency of a torque controlled SEA. Increasing torque tracking bandwidth from 30 Hz to 82 Hz

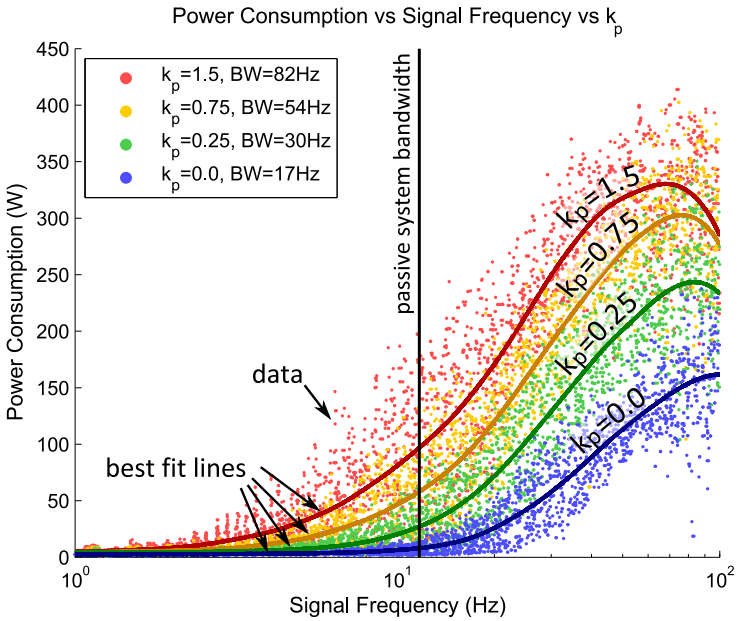

Figure 6.7: Power consumption of a Valkyrie knee actuator. Desired torques were tracked using a PD controller (P_c in Figure 4.3c) for different values of k_p, thus altering the closed-loop system bandwidth (BW). Power was measured using knowledge of motor current and motor velocity. As is shown, while the actuator is physically capable of tracking torque signals with frequency content above the bandwidth of the passive system (11 Hz), doing so requires large power consumption and thus increases motor heating. Therefore, a balance must be established between torque tracking performance and power consumption to avoid overheating the actuator.

results in an average increase of power consumption at the actuator's passive bandwidth frequency by a factor of 3.6, or 260% (97W versus 27W).

These power consumption considerations must be balanced with the needs of the multi-joint controller upstream of the torque controlled SEA. If this centralized controller requires high bandwidth torque tracking, the passive actuator cutoff frequency must be large enough to reduce the discrepancy

Table 6.1: Valkyrie Torque Control Performance Metrics

Metric	Value	Units
Controllable torque resolution	0.002	Nm
Nominal torque tracking error	1.0	Nm
Maximum torque bandwidth	70	Hz

between passive and active torque bandwidth, as this discrepancy determines peak power consumption (Bèlanger, 1995). Because the spring of an SEA heavily influences the passive actuator frequency, its selection should consider these energy-related issues. A large amount of work remains to better understand these relations, but the trends observed here are useful in directing future study.

Table 6.1 summarizes the torque control performance metrics of Valkyrie's elbow actuator discussed in this section. Valkyrie's other series elastic joints share a similar performance profile.

6.4 Experimental Results on Valkyrie

In this section, we put the disturbance rejection capability of the DOB to the test in the form of two experiments using Valkyrie's arm, a serial chain of four SEAs. From Chapter 5 we have learned that the DOB is capable of rejecting disturbances at frequencies below its Q filter cutoff frequency. In this section, the complexity of disturbances experienced by the DOB are increased beyond the variable inertia case considered in Chapter 5. Now, the controller

Table 6.2: Valkyrie Control Parameters used in Experiments 1 and 2

Parameter	**J1**	**J2**	**J3**	**J4**	**Units**	**Notes**
k_p	1.5	1.5	1.5	1.5	A/Nm	
ζ_d	0.9	0.9	0.9	0.9		
f_q	50	20	50	50	Hz	
K	100	100	50	50	Nm/rad	Experiment 2 only
B	10	10	5	5	Nm s/rad	Experiment 2 only

must attenuate disturbances from the full dynamics of a serial chain of four SEAs along with the disturbances introduced by control action of each of their motors.

The goal of the two experiments shown here is to assess the torque tracking capability of the distributed joint-level torque controllers, which possess no knowledge of neighboring joints. The only models used in these tests are the nominal high-output-impedance models used by the DOB as discussed in Section 6.3. The control parameters used in these experiments are listed in Table 6.2.

In the first experiment, joints one through four[1] are placed into torque control mode and commanded to track a constant torque value. A human then holds Valkyrie's forearm and applies motions such that all four joints move (see Figure 6.8). Figure 6.9 shows the data from this experiment.

In the second experiment, joints one through four are again placed into

[1]Refer back to Figure 6.2b for joint naming conventions.

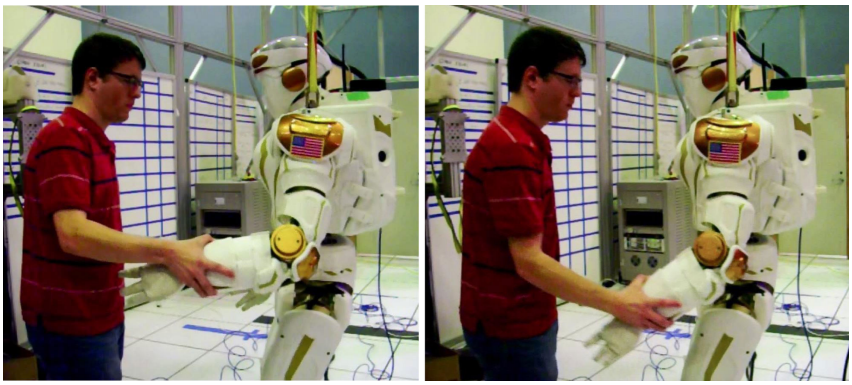

Figure 6.8: **Experiment 1: Human Interaction.** Valkyrie arm joints J1 through J4 are placed in torque control mode and are actively tracking a constant desired torque signal. A human applies motions to the arm causing all four joints to move.

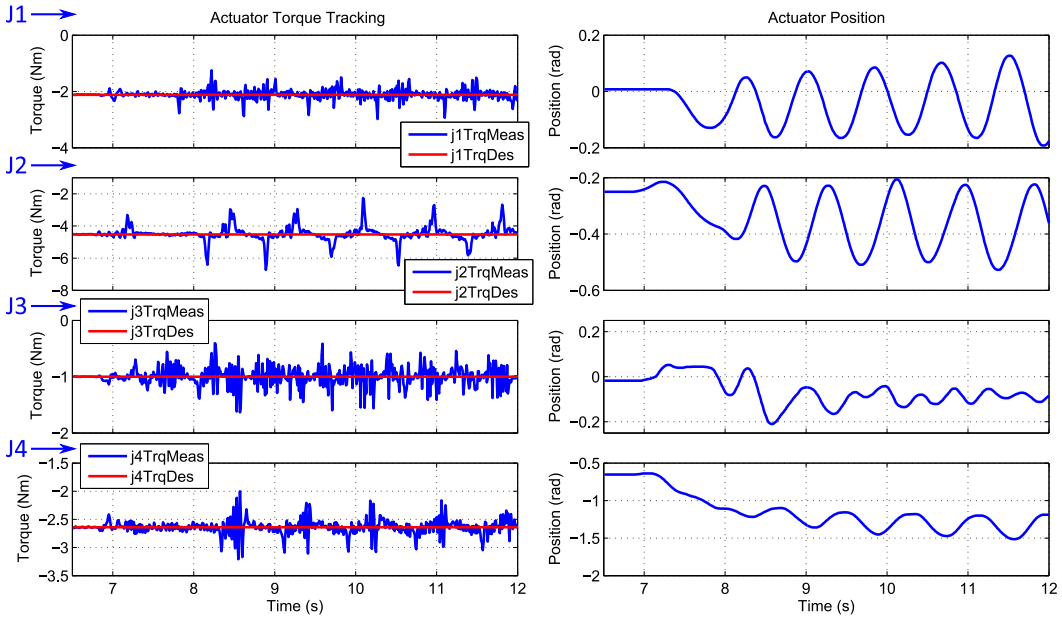

Figure 6.9: **Data from Experiment 1: Human Interaction.** Joint motion is caused by the human who is interacting with Valkyrie's arm. The peak tracking error for each joint is: J1: 0.85 Nm, J2: 2.2 Nm, J3: 0.63 Nm, J4: 0.64 Nm.

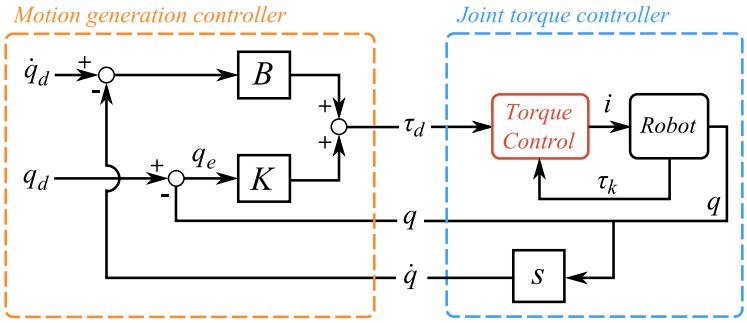

Figure 6.10: Motion generation controller used to test joint torque control in Experiment 2. The parameter K represents joint stiffness while B represents joint damping.

torque control mode. In this experiment, however, desired joint torques are generated by the controller shown in Figure 6.10. This is a simple proportional-derivative controller which creates an apparent joint stiffness (K) and damping (B) based on a desired joint position (q_d) and velocity ($\dot{q}_d$). We use this controller to generate motion in joints one through four; we are less concerned with the position tracking error (q_e) in this particular experiment.

Figure 6.11 shows the basic arm motion generated during Experiment 2. Correspondingly, Figure 6.12 shows the data from Experiment 2.

These two experiments demonstrate that the disturbance attenuation properties of the joint-torque controller are able to suppress the dynamics of neighboring joints, producing accurate torque tracking. However, by visual inspection of Figures 6.9 and 6.12 it is difficult to evaluate the achieved tracking performance based on torque error magnitude alone. A helpful metric in this scenario is the maximum torque tracking error relative to the rated maximum joint torque. Using this metric, we find that the torque error relative to the

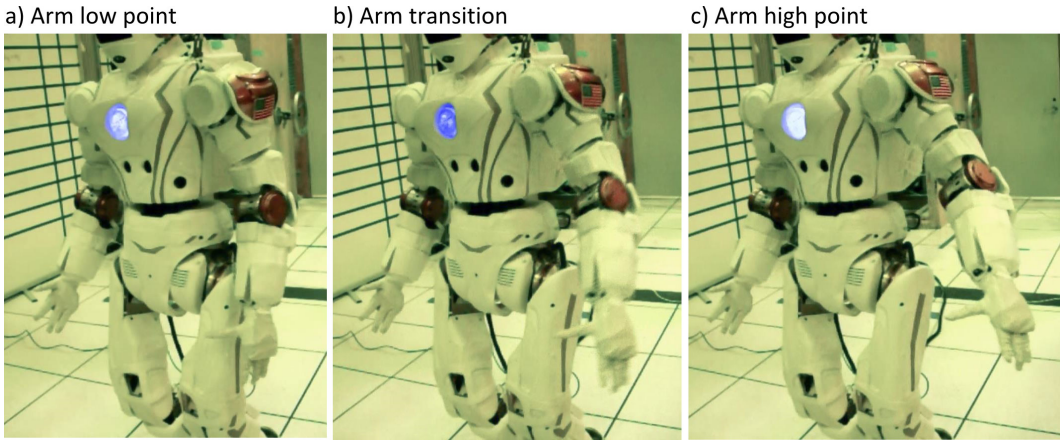

Figure 6.11: **Experiment 2: Coordinated Motion.** Valkyrie arm joints J1 through J4 are placed in torque control mode and are supplied desired torque signals according to the controller shown in Figure 6.10. As a result, each joint roughly tracks a sinusoidal position.

rated joint torque of each joint in Experiment 1 is: J1: 0.4%, J2: 0.76%, J3: 0.97%, and J4: 0.98%. For Experiment 2 these numbers are: J1: 1.16%, J2: 0.87%, J3: 1.38%, and J4: 1%. When considering that the accuracy of most *sensors* fair no better than this in terms of full-scale accuracy[2], we consider our results to be quite strong. These results also fare well in this metric compared to leading research in the field of accurate torque control of SEAs. To date, the best results the authors are aware of are tracking accuracies of 1.6% of full scale (Kong et al., 2009) and 15% of full scale (Kong et al., 2012).

The proposed control methodology has also been used by Valkyrie to perform useful tasks, both in the lab (Figure 6.13) and in the field during the DRC Trials 2013 (Figure 6.14).

[2]http://www.futek.com/files/pdf/Product%20Drawings/lcm300.pdf

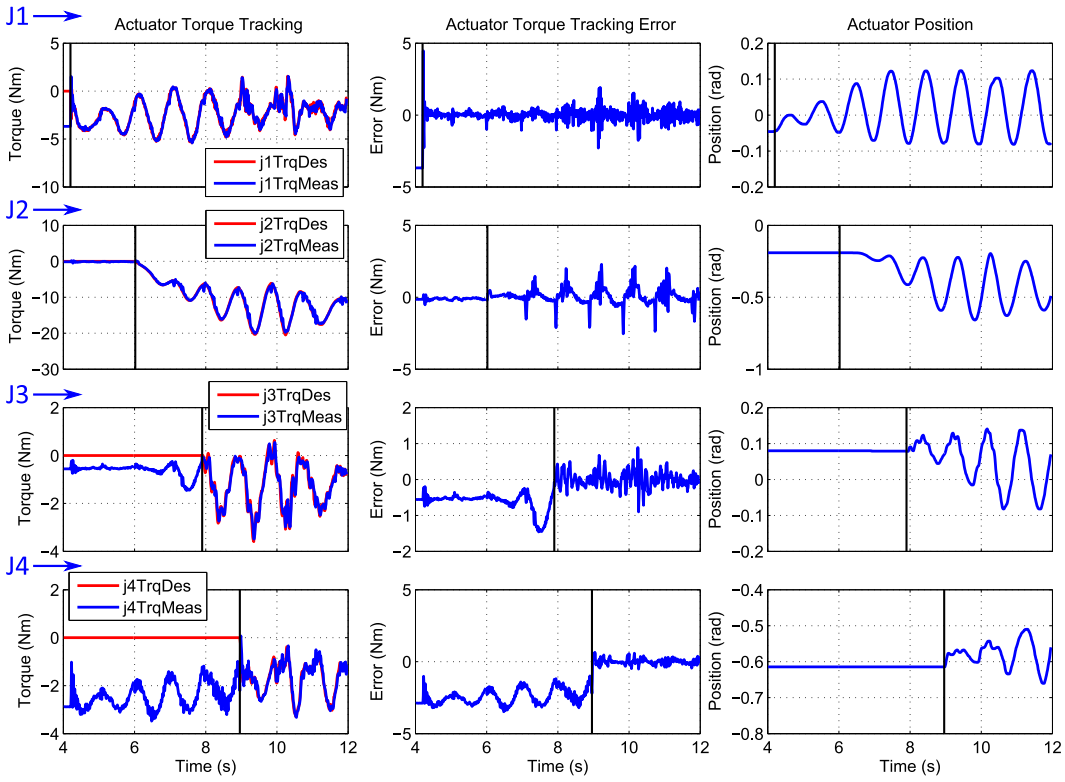

Figure 6.12: **Data from Experiment 2: Coordinated Motion.** The purpose of this experiment is to demonstrate the torque tracking accuracy of a serial chain of SEAs controlled using the methods presented in this thesis. The vertical black lines in the figure represent the time where each joint begins tracking the desired torque signal. The peak tracking error for each joint is: J1: 2.2 Nm, J2: 2.5 Nm, J3: 0.9 Nm, J4: 0.65 Nm.

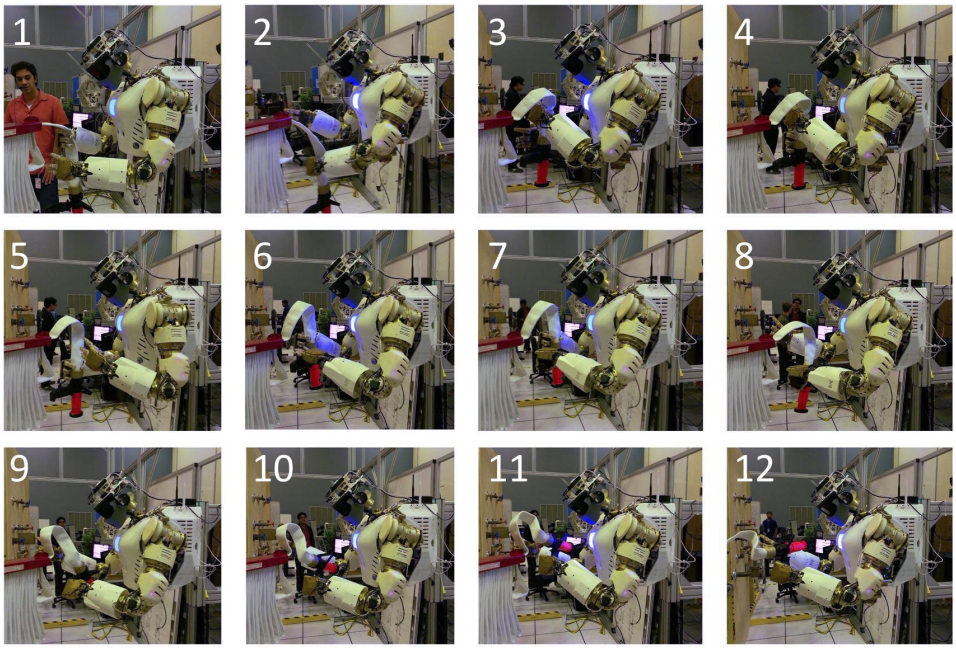

Figure 6.13: **Sequence of Semi-Autonomous Manipulation of a Fire Hose.** This sequence corresponds to a training exercise a few months before the DRC Trials in December 2013. The three torso actuators and eight arm actuators use the control policies presented in this thesis. Wrist pitch, roll, and yaw degrees of freedom use open-loop current control to implement the torque commands sent by the multi-joint controller.

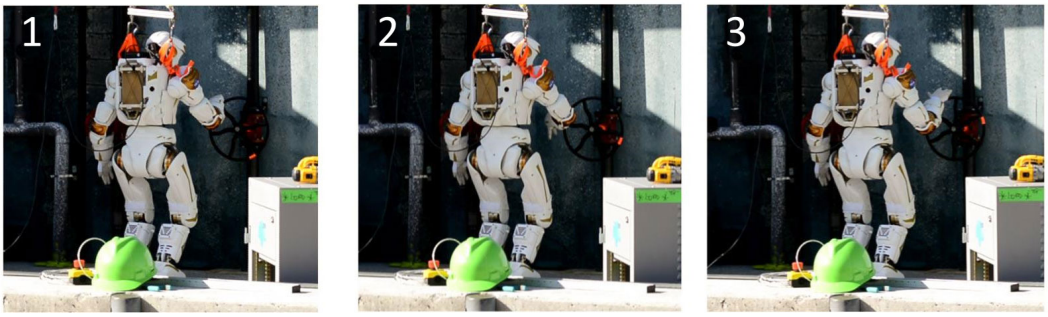

Figure 6.14: **Valve Turn at the DRC Trials.** Valkyrie using the control methods described in this thesis to turn a valve during the DRC Trials, December 2013 in Miami, Florida, USA.

Chapter 7

Maximizing Actuator Impedance with Delayed and Filtered Feedback

In previous chapters we have focused on force control techniques. Such techniques may be used to control actuator position through the use of inverse dynamics as discussed in Chapter 4. In contrast to those methods, this chapter focuses on high-accuracy position control, predominately through the use of feedback control effort. The motivation for this study came from our participation in the DRC Trials where several joints in NASA-JSC's Valkyrie robot required high mechanical output impedance to maximize positioning accuracy in the presence of unmodeled loads.

It is generally understood that increasing feedback effort reduces control loop stability, but what determines the maximum permissible values? This question is difficult to answer because it is largely influenced by the complexities of hardware implementation. Therefore, procedural methods for determining feedback gains are often cast aside in favor of ad-hoc guess-and-check methods. The ideas in this chapter were conceived to better address these

This chapter contains material from the following publication: (Paine and Sentis, 2014). Sentis contributed to this work by helping formulate the problem given his experience working with the Valkyrie robot.

fundamental questions.

This study was prompted by a collection of Valkyrie's joints which were constructed without series elastic components. Therefore, it should be noted that this is the only chapter in this thesis where a rigid actuator model is used. However, future work may leverage similar ideas for use with series elastic actuators.

7.1 Introduction

High mechanical output impedance is a common actuation requirement in robotics and related fields. Most robotic positioning systems, such as robotic arms, require high output impedance to minimize positioning error in the presence of unmodeled payloads (Kawamura et al., 1988). Other applications, often in the field of haptics, benefit from actuators capable of exhibiting a wide *range* of impedances (Colgate and Brown, 1994). Common to both of these fields are 1) the importance of identifying the upper bound of achievable actuator impedance and 2) their basic control plant model, composed of a force acting on a mass-damper.

The control laws used for haptics and robotic positioning systems share commonalities as well. Haptic devices typically employ impedance control (Hogan, 1985) to define a dynamic relationship between external (human) forces acting on the haptic display and the displacement of the haptic display in response to these forces. Robotic positioning systems often rely on the high mechanical output impedance produced by high-gain PID controllers to

maintain tracking accuracy in the presence of unmodeled payloads. If we consider cases of PID controllers where the integral term is set to zero, the action of the resulting PD control law depends on the error proportion (P) and its derivative (D). Perhaps the most common impedance model used in haptic displays is that of a virtual spring-damper (K-B) system, which similarly depends on error proportion (K) and its derivative (B). Therefore the analysis of the maximum impedance for a mass-damper plant controlled by a PD-type controller certainly has wide application. Any such analysis should aim to answer two questions: 1) What are the control parameters (i.e. P and D gains) given a set of system parameters (plant properties, feedback delay, etc.) which produce the maximum output impedance? and 2) How is maximum output impedance defined? The focus of this chapter is on providing direct and applicable answers to these questions.

Numerous methods for selecting controller parameters of PID-type systems exist in literature. Early work utilized knowledge of the system transient response to develop simple equations for determining PID controller values (Ziegler and Nichols, 1942). A number of methods have been developed since, most of which produce PI or PID controller parameters using either numerical or graphical techniques, or by relying on simplifying approximations (Åström, 1993; Lee, 2004; Poulin et al., 1996; Ho et al., 1998). Effects of nonlinearities due to actuator saturation on selection of PD gains have been studied in (Lawrence, 1989; Sourlas et al., 1994; Goldfarb and Sirithanapipat, 1999). More recent approaches apply to a wide range of plants and handle varying lev-

els of controller complexity, including those with filtered D terms, but require inclusion of an integral controller term (Yaniv and Nagurka, 2004; Åström et al., 1998). Optimization based methods may also be applied to tune PID controllers for particular performance criteria, but require the use of an optimization framework (Li et al., 2007; Wang and Li, 2011).

It is important to understand what assumptions or criteria apply to the methods used for selecting controller parameters. Perhaps the most conservative criterion is system passivity (Colgate and Schenkel, 1994). Despite producing conservative control parameters, this approach has been widely studied based on the simplicity and elegance of its main result: maximum feedback gains are found using a single closed-form expression. At the other end of the spectrum, several studies are based on a true or false stability criterion, some of which are empirically determined (Lawrence, 1988; Colgate and Brown, 1994; An and Kwon, 2004; Mehling et al., 2005; Hulin et al., 2006). Their observations are useful for system level mechanical design but do not provide quantitative measures needed for performance-critical controller design. Between these two extremes lie a wealth of methods which produce control parameters given a stability *margin* (Franklin et al., 1986; Ogata, 1990; Poulin et al., 1996; Ho et al., 1998; Suchomski, 2001; Yaniv and Nagurka, 2004; Li, 2013). Other criteria serving a similar purpose include the use of integrated absolute error, maximum sensitivity, Lyapunov-based stability, and others (Lee and Teng, 2003; Diolaiti et al., 2006; Li et al., 2007; Wang and Li, 2011).

While prior work provides partial solutions to the problem of deter-

mining the maximum control parameters, a complete and readily applicable answer has not been given. As a result, the tuning of the class of systems considered in this chapter is often performed using trial-and-error techniques, requiring many hours of tedious work and without any guarantee of finding the optimal parameters. A successful solution to this problem must satisfy a number of requirements. First, it must be simple to apply to enable adoption by a wide audience with varying levels of control systems background, as do the approaches in (Ziegler and Nichols, 1942; Ho et al., 1998; Colgate and Schenkel, 1994). Second, it must produce parameters which are not overly conservative, nor parameters which are unstable, so that minimal manual tuning is required, as do the approaches in (Franklin et al., 1986; Ogata, 1990; Poulin et al., 1996; Ho et al., 1998; Suchomski, 2001; Yaniv and Nagurka, 2004; Li, 2013; Li et al., 2007; Wang and Li, 2011). Third, it must accommodate the full complexity of real-world systems including time delay and filtering of the derivative term, as do the approaches in (Colgate and Schenkel, 1994; Colgate and Brown, 1994; An and Kwon, 2004; Yaniv and Nagurka, 2004; Diolaiti et al., 2006; Li et al., 2007; Wang and Li, 2011). None of the methods in existing literature meet all three of these requirements.

Is it possible to have a controller tuning method that is simple, meets performance and stability criteria, and captures the full complexity of real-world systems? We show that such a method is indeed possible if we reduce the scope of the problem by 1) considering only a single plant model, 2) considering only critically damped pairs of control parameters, and 3) limiting the range

of system parameter values to a finite set. Yet, as previously stated, many systems in practice fit these requirements. The goal of this chapter is therefore to provide a simple and effective procedure which applies to a large group of systems (robotic positioning systems and haptic displays, for example) but does not apply to *all* systems. Using these assumptions, we show that it is indeed possible to directly calculate maximum impedance control parameters if system properties are known.

Our method relies on fitting continuous curves to sets of data derived from simulations of the full control system including time delay and derivative filtering. We use parametric searches to locate critically damped pairs of stiffness and damping parameters meeting a phase margin criteria. As a result, we present a single closed-form equation which accurately maps system parameters to maximum impedance control parameters. Experimental results with a purpose-built actuation testbed demonstrate the efficacy of our method using two different actuators.

7.2 Problem Statement

Many conventional actuators may be modeled as a force (F) acting on a mass-damper (m-b) (see Figure 7.1b). The relationship between external forces (F_{ext}) and actuator displacement (X) may be considered to be a form of the output impedance (Z) of the system[1]. External disturbances may be applied

[1]While the mechanical impedance is typically defined as $Z = F/V$ (see (Hogan, 1985)), in this work we use the form of mechanical impedance defined by the relationship $Z = F/X$,

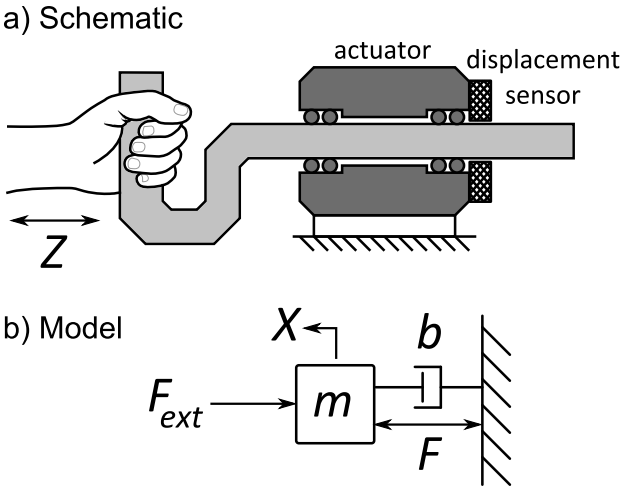

Figure 7.1: **a)** An actuator impedance interaction with a human. **b)** Model of actuator including forces from external sources.

by a human as depicted in Figure 7.1a (originally presented in (Colgate and Brown, 1994)), or from other sources such as unmodeled actuator loads. The system impedance is a function of both passive system properties ($Z_{passive}$) and impedance caused as a result of forces supplied by the actuator (Z_{active}):

$$\begin{aligned} Z(s) = \frac{F_{ext}(s)}{X(s)} &= Z_{passive}(s) + Z_{active} \\ &= -ms^2 - bs + Z_{active}. \end{aligned} \tag{7.1}$$

If a feedback control law is defined producing actuator forces as a function of proportional position error (K) and derivative position error (B),

following the convention used in (Pratt and Williamson, 1995).

$$F(s) = K(X_d - X) + B(sX_d - sX) \tag{7.2}$$

then, assuming the desired position (X_d) is set to zero, the system impedance becomes

$$Z(s) = Z_{passive}(s) + Z_{active}(s) = -ms^2 - (b + B)s - K. \tag{7.3}$$

Notice that the system impedance (7.3) matches that of a mass-spring-damper system, where the mass is that of the actuator inertia, the damping is the combined passive and active damping, and the stiffness is produced purely via feedback control effort. This "user defined" impedance is the central idea of impedance control (Hogan, 1985).

Suppose, as an example, a controller design criteria is given as a maximum tolerable position error (X_{err}) for some steady state disturbance force. In this case, the steady state impedance, $Z(s = 0) = K$, can be chosen as:

$$K = \frac{F_{ext}}{X_{err}}. \tag{7.4}$$

Then, B may be chosen given some desired damping ratio (ζ_d):

$$B = 2\zeta_d\sqrt{mK} - b. \tag{7.5}$$

For this example, this procedure will work for arbitrarily large values of K and B. However, the same is not true for real-world systems. The impedance

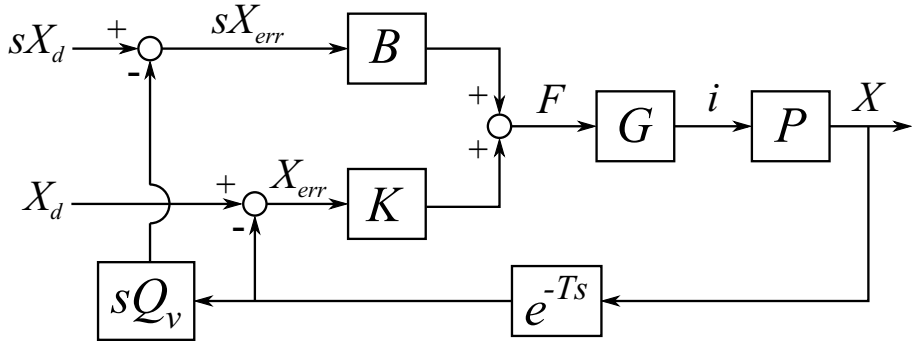

Figure 7.2: A spring-damper (K-B) impedance control model with delay (e^{-Ts}) and velocity filtering (Q_v). No force feedback is used, desired forces are simply translated into desired currents.

model in (7.3) does not accurately portray the effects of control loop delay and derivative signal filtering which together limit achievable system impedance (Colgate and Brown, 1994; Diolaiti et al., 2006).

If we create a model which includes the effects of control loop delay and derivative signal filtering, we obtain the controller depicted in Figure 7.2. In this diagram, the actuator model (P) represents the dynamic relation between input motor current (i) and output position (X) given knowledge of the speed reduction (N), motor torque constant (k_τ), and forward drivetrain efficiency (η):

$$P(s) = \frac{X(s)}{i(s)} = \frac{X(s)}{F(s)}\frac{F(s)}{i(s)} = \frac{Nk_\tau\eta}{ms^2 + bs}. \tag{7.6}$$

A scaling constant (G) is used to map desired forces into motor currents

$$G = \frac{i(s)}{F(s)} = \frac{1}{Nk_\tau \eta}. \tag{7.7}$$

To account for the effects of delay and filtering, two additional blocks are added to the control diagram. The e^{-Ts} block represents a time delay of T seconds while the sQ_v block represents the Laplace differentiation operator, s, multiplied by a first-order low-pass filter, Q_v, with a cutoff frequency of f_v Hertz:

$$Q_v(s) = \frac{\omega_v}{s + \omega_v} = \frac{2\pi f_v}{s + 2\pi f_v} = \frac{1}{\tau_v s + 1}. \tag{7.8}$$

The transfer function from desired position, X_d, to measured position, X, for the control system in Figure 7.2 can be found to be

$$\psi(s) = \frac{X(s)}{X_d(s)} = \frac{Bs + K}{ms^2 + bs + e^{-Ts}(BQ_v s + K)}. \tag{7.9}$$

As a check, for the case where there is no delay ($T = 0$) and where an unfiltered velocity signal is available ($Q_v = 1$), (7.9) simplifies to the familiar second-order mass-spring-damper equation.

What are the maximum values of K and B when delay and filtering are not ignored? This is a difficult question to answer for several reasons. First, the function describing maximum K and B values is affected by many system parameters. In the idealized system previously discussed, there was no maximum value of K, and B was a function of four parameters $B =$

$f(\zeta_d, m, b, K)$ where f is the well-known function (7.5). We would like to find such a closed-form expression for K_{max} and B_{max} using the non-idealized system model (7.9). In this case, we must find a function f which maps passive system parameters *and* non-idealized parameters such as delay (T) and filter cutoff frequency (f_v) to a single set of K's and B's:

$$[K_{max}, B_{max}] = f(\zeta_d, m, b, T, f_v). \tag{7.10}$$

A second complicating factor is that systems with pure delay cannot be analyzed using conventional pole/zero techniques unless approximations of the e^{-Ts} term are used. The delay term makes analysis of the system's behavior more difficult and necessitates the use of alternate stability analysis methods, such as the Nyquist Stability Criterion.

A third difficulty lies in how the word "maximum" is defined. Are the maximum values of K and B the largest values that ensure system passivity, system stability, or some system stability margin?

The following sections discuss how our parameter selection approach addresses these issues.

7.3 Our Approach

The goals of this section are to address the questions raised in Section 7.2, namely 1) to identify an appropriate selection criteria for K and B, 2) to understand the relationships between maximum controller parameters K, B

and the system parameters m, b, T, and f_v.

Our approach to answering these questions centers around analysis of stability margins of the complete system transfer function $\psi(s)$, including derivative filtering and time delay. We rely on parametric sweeps to identify trends and gauge the sensitivity of $\psi(s)$ to each system parameter. We use a phase margin threshold instead of a passivity or stability threshold as a trade-off between system performance and robustness.

With an understanding of how $\psi(s)$ depends on the various system parameters, we perform a thorough search of the system parameter space which is intended to encompass a large portion of actuation systems in use today. Based on the data collected from this experiment, we then fit continuous equations to the experimental data. As a result, we obtain a generic closed-form expression which may be used to select values for K and B which yield the maximum actuator impedance meeting our phase margin criteria.

An important theme in the following discussion is the need to reduce the complexity and coupling between parameters of (7.10). We simplify the problem by considering a dependent set of K and B values where one may be calculated if the other is known using critically damped assumptions. Using the phase margin thresholding criteria, we discover a useful relationship between maximum closed-loop impedance and the passive system corner frequency which allows us to treat m and b as a single term.

Our approach relies heavily on sampled simulation data and not on

theoretical proofs. Therefore, the results shown in this chapter should be taken as evidence, not proof, that such relationships exist.

7.3.1 A Critically Damped Constraint

Let us assume a critically damped impedance response is a design constraint. A critically damped response has several desirable properties such as a near-minimum settling time (assuming a small error-band tolerance) and a high stability margin. In addition, a critically damped constraint reduces the solution space of (7.10) from a two degree of freedom plane of K's and B's to a single array of critically damped K-B pairs. These K-B pairs can be usefully characterized by their natural frequency, f_n, where

$$\omega_n = 2\pi f_n = \sqrt{\frac{K}{m}} \tag{7.11}$$

and B is selected using (7.5) with $\zeta_d = 1$. In the remainder of this chapter, it is assumed that when we discuss varying f_n, we do so by altering K (and B to maintain the critically damped constraint), and not by choosing a different m. Therefore, there is a direct correlation between f_n, the values of K and B, and actuator impedance.

It is important to note here that while we place a critical damping constraint on the selection of K and B, this does not guarantee the dynamic response of the system transfer function ($\psi(s)$) will be critically damped. The reason for this is because (7.5) only guarantees critical damping for second-

order systems and $\psi(s)$ is not a second-order system. Therefore, the main purpose for the critically damped constraint is to reduce the solution space, not to guarantee absence of overshoot or other critically damped properties in the dynamic response of $\psi(s)$. The selection of parameters to guarantee a desired response of $\psi(s)$ is the topic of the remainder of this chapter.

Using the critically damped constraint, the solution space of (7.10) is reduced to a single degree of freedom and is a function of four variables:

$$f_{n_{max}} = f(m, b, T, f_v). \tag{7.12}$$

Suppose that, for the system described by (7.9), f_n is varied while m, b, T, and f_v remain constant. If we measure the step response for each of these permutations of $\psi(s)$ we obtain the results shown in Figure 7.3. As f_n is increased (and thus K and B are increased) the response becomes faster. If f_n is increased too far, the response begins to deviate from the shape of a critically damped response and begins to oscillate. The figure also shows the system's phase margin for each response (obtained using the Matlab `margin` command). The phase margin is a measure of stability, higher being more stable and more damped and lower being less stable and more vibratory. The phase margin is a useful metric in this case because it may be calculated for systems with pure delay (see (Franklin et al., 1986; Ogata, 1990) for more information).

The oscillation shown in Figure 7.3 is caused by the combined effects

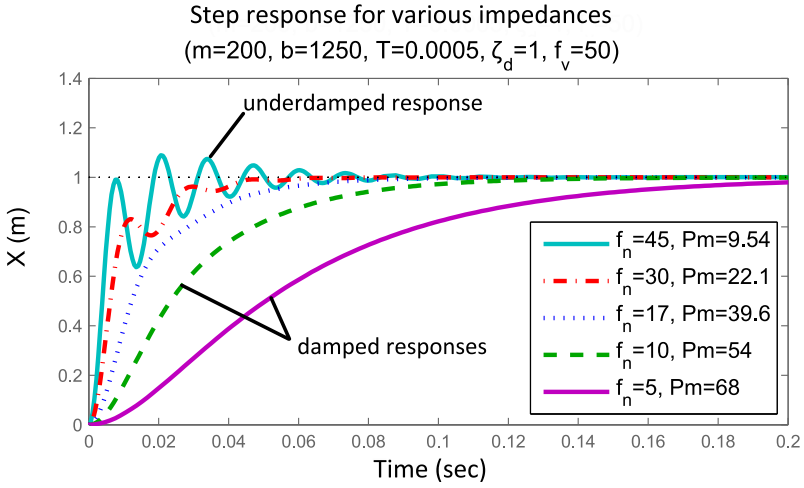

Figure 7.3: Step response of $\psi(s)$ for various values of f_n. The phase margin (Pm) of each response is shown. Response deformation begins to occur at a phase margin of 39.6 degrees and large oscillations are visible for a phase margin of 9.54 degrees.

of derivative feedback filtering frequency, f_v, and time delay, T. To gain more insight into this phenomenon, we can perform a similar experiment as before, but this time we vary both f_v and f_n while the other parameters remain constant. If instead of plotting the time domain response, we plot each response's phase margin, we obtain the results shown in Figure 7.4. Here, we can see how the system phase margin, and thus its dynamic response, is a nonlinear function of f_n and f_v.

These two experiments yield useful insights. First, they demonstrate a relationship, albeit complex, between system stability, feedback gains, and the level of filtering applied to the derivative term. They also demonstrate a connection between the effective damping ratio of the system's step response

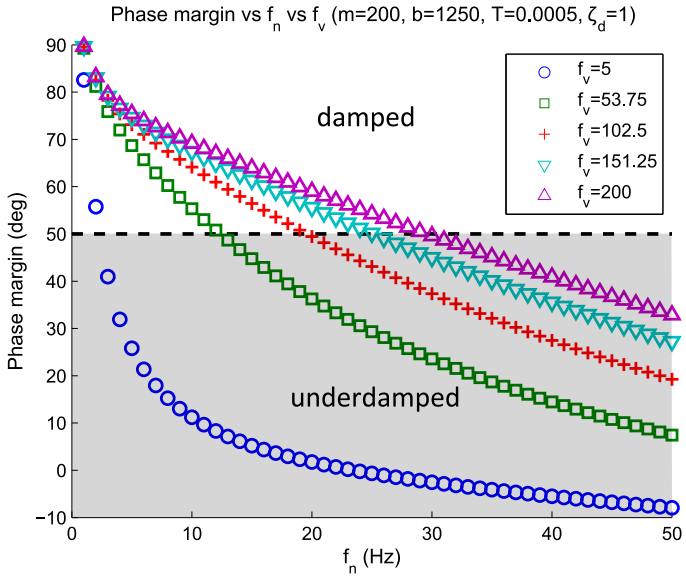

Figure 7.4: Phase margin of $\psi(s)$ for various values of f_n and f_v. The system is destabilized either by heavily filtering the derivative term (lower f_v values) or by increasing feedback gains (higher f_n values). A phase margin threshold is shown at 50 degrees. This threshold is determined by observing the minimum phase margin step response which does not exhibit oscillatory distortion in Figure 7.3. Parameter combinations producing phase margins above this line are represented by an 'o' in Figure 7.5 while those below are represented by an 'x'.

and the system's phase margin. Clearly, a phase margin of 9.54 degrees produces an underdamped step response for these parameters while phase margins above 54 degrees appear to possess a more damped response.

7.3.2 A Phase Margin Criteria

Phase margins are known to relate to system damping (Franklin et al., 1986; Ogata, 1990). Suppose, based on the results shown in Figures 7.3 and 7.4, we consider any system with a phase margin greater than 50 degrees to

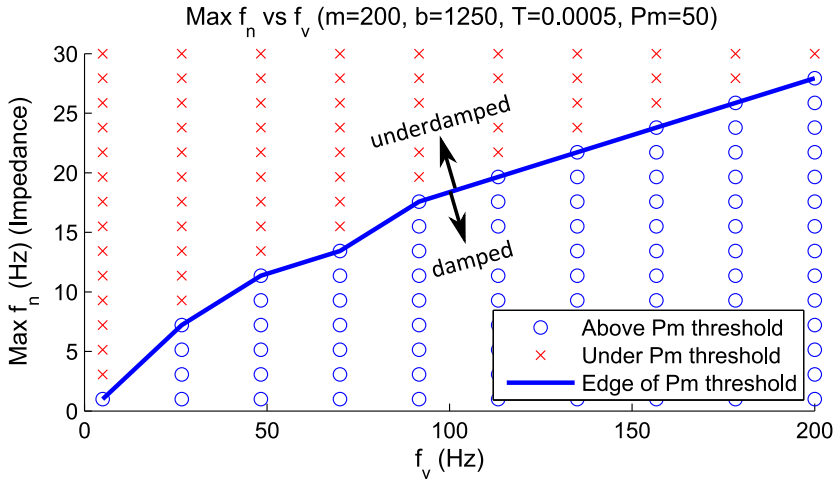

Figure 7.5: A parametric search of phase margins of $\psi(s)$ across values of f_n and f_v. Combinations producing phase margins (Pm) above the Pm threshold are represented by an 'o' while those below are represented by an 'x'. The line represents the maximum values of f_n which pass the phase margin criteria and is analogous to the dashed line seen in Figure 7.4.

possess a damped, non-vibratory step response. Note that when $\zeta_d = 1$ this characteristic dynamic response is similar to a second-order system's critically damped response in that it has a fast rise time yet does not exhibit periodic signal content. However, the response of $\psi(s)$ cannot be directly compared to a second-order system due to the effects of filtering and delay. The phase margin criteria is independent of the source of instability, whether it is caused by f_v being small, T being large, or some other cause. Regardless of the numerical value of this phase margin threshold, the following analysis may still be performed.

Performing the same search as was used in the experiment for Figure 7.4, where f_v and f_n are varied while other parameters remain constant, we

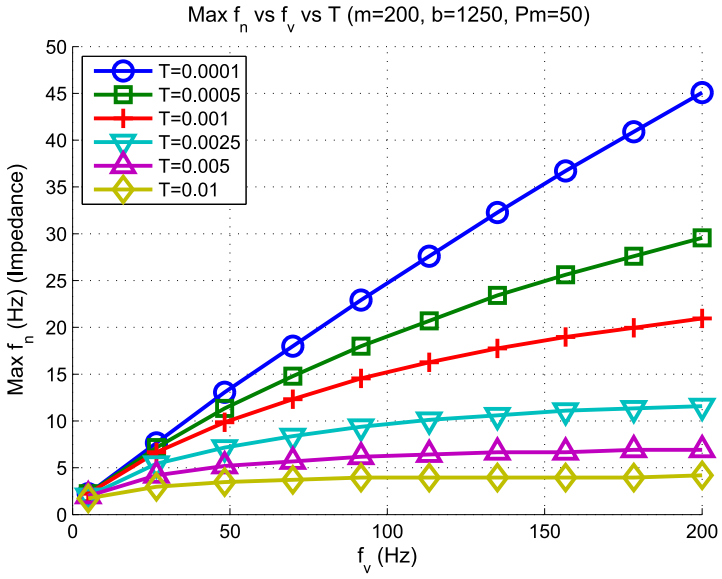

Figure 7.6: A parametric search similar to the search in Figure 7.5 except with higher vertical resolution and an added dimension showing sensitivity to time delay, T. Because f_n represents specific values of K and B, this plot can be used to find the maximum values of K and B, given f_v and T, which produces an impedance controller with a phase margin of 50 degrees.

may instead plot the result of a boolean comparison of each system's phase margin with a phase margin threshold of 50 degrees. Figure 7.5 shows such a search. For each combination of f_n (y-axis) and f_v (x-axis), an 'o' represents a phase margin greater than 50 degrees while an 'x' represents a phase margin below 50 degrees. The line connecting the maximum values of f_n represents the highest impedance, or the highest values of K and B, which satisfy our phase margin requirement. As before, we see that larger values of f_v produce a more stable system, but now we are able to identify the maximum f_n corresponding to each value of f_v.

Going one step further, we can add T as an additional search parameter. Figure 7.6 reveals these results. For clarity, only the boundary lines are shown defining maximum values of f_n. We gain several insights from Figure 7.6:

1. Achievable impedance increases with f_v (heavily filtering velocity feedback reduces stability).

2. Achievable impedance increases as T decreases (large delay reduces stability).

3. The numerical value of f_v becomes less important as T increases.

The results shown in Figure 7.6 are enough to fully describe a single system, characterized by its values of m and b. We can easily select K and B values by using Figure 7.6 as a look-up table to obtain f_n given f_v and T. With f_n found, K and B can be calculated using (7.11) and (7.5). However, if we select a different m or b value, the results of Figure 7.6 will no longer be valid. The next section focuses on the generalization of these results to apply to a wider range of systems.

7.4 Generalization

The goal of this section is to take the results shown in Figure 7.6, which only apply to a single pair of m-b values, and generalize them to a wide range of actuators characterized by different m's and b's. Our goal is to find an equation which closely matches a discrete set of data points generated from

parametric searches. The data points are the maximum values of f_n ($f_{n_{max}}$) as a function of m, b, T, and f_v which produce phase margins of 50 degrees. Therefore, a continuous function should be in the form of the function f from (7.12). Recall from Section 7.3.1 that f_n is directly related to closed-loop actuator impedance, and therefore maximum system impedance is obtained through $f_{n_{max}}$.

Let us consider the closed-loop impedance transfer function (7.3). If we plot impedance magnitude versus frequency for various values of f_n we obtain the results shown in Figure 7.7. If K and B are set to zero, then the impedance of the actuator reduces to the impedance of the passive system. Since the passive impedance is purely a function of velocity and acceleration, the impedance drops to zero as the frequency decreases. That is, the passive system does not respond to constant position offsets; it has no "stiffness." If we select non-zero values of K and B we see how impedance at low frequencies increases, as does the bandwidth of the impedance response.

This frequency representation is useful as it allows us to visualize the closed-loop natural frequency, f_n, as well as the corner frequency of the passive system, f_p. f_p can be found given m and b as follows:

$$\frac{V(s)}{F(s)} = \frac{1}{ms+b} = \frac{1/b}{\tau_p s + 1} = \frac{1/b}{\dfrac{s}{2\pi f_p} + 1} \tag{7.13a}$$

$$f_p = \frac{b}{2\pi m} = \frac{\omega_p}{2\pi}. \tag{7.13b}$$

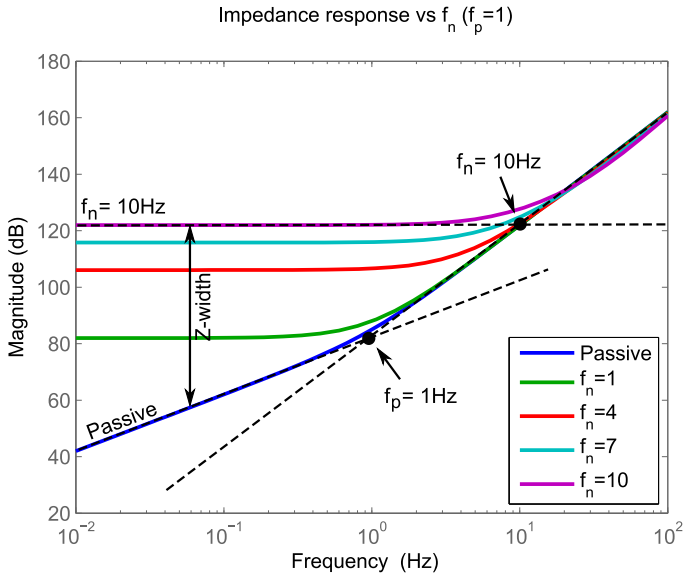

Figure 7.7: Impedance frequency response (F_{ext}/X) of an actuator with various closed-loop gains (K,B,$\zeta_d = 1$) determined by f_n. With K and B set to zero, the impedance response is that of the passive actuator. The open loop passive corner frequency (f_p) is shown as well as the closed-loop natural frequency (f_n) for $f_n = 10$Hz. The difference between maximum and passive impedance, Z-width, is also illustrated.

We use these relationships in the following section to understand the influence b and m have on $f_{n_{max}}$.

7.4.1 The Effects of b and m on Maximum Impedance

In earlier work, b has often been solely considered as the system parameter affecting achievable closed-loop impedance (Colgate and Schenkel, 1994; An and Kwon, 2004; Mehling et al., 2005; Weir et al., 2008; Rossa et al., 2013). A conclusion was drawn, then, that to maximize impedance, a large b is necessary. However, these studies do not directly analyze the influence of m on

system impedance.

If we solve (7.11) for K, substitute K into (7.5), refactor (7.9) based on these expressions of K and B with $\zeta_d = 1$, and divide by m, we obtain another representation of $\psi(s)$ which is a function of natural frequency (ω_n), passive corner frequency (ω_p), delay (T), and derivative filter frequency (ω_v):

$$\psi(s) = \frac{(2\omega_n - \omega_p)s + \omega_n^2}{s^2 + \omega_p s + e^{-Ts}[(2\omega_n - \omega_p)Q_v s + \omega_n^2]}. \tag{7.14}$$

From (7.14) it is apparent that the stability of $\psi(s)$ does not change with b as long as ω_p (or the ratio of b/m) is held constant. It is then ω_p (or $f_p \propto b/m$), not b alone, which determines the maximum feedback gains, and thus the maximum impedance of the closed-loop actuation system. Note that this observation does not contest the conclusions drawn in (Colgate and Brown, 1994; An and Kwon, 2004; Mehling et al., 2005; Weir et al., 2008; Rossa et al., 2013). Indeed, if m is held constant then achievable closed-loop impedance increases with b.

The observation of maximum impedance dependency on the ratio of b/m rather than b alone could offer useful insight towards designing devices for high output impedance. For example, device designers attempting to maximize output impedance who were operating based on the recommendations of (Colgate and Schenkel, 1994; An and Kwon, 2004; Mehling et al., 2005; Weir et al., 2008; Rossa et al., 2013) would attempt to maximize sampling rate (minimizing T) and may either maximize b or make b controllable or tunable.

The prior work in this area makes no recommendation on how to select m. It is true that m may vary, for example, based on human grip type in haptic devices or based on kinematic configuration in redundant manipulators, but its influence is still important nonetheless. One interesting example of the influence of m on device design could be the selection of actuator speed reduction, N, as it critically influences the magnitude of the actuator load inertia seen by the motor; larger values of N correspond to lower actuator sensitivity to load inertia based on the $1/N^2$ relationship between load inertia and the effective inertia seen by the motor.

Dependency on the ratio b/m greatly simplifies the generalization process. If, for example, $f_{n_{max}}$ were a function of b alone, the units of b would affect the calculation of $f_{n_{max}}$. In SI units, rotary damping is represented by Nms/rad where linear damping is represented by N/m. An equation which accepts units of rotary damping would have to be altered to produce the same results for units of linear damping. Or, the units could be normalized to some predetermined reference frame before being passed to an equation. Dependency on the ratio of b/m removes this issue as both linear and rotary systems may be represented by their respective time constants, τ_p, or corner frequencies, f_p.

In contrast to (7.12), dependency on f_p allows us to further simplify the function describing $f_{n_{max}}$:

$$f_{n_{max}} = f(m, b, T, f_v) \Rightarrow f(f_p, T, f_v). \tag{7.15}$$

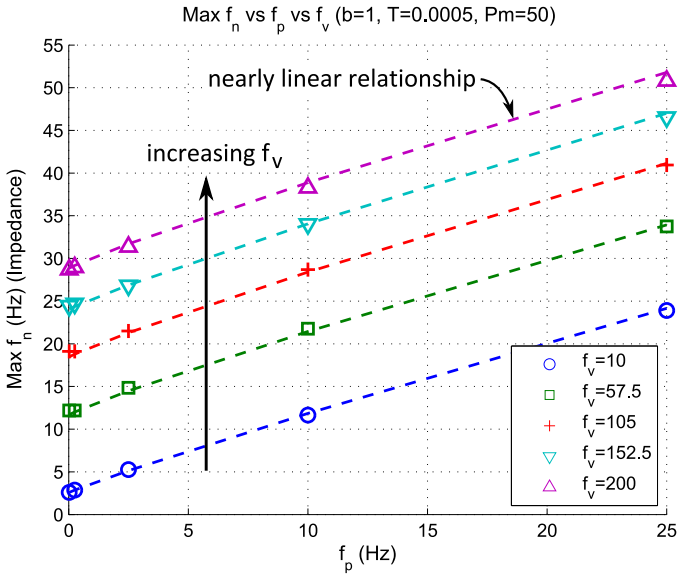

Figure 7.8: Parametric search of maximum impedance for various values of f_p and f_v. A relation which is strongly linear can be seen between $f_{n_{max}}$ and f_p. The markers represent simulation data while the dashed lines represent values calculated using (7.16).

7.4.2 Fitting Curves to Sampled Data

To study the relationship between f_n, f_p, and f_v further, we perform a search of maximum impedance parameters across values of f_p and f_v (see Figures 7.8 and 7.9). The results show a relationship between maximum f_n and f_p which is very nearly linear. f_v then acts as an offset to this linear relationship. This observation creates the basis of our generalization strategy.

Based on the results shown in Figure 7.8, let us assume that there is a relationship between maximum f_n and f_p which is *almost* linear. For some unknown variables c, d, and e, we can represent this relationship with the following equation:

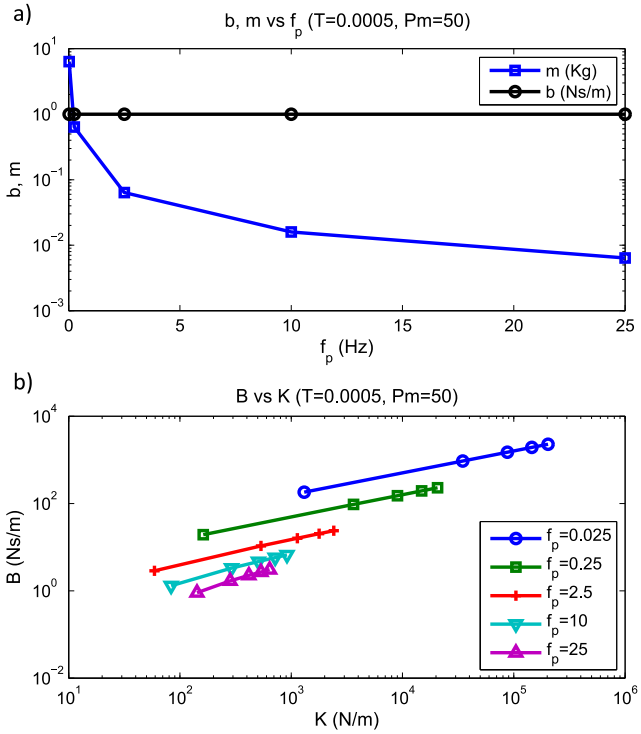

Figure 7.9: **a)** Search space of m and b for the experiment shown in Figure 7.8. **b)** Solution space of K and B for the experiment shown in Figure 7.8. Each point represents a K, B pair producing a system with the target phase margin.

$$f_{n_{max}}(f_p, f_v, T) = c(f_p)^d + e. \tag{7.16}$$

The problem of finding a closed-form expression describing maximum f_n is now reduced to a surface fitting problem of c, d, and e which are each a function of f_v and T.

To fit a surface to c, d, and e, we perform a series of searches which

Table 7.1: Fitting bounds of $f_{n_{max}}$ equation

Parameter	Search space	Units	Orders of magnitude
f_p	$0.025 \rightarrow 25$	Hz	3
f_v	$10 \rightarrow 200$	Hz	1.3
T	$0.0001 \rightarrow 0.01$	Seconds	2

cover our desired parameter space. The wider the search space, the more generally applicable a closed-form solution will be. Keeping in mind the scale-independent nature of the f_p parameter, the authors believe a large percentage of actuators in existence today fit within the search space shown in Table 7.1. The most viscous joints on NASA-JSC's humanoid Valkyrie robot[2] (NASA-JSC, 2013) with grease-lubricated harmonic drive actuators have an f_p value of 9.9 Hz, while the low friction ball screw actuator shown in Section 7.6 has an f_p value of 0.77 Hz. A standalone 200W electric motor (Maxon EC-powermax 30) only has friction from its bearings and therefore has a lower f_p value of 0.18 Hz. Values for f_v depend primarily on the amount of sensor noise present in the position encoder signal, but based on the author's experience values in the 20-70 Hz range are common. Sampling rates in the range of 1000 Hz (T = 0.0005 seconds) are widely used in robotic servo systems.

After performing a series of searches using automated scripts we were able to find suitable fits for c, d, and e with low error bounds. Polynomial fits

[2]The authors spent a year working with the Valkyrie robot during the 2012-2013 portion of the DARPA Robotics Challenge.

produced the best results for the c and d parameters, while a more exotic form based on power fitting produced the best results for e. The equations for each of these parameters are as follows:

$$c(f_v, T) = c_1 + c_2 f_v + c_3 T + c_4 {f_v}^2 + c_5 f_v T + c_6 T^2 + c_7 {f_v}^3 + c_8 {f_v}^2 T + c_9 f_v T^2 + c_{10} T^3 \tag{7.17a}$$

$$d(f_v, T) = d_1 + d_2 f_v + d_3 T + d_4 {f_v}^2 + d_5 f_v T + d_6 T^2 + d_7 {f_v}^3 + d_8 {f_v}^2 T + d_9 f_v T^2 + d_{10} T^3 \tag{7.17b}$$

$$e(f_v, T) = e_1 T^{e_2} + e_3 + (f_v + e_4)(e_5 {f_v}^{e_6} T^{e_7} + e_8 f_v T + e_9). \tag{7.17c}$$

The numeric values of the 29 parameters used in equations (7.17a - 7.17c) can be found in Table 7.2.

An example of the surface fit for the d parameter is shown in Figure 7.10. As can be seen, the continuous surface closely matches the simulation values. The coefficient of determination (R^2) for this fit is 0.986.

The accuracy of this closed-form approximation can be seen for $T = 0.0005$ seconds in Figure 7.8 where the dashed lines represent the closed-form approximation and the markers represent the real simulation values. Additionally, Figure 7.11 shows the fitting accuracy of the $f_{n_{max}}$ equation over the entire parameter space, using averaged values for the f_v parameter. As can be seen, for the majority of the parameter space the $f_{n_{max}}$ equation closely fits the ideal simulation values (2%-4% error). Peak error of 21% occurs in the corner case when $T = 0.01$ seconds and $f_p = 0.025$ Hz.

Table 7.2: $f_{n_{max}}$ equation parameters

Index	c	d	e
1	1.093	0.9544	-14.77
2	0.004883	-0.001039	0.4916
3	-54.2	-51.65	2.908
4	-3.694e-5	9.111e-6	-10
5	-0.2871	0.0638	-0.5162
6	1.541e4	6918	0.2257
7	9.201e-8	-2.451e-8	0.2566
8	-4.08e-4	1.559e-4	0.08373
9	49.89	-13.29	0.3725
10	-9.713e5	-3.869e5	n/a

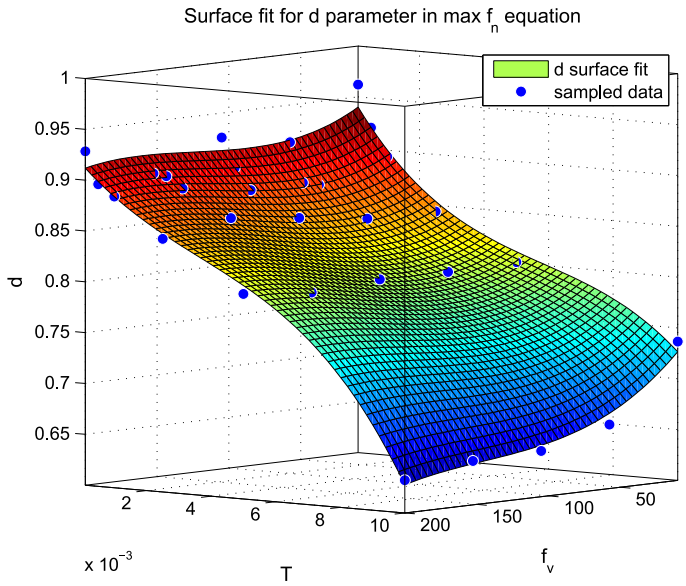

Figure 7.10: An example of the fitting process used to match the continuous $f_{n_{max}}$ equation to data points gathered from simulation. This fit represents the d term in (7.16), (7.17b).

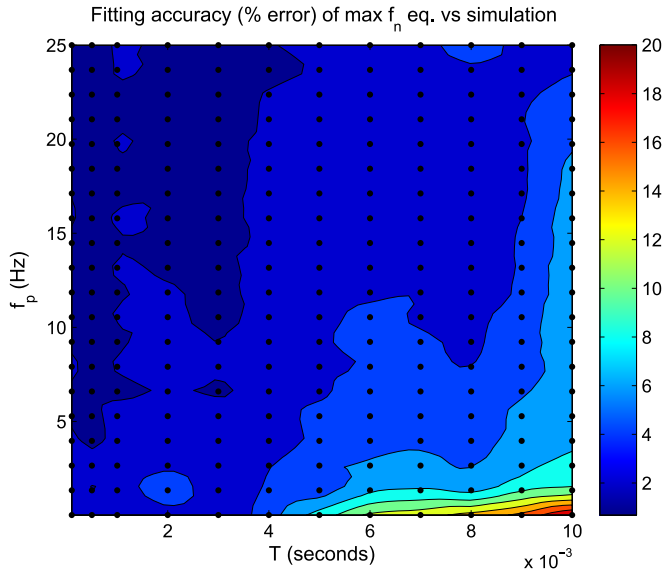

Figure 7.11: A graphical representation of the fitting accuracy of the $f_{n_{max}}$ equation compared to ground truth simulation values. The 240 sample points are marked on the contour plot. Error values are averaged along the f_v dimension to simplify data representation. Error percentage remains below 5% for values of T <0.005 seconds. Maximum error (21%) occurs in the corner case where T=0.01 seconds and f_p=0.025 Hz.

7.4.3 An Example: Applying the $f_{n_{max}}$ Equation

One of the primary benefits of equation (7.16) is that it transforms gain tuning of rigid actuators with position feedback from a guess-and-check process to a simple and deterministic process. Here we demonstrate how this procedure is applied.

First, m and b must be identified. Section 7.6.1 provides one method for accomplishing this using a linear actuator of our design.

With m and b known, we can calculate the passive corner frequency,

f_p:

$$f_p = \frac{b}{2\pi \cdot m} = \frac{1250}{2\pi \cdot 256} = 0.77\,\text{Hz}. \tag{7.18}$$

We then need T and f_v to solve for $f_{n_{max}}$. T should be calculated as T = round trip communication latency + sampling period / 2. In our case, we are using a computer directly connected to the actuation hardware and so the effective delay is dominated by the sampling period ($T = 0.0005$ seconds for a sampling rate of 1kHz). f_v can be found by implementing the B term of the impedance controller in hardware for some large but stable value of B and slowly increasing f_v until just before noise is felt or heard when backdriving the output. In our case, we chose $f_v = 50$ Hz. If noise is felt or heard for the final values of K and B, f_v should be reduced and $f_{n_{max}}$ recalculated.

Once f_p, f_v and T have been found, we may then solve for $f_{n_{max}}$ using (7.16):

$$f_{n_{max}}(f_p = 0.77, f_v = 50, T = 0.0005) = 11.4\,\text{Hz}. \tag{7.19}$$

K can then be found using (7.11):

$$\begin{aligned} K_{max} &= (2\pi \cdot f_{n_{max}})^2 \cdot m = (2\pi \cdot 11.4)^2 \cdot 256 \\ &= 1,313,437\,\text{N/m}. \end{aligned} \tag{7.20}$$

Finally, B is found using (7.5) with critical damping ($\zeta_d = 1$):

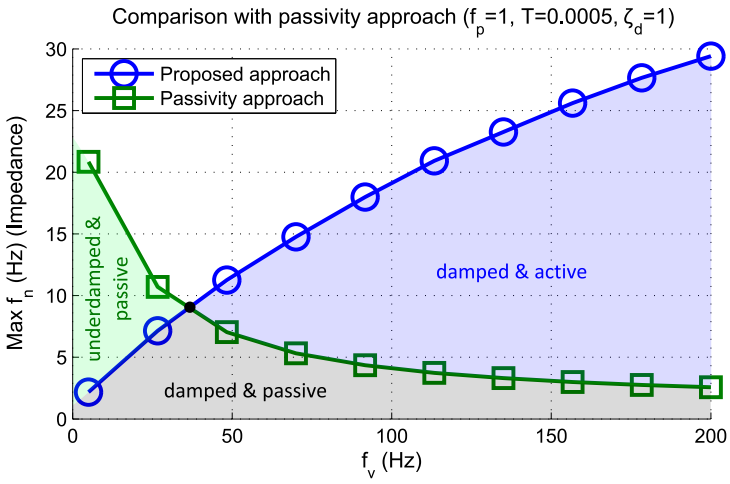

Figure 7.12: A comparison between maximum impedance for gains selected by the proposed approach and gains selected by a passivity approach (Colgate and Brown, 1994).

$$B_{max} = 2\zeta_d\sqrt{mK_{max}} - b = 2 \cdot 1\sqrt{256 \cdot 1313437} - 1250$$
$$= 35,423\,\mathrm{N}\cdot\mathrm{s/m}. \tag{7.21}$$

We show the experimental results of applying these values in Section 7.6.

7.5 Comparisons with Passivity Approach

In (Colgate and Brown, 1994), equations are provided for determining maximum values for K and B based on a passivity criterion:

$$b > \frac{KT}{2} + \frac{BT}{2\tau_v + T} \qquad (B \geq 0) \tag{7.22a}$$
$$b > \frac{KT}{2} - B \qquad (B < 0). \tag{7.22b}$$

If we apply the critically damped constraint from Section 7.3.1, we can solve (7.22) for a single pair of values for K and B producing the maximum impedance with passivity properties.

Figure 7.12 shows a comparison of maximum impedance as a function of f_v for our proposed approach (7.16) and for the passivity approach. A few interesting observations are readily apparent. First, impedance is proportional to f_v from a phase margin standpoint and inversely proportional from a passivity standpoint. Second, the values of the passivity line lying above the proposed line (f_v=10→30 Hz) will produce an underdamped response, while the values of the proposed line lying above the passivity line (f_v=50→200 Hz) will produce an active response, meaning the phase response penetrates the ±90 degree boundary (Bao and Lee, 2007). Third, the intersecting point of the two lines is the maximum damped impedance producing a passive response. Previous work suggests an optimal value for f_v exists but does not propose a way to determine such a value (Colgate and Brown, 1994). This point of intersection could prove to be a promising candidate.

7.6 Experimental Validation

In this section we apply our approach of finding maximum impedance parameters to two different actuation systems. Up to this point, all of our analysis has been performed in a simulation environment. The goal of this study is to check how closely matched the responses of real-world actuators

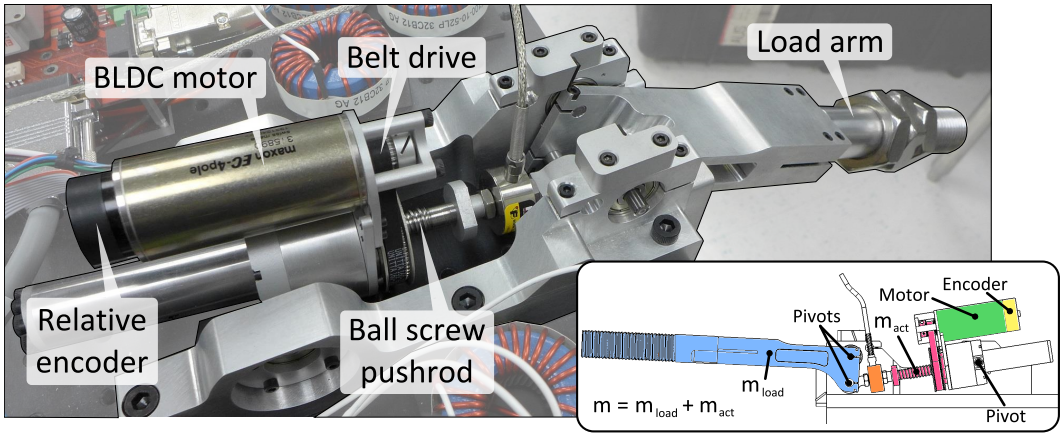

Figure 7.13: Ball screw pushrod actuator used in experimental tests. For the standalone motor experiment, the belt was removed so that the motor could spin freely. In the full actuator experiment, the belt was connected, coupling motor motion to ball screw and load arm motion. The depicted load cell is unused in these tests.

are to actuators we have simulated. Our approach is validated if the responses of the real-world actuators become underdamped for the same parameters that cause the simulated actuators to become underdamped.

The properties of the two actuation systems differ significantly in order to demonstrate the wide applicability of our approach. The first system is a standalone brushless direct current (BLDC) motor. Using the system shown in Figure 7.13, we disconnect the belt from the ball screw drive and test the motor with no load attached. The motor is a small scale rotary actuator with very little passive damping. The second system is the full actuator shown in Figure 7.13. The output of the ball screw pushrod is connected to the load arm. Small displacements are used to operate the actuator in an approximately linear region of its load inertia. The actuator is a medium scale linear actuator with

Table 7.3: Identified system parameters of rigid actuator

Parameter	Standalone motor	Full actuator
m	3.0e-6 $[Kg \cdot m^2]$	256 $[Kg]$
b	3.5e-6 $[Nm \cdot s/rad]$	1250 $[N \cdot s/m]$

a large, two-stage speed reduction (209:1 at the arm output). This actuator is a new design that iterates on the actuator presented in (Paine et al., 2014b) and is intended for use in performance studies. It was designed to further reduce weight and size and improve maintainability compared to the previous iteration. The new design uses a modular spring assembly. For the experiments in this chapter, a spring was not used in order to match the model shown in Figure 7.1b.

7.6.1 System Identification

To apply (7.16) we must have an accurate model of m and b. While inertial values may be obtained through datasheets and computer-aided design (CAD), it is prudent to ensure the same values are obtained empirically. Damping values are difficult to anticipate and therefore are best measured directly. Refer to Appendix A.2 for one method of empirically measuring m and b. Using these techniques, we obtained system parameters for the two actuation systems as shown in Table A.1. The data for the full actuator is in the linear output reference frame.

For systems with variable inertia, control parameters may be selected in

several different ways. One option is to select parameters for the scenario where gains are minimized (when m is minimum). This method will produce the most stable system, but will also become underdamped when m is increased. Another option is to select parameters for some nominal value of m. This method will produce a system with trade-offs between control loop stability and performance considered from a damping perspective. A third option is to recalculate stiffness and damping parameters whenever m changes significantly. This method requires higher computational load, but will maintain optimal control parameters throughout the system workspace.

7.6.2 Standalone Motor Experiment

In our first experiment, we observed the dynamic response to a step input of the standalone motor using two sets of impedance control parameters. The first set of parameters was chosen using our proposed approach (7.16). The second set of parameters was chosen to be greater than the values in the first set, with the hypothesis that higher values should produce an underdamped response (as demonstrated in Figure 7.3). Specifically, parameters were chosen by doubling the B value obtained from (7.16), and then selecting K using the critically damped constraint. Figure 7.14 shows the results. In Figure 7.14a, we first observe that, for both sets of parameters, the experimental data closely matches the simulation data. Second, it shows that the parameters obtained using the proposed method produce a damped response while the second set of parameters produce a response with small oscillations. This result confirms

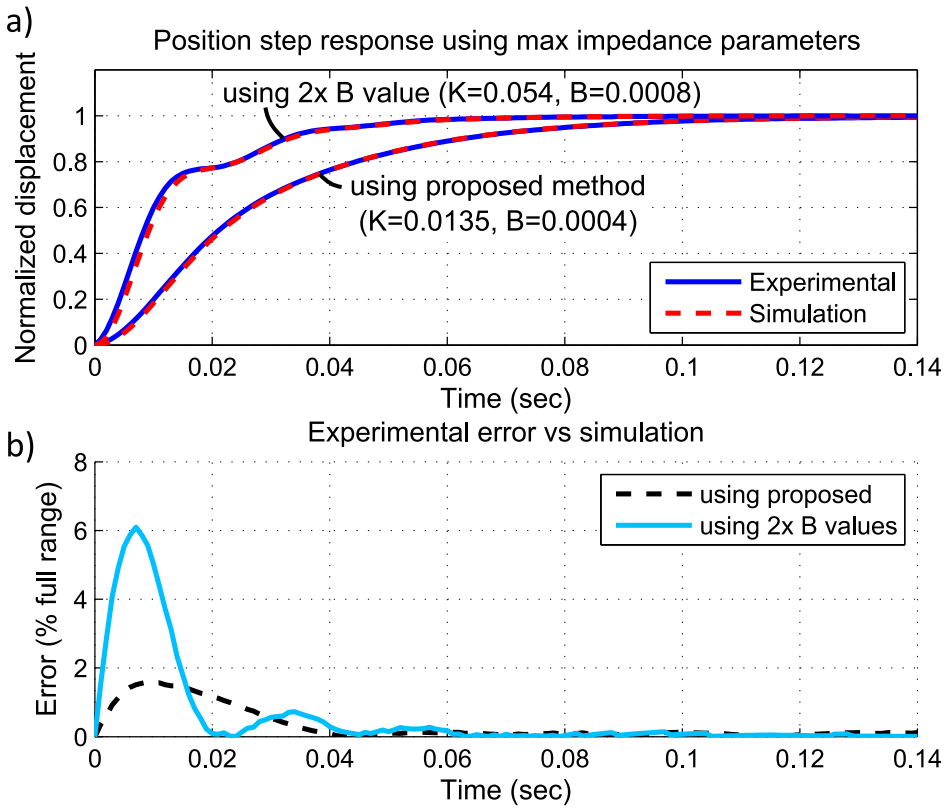

Figure 7.14: **Standalone Motor Experiment. a)** Step responses for two different sets of parameters are shown. The first set (proposed method) was obtained using the $f_{n_{max}}$ equation (7.16). The second set (2x B value) used double the B parameter from the first set, and selected K using the critically damped constraint. The higher gains produce a deformed step response which exhibits small oscillation and therefore exceeds the maximum achievable actuator impedance with a phase margin of 50 degrees. The step displacement for this test was four motor rotations. Due to the high gains used, a higher displacement would cause current saturation to occur (30 amp limit). **b)** Discrepancy between simulation and experimental results. Error peaks at 6% showing the simulation accurately represents real-world effects.

that, for this case, our proposed method functions as expected and accurately provides the maximum critically damped parameters in practice.

Figure 7.14b shows the deviation of the experimental results from the simulation results. In the worst case, error reaches just 6% of the full range of the step amplitude indicating that the experimental results closely match the model.

7.6.3 Full Actuator Experiment

Our second experiment is similar to the motor experiment, except that the full actuator was used. Two sets of impedance parameters were again tested, the first chosen using our proposed approach (7.16) and the second set chosen to be greater than the values in the first set. Figure 7.15 shows the results. Here, the results are similar to the motor experiment, except that the deviations between the experimental data and the model are greater. The extra deviation is likely caused by the dynamics of the drivetrain, particularly the belt, which may deflect slightly under large torques. Motor current in this test reached values of 15 amps, or more than three times the motor's rated continuous value. Similar deviations may be expected on drivetrains with limited stiffness, such as Harmonic Drives and cable drives. Despite these deviations, error between experiment and simulation remains small, peaking at 7% of full scale.

Perhaps of greater significance than error magnitude is the fact that our hypothesis holds true. Our proposed method produces a response which, upon

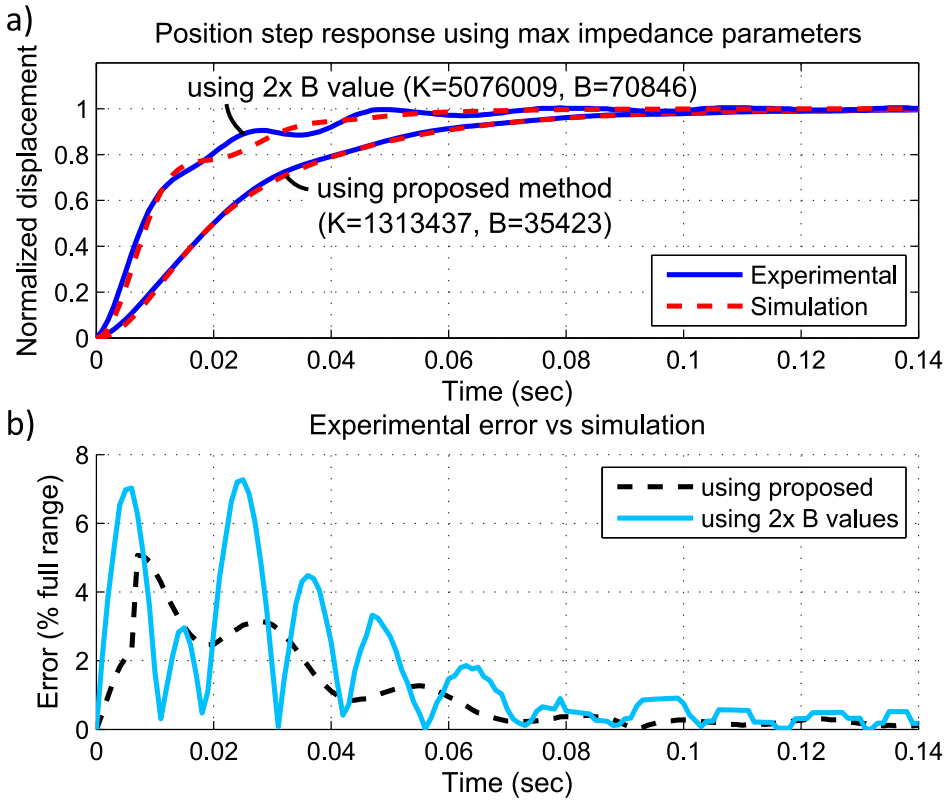

Figure 7.15: **Full Actuator Experiment.** The same experiment was performed as was described in Figure 7.14 except with the full actuator. While the experimental data closely matches simulation data, larger discrepancies can be seen compared to Figure 7.14. The cause of this increase is likely due to the drivetrain dynamics (particularly the belt). As was the case in the motor experiment, our method again correctly chooses the maximum critically damped control parameters with a phase margin of 50 degrees. The step displacement for this test was 2mm. Due to the high gains used, a higher displacement would cause current saturation to occur (30 amp limit).

visual inspection, is well damped, while the higher gains produce a response which is clearly underdamped. These results demonstrate the efficacy of the work presented in this chapter. With almost no manual effort and with a single set of equations, we were able to find nearly optimal stiffness and damping feedback parameters for two very different actuations systems.

Chapter 8

Conclusions

This work has sought to propel forward the field of high-performance series elastic actuation, a technology which has direct application in many of today's active areas of robotics. An actuator is the most basic building block of a robotic system. It incorporates sensing and control, and is the unification of the various engineering disciplines, mechanical, electrical, control, and computer science, which comprise robotics as a field. Its benefits and pitfalls are inherited by the robotic system as a whole. Its optimization, therefore, is of the utmost importance.

We have addressed this need with a number of contributions. By introducing a lightweight and compact ball-screw-based SEA design in Chapter 2, we minimize mass and bounding volume while maximizing efficiency and power output of SEAs. Coupled with the methods for increasing motor power output with minimal loss in Chapter 3, these contributions improve the *mechanical capability* of SEAs. By using disturbance-observer-force based controllers in Chapters 4 and 5, we seek to fully utilize this mechanical capability. The *realized performance* of our system was shown to compare well with other documented work, improving on empirical SEA power-to-weight ratio by 47%. In

Chapter 6, we achieved strong results in terms of force controllability, achieving worst case joint tracking errors of of 1.38% of full-scale actuator output on the Valkyrie robot.

It is encouraging that our work has already substantially influenced the robotics community. To date, at least 8 copies of the UT-SEA have been produced, not including the 13 or so actuators (many within Valkyrie robots) directly derived from the UT-SEA design. The applications of these actuators range from legged robotics, to exoskeletons, to human orthotics, to stand-alone commercial products. Aside from producing leading results in terms of actuator performance, the control algorithms presented here have found use outside the lab as well. As mentioned in Chapter 6, our SEA force controller was used on every series elastic joint of Valkyrie during the DRC Trials 2013.

While not yet used outside of the laboratory environment, we expect our method for tuning rigid actuators to be beneficial in the various fields of robotics. It is simple, deterministic procedures, such as our proposed method, which have traditionally experienced the strongest adoption in practice.

8.1 Future Work

Several promising studies await further thought. First, using the forced cooling techniques outlined in Chapter 3, the power output per unit weight of the UT-SEA can easily be enhanced over what was achieved here. Slight modifications to the mechanical design of the actuator may be required given larger motor torque and speed, but the same basic principles apply. Special

attention should be played to the balance of loss between the motor servo drive and the motor to ensure that losses are minimized given the increased power consumption. The equations presented in Chapter 3 will be helpful in this regard.

Second, the maximum impedance study from Chapter 7 may be extended to SEAs. Some work is required to determine a tractable solution to this high-dimensioned problem due to the added dynamics of the control plant. It would also be best to find a *control-independent* approach, or perhaps an optimal controller for producing maximum impedance with an SEA.

Finally, while actuator-level tests of the UT-SEA are promising, a full robot has yet to be constructed based on this actuation technology. It is our hope to build such a system in the future which will demonstrate both the dynamic performance capability and overall system efficiency inherent to the UT-SEA design.

Appendix

Appendix A

System Identification Techniques

A.1 Series Elastic Actuators

The plant parameters in a series elastic actuator can be identified by fixing the output of the actuator, commanding an input signal, and measuring the actuator output. In our case, we chose to apply a desired motor current as an input signal and measure spring force as the actuator output. Note that these signals correspond to the plant input and output used in Chapter 4. This control plant is shown in Figure A.1 and is represented by

$$P(s) = \frac{F_o(s)}{i(s)} = \frac{N k_\tau \eta k}{m_k s^2 + b_{eff} s + k} = \frac{\beta k}{m_k s^2 + b_{eff} s + k}. \tag{A.1}$$

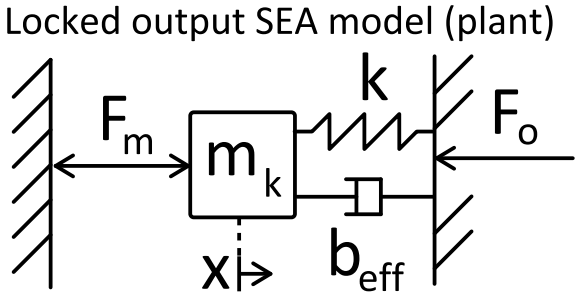

Figure A.1: SEA plant model used for system identification and control design. Note that a locked-output assumption is used.

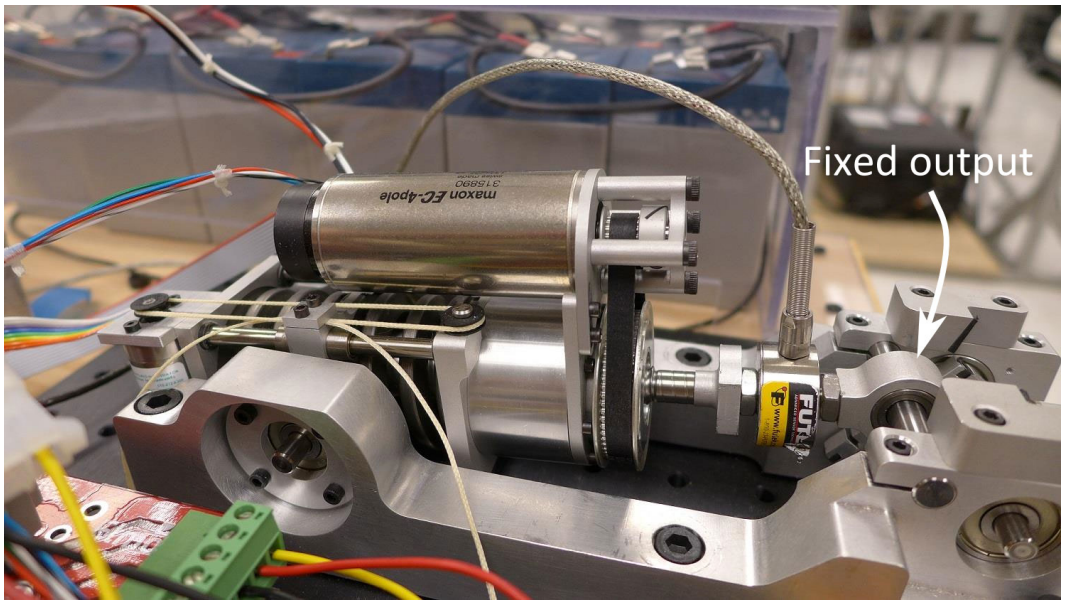

Figure A.2: UT-SEA system identification setup.

Therefore, assuming k is measured beforehand, N is known, and $k_\tau\eta$ can be lumped together as one parameter, there are three parameters that must be identified: $k_\tau\eta$, m_k, and b_{eff}.

A picture of our system identification setup is shown in Figure A.2. Note that the actuator output is attached to ground so that it cannot move. Using this setup, we command a chirp signal as the desired current and measure spring force. The input signal should be large enough to obtain significant measured actuator output and small enough to avoid saturating the actuator (current limiting or voltage limiting for example). The results of this experiment are shown in Figure A.3. A bode plot is created from this data (see Figure A.4) using the following Matlab code.

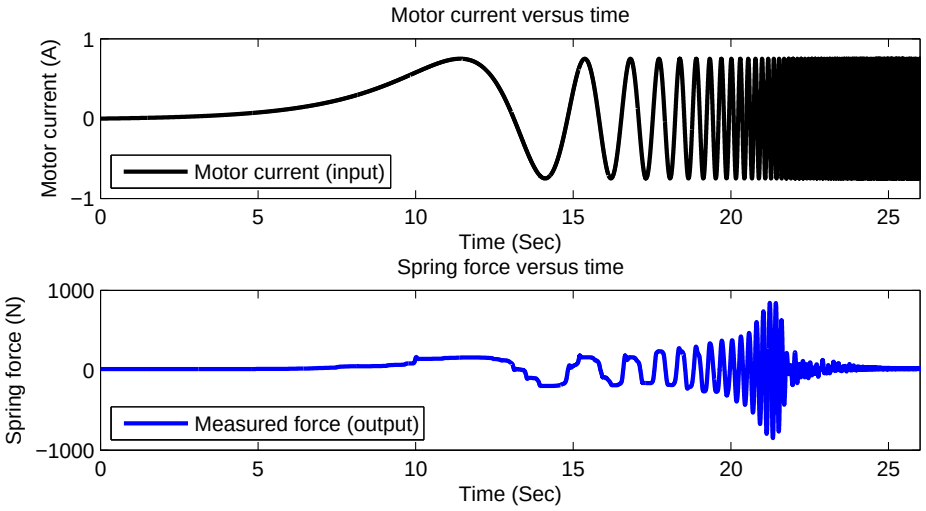

Figure A.3: UT-SEA system identification data.

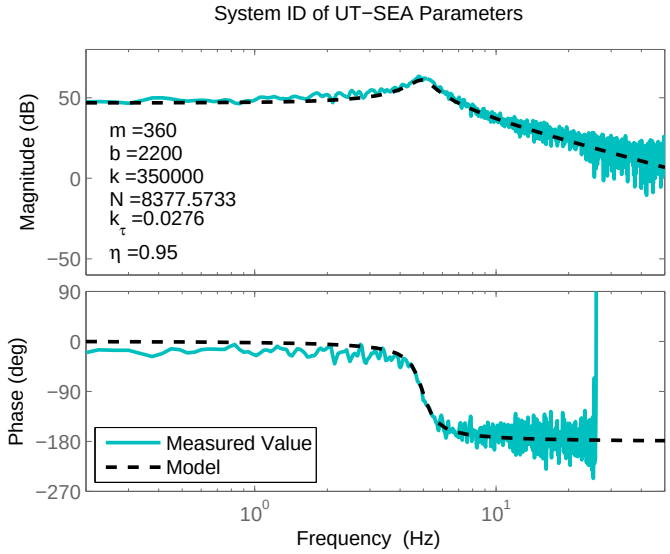

Figure A.4: UT-SEA system identification fitting.

```
function h = getFreqResponse(out, in)

 %fft
 fdout = fft(out);
 fdin = fft(in);

 %divide
 h = fdout./fdin;

end
```

The three parameters are manually tuned to fit the data. Increasing the $k_\tau \eta$ term increases the steady state value, increasing m_k lowers the resonant frequency, and increasing b_{eff} increases the damping ratio.

A.2 Rigid Actuators

Measuring m and b (see Figure A.5) requires measurement of force and displacement, or one of its derivatives. Force measurement may be approximated using current control of electric motors with knowledge of the mapping function between currents and forces (G in (7.7) in Chapter 7), or it may be measured accurately using a load cell.

We begin by implementing an arbitrarily stiff impedance controller (that of Figure 7.2 in Chapter 7 for example). We then send a chirp signal as a desired position setpoint. By measuring both force and velocity (ensuring the velocity cutoff frequency is greater than the chirp frequency), we can plot the frequency response of these two signals. Then, as shown in Figure A.6,

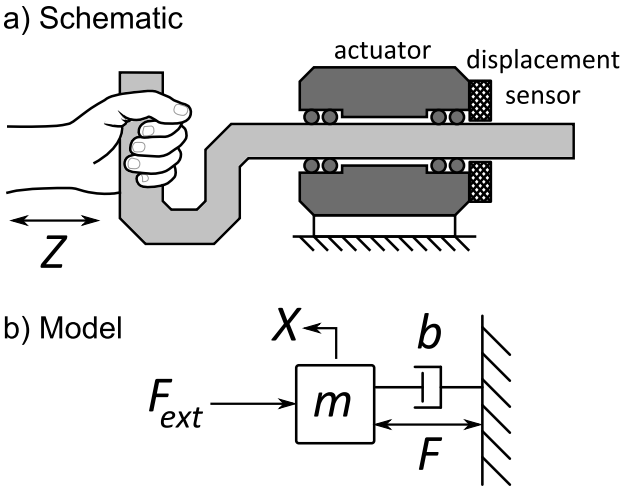

Figure A.5: **a)** An actuator impedance interaction with a human. **b)** Model of actuator including forces from external sources.

an inertia model may be fit to the measured data using the transfer function $\frac{V(s)}{F(s)} = \frac{1}{ms}$.

The b parameter may be measured in a similar way, but depending on the amount of damping and the permissible range of motion, it can be difficult to obtain acceptable data in the lower frequency range. As an alternative we use a closed-loop approach. As before, we implement an impedance controller, but this time we set $B = 0$ and set K to some arbitrarily large value such that it produces an underdamped response. If we apply a step in desired position, we may observe the underdamped response and fit this response to a model of the closed-loop system (see Figure A.7). The fitting goal should be to produce the same level of overshoot in the model as observed in the experiment.

Using these techniques, we obtained system parameters for the two

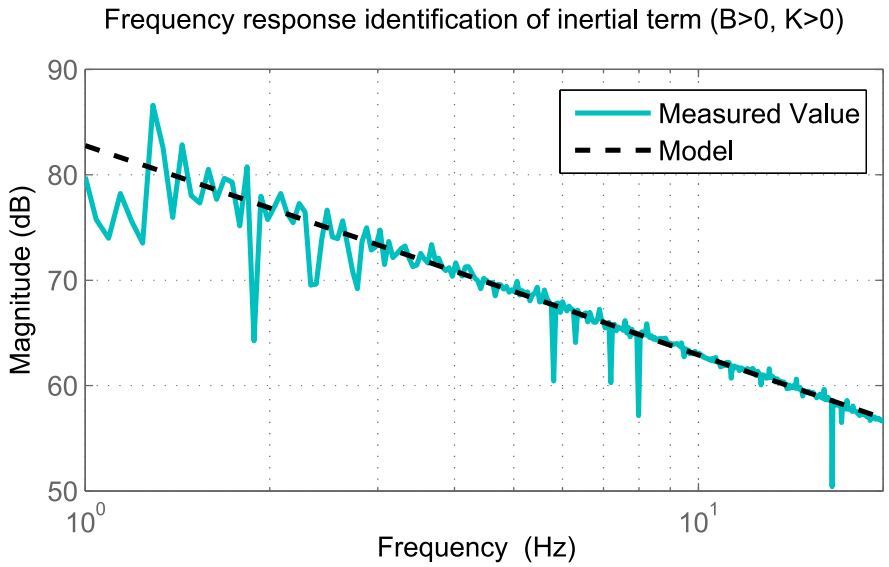

Figure A.6: Experimental identification of inertial (m) model parameter.

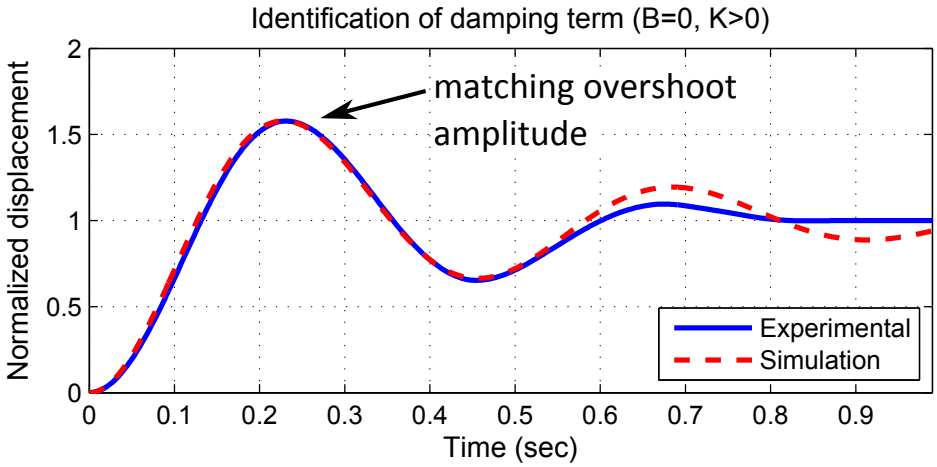

Figure A.7: Experimental identification of damping (b) model parameter using a step response. The deviation seen after 0.5 seconds may be attributed to the presence of unmodeled friction.

Table A.1: Identified system parameters of rigid actuator

Parameter	Standalone motor	Full actuator
m	3.0e-6 $[Kg \cdot m^2]$	256 $[Kg]$
b	3.5e-6 $[Nm \cdot s/rad]$	1250 $[N \cdot s/m]$

actuation systems as shown in Table A.1. The data for the full actuator is in the linear output reference frame.

Bibliography

F. Aghili, J. Hollerbach, and M. Buehler. A modular and high-precision motion control system with an integrated motor. *Mechatronics, IEEE/ASME Transactions on*, 12(3):317–329, June 2007. ISSN 1083-4435. doi: 10.1109/TMECH.2007.897273.

A. Albu-Schäffer, S. Haddadin, C. Ott, A. Stemmer, T. Wimbck, and G. Hirzinger. The DLR lightweight robot: design and control concepts for robots in human environments. *Industrial Robot: An International Journal*, 34(5):376 – 385, 2007a.

A. Albu-Schäffer, C. Ott, and G. Hirzinger. A unified passivity-based control framework for position, torque and impedance control of flexible joint robots. *Int. J. Robot. Res.*, 26(1):23–29, Jan. 2007b.

R. O. Ambrose, R. R. Burridge, W. Bluethmann, H. Aldridge, R. S. Askew, F. Rehnmark, M. Diftler, D. Magruder, and C. Lovchik. Robonaut: NASA's space humanoid. *IEEE Intelligent Systems*, 15(4):57–63, 2000.

J. An and D.-S. Kwon. In haptics, the influence of the controllable physical damping on stability and performance. In *Intelligent Robots and Systems, Proceedings. IEEE/RSJ International Conference on*, volume 2, pages 1204–1209, Sept 2004. doi: 10.1109/IROS.2004.1389560.

S. Au, J. Weber, and H. Herr. Biomechanical design of a powered ankle-foot prosthesis. In *Rehabilitation Robotics, ICORR. IEEE 10th International Conference on*, pages 298–303, June 2007. doi: 10.1109/ICORR.2007.4428441.

J. Bao and P. L. Lee. *Process Control: The Passive Systems Approach*. Springer, 2007.

E. Basafa, M. Sheikholeslami, A. Mirbagheri, F. Farahmand, and G. R. Vossoughi. Design and implementation of series elastic actuators for a haptic laparoscopic device. In *Engineering in Medicine and Biology Society, EMBC. Annual International Conference of the IEEE*, pages 6054–6057, Sept 2009. doi: 10.1109/IEMBS.2009.5332616.

B. Bauml, O. Birbach, T. Wimbock, U. Frese, A. Dietrich, and G. Hirzinger. Catching flying balls with a mobile humanoid: System overview and design considerations. In *Humanoid Robots (Humanoids), 11th IEEE-RAS International Conference on*, pages 513–520, Oct 2011. doi: 10.1109/Humanoids.2011.6100837.

A. Beelen and A. J. Sargeant. Effect of fatigue on maximal power output at different contraction velocities in humans. *Journal of Applied Physiology*, 71(6):2332–2337, 1991.

P. R. Bèlanger. *Control Engineering: A Modern Approach*. Saunders College Publishing, Orlando, Florida, 1995.

W. Bluethmann, R. Ambrose, M. Diftler, S. Askew, E. Huber, M. Goza, F. Rehnmark, C. Lovchik, and D. Magruder. Robonaut: A robot designed to work with humans in space. *Autonomous Robots*, 14(2-3):179–197, 2003.

Boston Dynamics. Atlas - the agile anthropomorphic robot. 2013. `http://www.bostondynamics.com/robot_Atlas.html` Accessed: 2014-07-1.

L. B. Bridgwater, C. A. Ihrke, M. A. Diftler, M. E. Abdallah, N. A. Radford, J. Rogers, S. Yayathi, R. S. Askew, and D. M. Linn. The Robonaut 2 hand-designed to do work with tools. In *Robotics and Automation (ICRA), IEEE International Conference on*, pages 3425–3430. IEEE, 2012.

C. R. Carignan and K. R. Cleary. Closed-loop force control for haptic simulation of virtual environments. *Haptics-e*, 1(2):01–1, 2000.

J. E. Colgate and J. M. Brown. Factors affecting the Z-Width of a haptic display. In *Robotics and Automation, Proceedings., IEEE International Conference on*, volume 4, pages 3205–3210, May 1994. doi: 10.1109/ROBOT.1994.351077.

J. E. Colgate and G. Schenkel. Passivity of a class of sampled-data systems: Application to haptic interfaces. In *American Control Conference*, volume 3, pages 3236–3240, June 1994. doi: 10.1109/ACC.1994.735172.

S. Curran and D. E. Orin. Evolution of a jump in an articulated leg with series-elastic actuation. In *Robotics and Automation, ICRA IEEE International*

Conference on, pages 352 –358, May 2008. doi: 10.1109/ROBOT.2008.4543233.

S. Curran, B. T. Knox, J. P. Schmiedeler, and D. E. Orin. Design of series-elastic actuators for dynamic robots with articulated legs. *Journal of Mechanisms and Robotics*, 1(1), 2008.

DARPA. How do you turn 10 minutes of power into 200? Efficiency, efficiency, efficiency. 2012. `http://www.darpa.mil/NewsEvents/Releases/2012/07/02.aspx` Accessed: 2014-06-29.

DARPA. The DARPA robotics challenge. 2013. `http://www.theroboticschallenge.org` Accessed: 2014-06-27.

M. A. Diftler, J. S. Mehling, M. E. Abdallah, N. A. Radford, L. B. Bridgwater, A. M. Sanders, R. S. Askew, D. M. Linn, J. D. Yamokoski, F. A. Permenter, B. K. Hargrave, R. Piatt, R. T. Savely, and R. O. Ambrose. Robonaut 2 - the first humanoid robot in space. In *Robotics and Automation (ICRA), IEEE International Conference on*, pages 2178 –2183, May 2011. doi: 10.1109/ICRA.2011.5979830.

N. Diolaiti, G. Niemeyer, F. Barbagli, and J. K. Salisbury. Stability of haptic rendering: Discretization, quantization, time delay, and coulomb effects. *Robotics, IEEE Transactions on*, 22(2):256–268, April 2006. ISSN 1552-3098. doi: 10.1109/TRO.2005.862487.

J. Dodge. Power MOSFET tutorial. Technical report, Advanced Power Technology, March 2006.

W. M. dos Santos, G. A. P. Caurin, and A. A. G. Siqueira. Torque control characterization of a rotary series elastic actuator for knee rehabilitation. In *Advanced Robotics (ICAR), 16th International Conference on*, pages 1–6, Nov 2013. doi: 10.1109/ICAR.2013.6766567.

A. Edsinger-Gonzales and J. Weber. Domo: a force sensing humanoid robot for manipulation research. In *Humanoid Robots, 4th IEEE/RAS International Conference on*, volume 1, pages 273 – 291, Nov. 2004. doi: 10.1109/ICHR.2004.1442127.

Energizer. Battery internal resistance. Technical report, December 2005.

Fairchild Semiconductor. Understanding power MOSFETS. Technical report, 1999.

Festo. Fluidic muscle DMSP/MAS. 2014. `http://www.festo.com/cat/en-us_us/data/doc_enus/PDF/US/DMSP-MAS_ENUS.PDF` Accessed: 2014-06-30.

C. Fitzgerald. Developing baxter. In *Technologies for Practical Robot Applications (TePRA), IEEE International Conference on*, pages 1–6, April 2013. doi: 10.1109/TePRA.2013.6556344.

G. F. Franklin, J. D. Powell, and A. E. Baeini. *Feedback Control of Dynamic Systems*. Addison-Wesley, Reading, MA, 1986.

G. Gabrielli and T. H. von Karman. What price speed? *Mech. Eng.*, 72(10): 775781, 1950.

E. Garcia, J. C. Arevalo, F. Sanchez, J. F. Sarria, and P. Gonzalez-de Santos. Design and development of a biomimetic leg using hybrid actuators. In *Intelligent Robots and Systems (IROS), IEEE/RSJ International Conference on*, pages 1507 –1512, Sept. 2011. doi: 10.1109/IROS.2011.6094429.

P. Garrec. Design of an anthropomorphic upper limb exoskeleton actuated by ball-screws and cables. In *Scientific Bulletin, Series D*, volume 72, pages 23–34, 2010.

R. B. Gillespie, D. Kim, J. M. Suchoski, Y. Bo, and J. D. Brown. Series elasticity for free free-space motion for free. In *Haptics Symposium (HAPTICS), IEEE*, pages 609–615, Feb 2014. doi: 10.1109/HAPTICS.2014.6775525.

I. Godler, M. Inoue, T. Ninomiya, and T. Yamashita. Robustness comparison of control schemes with disturbance observer and with acceleration control loop. In *Industrial Electronics, ISIE Proceedings of the IEEE International Symposium on*, volume 3, pages 1035–1040, 1999a. doi: 10.1109/ISIE.1999.796770.

I. Godler, T. Ninomiya, M. Horiuchi, and M. Hashimoto. Ripple compensation of harmonic drive built-in torque sensing. In *Industrial Electronics, ISIE. Proceedings of the IEEE International Symposium on*, volume 2, pages 888–892, 1999b. doi: 10.1109/ISIE.1999.798731.

M. Goldfarb and T. Sirithanapipat. The effect of actuator saturation on the performance of PD-controlled servo systems. *Mechatronics*, 9:497–511, 1999.

P. Gregorio, M. Ahmadi, and M. Buehler. Design, control, and energetics of an electrically actuated legged robot. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 27(4):626 –634, Aug 1997. ISSN 1083-4419. doi: 10.1109/3477.604106.

J. A. Grimes and J. W. Hurst. The design of atrias 1.0 a unique monopod, hopping robot. In *In International Conference on Climbing and Walking Robots*, 2012.

M. Grun, R. Muller, and U. Konigorski. Model based control of series elastic actuators. In *Biomedical Robotics and Biomechatronics (BioRob), 4th IEEE RAS EMBS International Conference on*, pages 538–543, June 2012. doi: 10.1109/BioRob.2012.6290300.

B. Hannaford and J. Winters. Actuator properties and movement control: Biological and technological models. In J. Winters and S.-Y. Woo, editors, *Multiple Muscle Systems*, pages 101–120. Springer New York, 1990. ISBN 978-1-4613-9032-9. doi: 10.1007/978-1-4613-9030-5_7. URL `http://dx.doi.org/10.1007/978-1-4613-9030-5_7`.

Harmonic Drive. Harmonic drive catalog. 2014. `http://harmonicdrive.net/media/support/catalogs/pdf/csf-csg-catalog.pdf` Accessed: 2014-06-27.

M. D. Hasankola, M. M. Moghaddam, A. Dashkhaneh, and A. M. Saba. Design of a rotary elastic actuator for use as torque-actuator in rehabilitation robots. In *Robotics and Mechatronics (ICRoM), First RSI/ISM International Conference on*, pages 505–510, Feb 2013. doi: 10.1109/ICRoM.2013.6510157.

A. Hereid, S. Kolathaya, M. S. Jones, J. V. Why, J. W. Hurst, and A. D. Ames. Dynamic multi-domain bipedal walking with atrias through slip based human-inspired control. In *Hybrid Systems and Control Conference (HSCC)*, April 2014.

W. K. Ho, K. W. Lim, and W. Xu. Optimal gain and phase margin tuning for PID controllers. *Automatica*, 34(8):1009–1014, 1998.

N. Hogan. Impedance control: An approach to manipulation. In *American Control Conference*, pages 304 –313, june 1984.

N. Hogan. Impedance control: An approach to manipulation: Part I-Theory. *Journal of Dynamic Systems, Measurement, and Control*, 107:1–7, 1985.

J. M. Hollerbach, I. W. Hunter, and J. Ballantyne. A comparative analysis of actuator technologies for robotics. In O. Khatib, J. J. Craig, and T. Lozano-Pérez, editors, *The Robotics Review 2*. MIT Press, Cambridge, MA, USA, 1992.

Honda. Honda unveils all-new ASIMO with significant advancements.

2011. http://world.honda.com/news/2011/c111108All-new-ASIMO/index.html Accessed: 2014-06-27.

T. Hulin, C. Preusche, and G. Hirzinger. Stability boundary for haptic rendering: Influence of physical damping. In *Intelligent Robots and Systems, IEEE/RSJ International Conference on*, pages 1570–1575, Oct 2006. doi: 10.1109/IROS.2006.282043.

I. W. Hunter and S. Lafontaine. A comparison of muscle with artificial actuators. In *Solid-State Sensor and Actuator Workshop, 5th Technical Digest., IEEE*, pages 178–185, June 1992. doi: 10.1109/SOLSEN.1992.228297.

J. W. Hurst and A. A. Rizzi. Series compliance for an efficient running gait. *Robotics Automation Magazine, IEEE*, 15(3):42–51, September 2008. ISSN 1070-9932. doi: 10.1109/MRA.2008.927693.

J. W. Hurst, J. E. Chestnutt, and A. A. Rizzi. An actuator with physically variable stiffness for highly dynamic legged locomotion. In *Robotics and Automation, Proceedings. ICRA. IEEE International Conference on*, volume 5, pages 4662–4667, April 2004. doi: 10.1109/ROBOT.2004.1302453.

J. W. Hurst, J. E. Chestnutt, and A. A. Rizzi. The actuator with mechanically adjustable series compliance. *Robotics, IEEE Transactions on*, 26(4):597 – 606, Aug. 2010. ISSN 1552-3098. doi: 10.1109/TRO.2010.2052398.

M. Hutter. *StarlETH & Co. - Design and Control of Legged Robots with Compliant Actuation*. PhD thesis, ETH Zurich, 2013.

M. Hutter, C. Remy, and R. Siegwart. Design of an articulated robotic leg with nonlinear series elastic actuation. In *Proc. of The 12th International Conference on Climbing and Walking Robots and the Support Technologies for Mobile Machines (CLAWAR)*, pages 645–652, September 2009.

M. Hutter, C. Remy, M. Hoepflinger, and R. Siegwart. High compliant series elastic actuation for the robotic leg ScarlETH. In *Proc. of the International Conference on Climbing and Walking Robots (CLAWAR)*, 2011a.

M. Hutter, C. Remy, M. Hoepflinger, and R. Siegwart. ScarlETH: Design and control of a planar running robot. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2011b.

M. Hutter, C. D. Remy, M. A. Hoepflinger, and R. Siegwart. Efficient and versatile locomotion with highly compliant legs. *Mechatronics, IEEE/ASME Transactions on*, 18(2):449–458, April 2013. ISSN 1083-4435. doi: 10.1109/TMECH.2012.2222430.

IEEE Spectrum. DARPA robotics challenge trials: Final results. 2013. `http://spectrum.ieee.org/automaton/robotics/humanoids/darpa-robotics-challenge-trials-results` Accessed: 2014-06-29.

N. Ito, J. Urata, Y. Nakanishi, K. Okada, and M. Inaba. Development of very small high output motor driver for realizing forceful musculoskeletal humanoids. In *Humanoid Robots (Humanoids), 10th IEEE-RAS International Conference on*, pages 385–390, Dec 2010. doi: 10.1109/ICHR.2010.5686268.

Y. Ito, S. Nozawa, J. Urata, T. Nakaoka, K. Kobayashi, Y. Nakanishi, K. Okada, and M. Inaba. Development and verification of life-size humanoid with high-output actuation system. In *Robotics and Automation (ICRA), IEEE International Conference on*. IEEE, 2014.

K. Kaneko, F. Kanehiro, M. Morisawa, K. Akachi, G. Miyamori, A. Hayashi, and N. Kanehira. Humanoid robot HRP-4 - humanoid robotics platform with lightweight and slim body. In *Intelligent Robots and Systems (IROS), IEEE/RSJ International Conference on*, pages 4400–4407, Sept 2011. doi: 10.1109/IROS.2011.6094465.

S. Kawamura, F. Miyazaki, and S. Arimoto. Is a local linear PD feedback control law effective for trajectory tracking of robot motion? In *Robotics and Automation, Proceedings., IEEE International Conference on*, volume 3, pages 1335–1340, Apr 1988. doi: 10.1109/ROBOT.1988.12253.

O. Khatib. A unified approach for motion and force control of robot manipulators: The operational space formulation. *Robotics and Automation, IEEE Journal of*, 3(1):43 –53, february 1987. ISSN 0882-4967. doi: 10.1109/JRA.1987.1087068.

J. Kim, H. Kwak, H. Lee, K. Seo, B. Lim, M. Lee, J. Lee, and K. Roh. Balancing control of biped robot. In *IEEE international Conference on Systems, Man, and Cybernetics*, Oct 2012.

K. Kong and M. Tomizuka. Nominal model manipulation for enhancement

of stability robustness for disturbance observer-based control systems. *International Journal of Control, Automation, and Systems*, 11(1):12 – 20, 2013.

K. Kong, J. Bae, and M. Tomizuka. Control of rotary series elastic actuator for ideal force-mode actuation in human-robot interaction applications. *Mechatronics, IEEE/ASME Transactions on*, 14(1):105 –118, Feb. 2009. ISSN 1083-4435. doi: 10.1109/TMECH.2008.2004561.

K. Kong, J. Bae, and M. Tomizuka. A compact rotary series elastic actuator for knee joint assistive system. In *Robotics and Automation (ICRA), IEEE International Conference on*, pages 2940 –2945, May 2010. doi: 10.1109/ROBOT.2010.5509227.

K. Kong, J. Bae, and M. Tomizuka. A compact rotary series elastic actuator for human assistive systems. *Mechatronics, IEEE/ASME Transactions on*, 17(2):288 –297, April 2012. ISSN 1083-4435. doi: 10.1109/TMECH.2010.2100046.

A. Kuo. Choosing your steps carefully. *Robotics Automation Magazine, IEEE*, 14(2):18–29, June 2007. ISSN 1070-9932. doi: 10.1109/MRA.2007.380653.

J. Kuria and P. Hwang. Modeling power losses in electric vehicle bldc motor. *Journal of Energy Technologies and Policy*, 1(4), 2011.

H. K. Kwa, J. H. Noorden, M. Missel, T. Craig, J. E. Pratt, and P. D. Neuhaus. Development of the ihmc mobility assist exoskeleton. In *Robotics and Au-*

tomation, ICRA. IEEE International Conference on, pages 2556–2562, May 2009. doi: 10.1109/ROBOT.2009.5152394.

C. Lagoda, A. C. Schouten, A. H. A. Stienen, E. E. G. Hekman, and H. van der Kooij. Design of an electric series elastic actuated joint for robotic gait rehabilitation training. In *Biomedical Robotics and Biomechatronics (BioRob), 3rd IEEE RAS and EMBS International Conference on*, pages 21 –26, Sept. 2010. doi: 10.1109/BIOROB.2010.5626010.

D. Lancaster. Some energy fundamentals. *The Blatant Opportunist*, pages 1–7, Oct. 2002.

D. A. Lawrence. Impedance control stability properties in common implementations. In *Robotics and Automation, Proceedings., IEEE International Conference on*, volume 2, pages 1185–1190, Apr 1988. doi: 10.1109/ROBOT.1988.12222.

D. A. Lawrence. Actuator limitations on achievable manipulator impedance. In *Robotics and Automation, Proceedings., IEEE International Conference on*, volume 1, pages 560–565, May 1989. doi: 10.1109/ROBOT.1989.100044.

B. Lee, V. Orekhov, D. Lahr, and D. Hong. Design and measurement error analysis of a low-friction, lightweight linear series elastic actuator. In *International Design Engineering Technical Conferences & Computers and Information in Engineering Conference (IDETC/CIE)*, 2013.

C.-H. Lee. A survey of PID controller design based on gain and phase margins. *International Journal of Computational Cognition*, 2:63–100, 2004.

C.-H. Lee and C.-C. Teng. Calculation of PID controller parameters by using a fuzzy neural network. *ISA Transactions*, 42:391–400, 2003.

T. Lens and O. von Stryk. Design and dynamics model of a lightweight series elastic tendon-driven robot arm. In *Robotics and Automation (ICRA), IEEE International Conference on*, pages 4512–4518, May 2013. doi: 10.1109/ICRA.2013.6631218.

D. Li, F. Gao, Y. Xue, and C. Lu. Optimization of decentralized PI/PID controllers based on genetic algorithm. *Asian Journal of Control*, 9(3):306–316, 2007. ISSN 1934-6093. doi: 10.1111/j.1934-6093.2007.tb00416.x. URL `http://dx.doi.org/10.1111/j.1934-6093.2007.tb00416.x`.

K. Li. PID tuning for optimal closed-loop performance with specified gain and phase margins. *Control Systems Technology, IEEE Transactions on*, 21(3):1024–1030, May 2013. ISSN 1063-6536. doi: 10.1109/TCST.2012.2198479.

Z. Li, N. G. Tsagarakis, and D. G. Caldwell. A passivity based admittance control for stabilizing the compliant humanoid coman. In *Humanoid Robots (Humanoids), 12th IEEE-RAS International Conference on*, pages 43–49, Nov 2012. doi: 10.1109/HUMANOIDS.2012.6651497.

Y.-S. Lu, S.-M. Lin, M. Hauschild, and G. Hirzinger. A torque-ripple compensation scheme for harmonic drive systems. *Electrical Engineering*, 95

(4):357–365, 2013. ISSN 0948-7921. doi: 10.1007/s00202-012-0264-4. URL http://dx.doi.org/10.1007/s00202-012-0264-4.

T. H. Massie and J. K. Salisbury. The phantom haptic interface: A device for probing virtual objects. In *Proceedings of the ASME Dynamic Systems and Control Division*, pages 295–301, 1994.

Maxon. EC-4pole 30. 2014a. http://www.maxonmotorusa.com/medias/sys_master/8813562167326/14-214-EN.pdf Accessed: 2014-06-27.

Maxon. Product catalog 2014-2015. 2014b. http://epaper.maxonmotor.ch/en/ Accessed: 2014-07-1.

Maxon. EC 22. 2014c. http://www.maxonmotorusa.com/medias/sys_master/8813542670366/14-190-EN.pdf Accessed: 2014-07-1.

J. S. Mehling, J. E. Colgate, and M. A. Peshkin. Increasing the impedance range of a haptic display by adding electrical damping. In *Eurohaptics Conference, Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 257–262, March 2005. doi: 10.1109/WHC.2005.79.

K. Nagatani, S. Kiribayashi, Y. Okada, K. Otake, K. Yoshida, S. Tadokoro, T. Nishimura, T. Yoshida, E. Koyanagi, M. Fukushima, and S. Kawatsuma. Emergency response to the nuclear accident at the Fukushima Daiichi Nuclear Power Plants using mobile rescue robots. *Journal of Field*

Robotics, 30(1):44–63, 2013. ISSN 1556-4967. doi: 10.1002/rob.21439. URL http://dx.doi.org/10.1002/rob.21439.

M. Nakao, K. Ohnishi, and K. Miyachi. A robust decentralized joint control based on interference estimation. In *Robotics and Automation. Proceedings. IEEE International Conference on*, volume 4, pages 326–331, 1987. doi: 10.1109/ROBOT.1987.1087996.

NASA-JSC. NASA-JSC DRC team valkyrie. 2013. http://www.theroboticschallenge.org/node/59 Accessed: 2014-02-10.

R. Niiyama, A. Nagakubo, and Y. Kuniyoshi. Mowgli: A bipedal jumping and landing robot with an artificial musculoskeletal system. In *Robotics and Automation, IEEE International Conference on*, pages 2546–2551, April 2007. doi: 10.1109/ROBOT.2007.363848.

R. Niiyama, S. Nishikawa, and Y. Kuniyoshi. Athlete robot with applied human muscle activation patterns for bipedal running. In *Humanoid Robots (Humanoids), 10th IEEE-RAS International Conference on*, pages 498–503, Dec 2010. doi: 10.1109/ICHR.2010.5686316.

J. Oblak and Z. Matjačić. On stability and passivity of haptic devices characterized by a series elastic actuation and considerable end-point mass. In *Rehabilitation Robotics (ICORR), IEEE International Conference on*, pages 1–5, June 2011. doi: 10.1109/ICORR.2011.5975497.

K. Ogata. *Modern Control Engineering*. Prentice-Hall, Englewood Cliffs, NJ, 1990.

C. Ott, A. Albu-Schäffer, A. Kugi, and G. Hirzinger. Decoupling based cartesian impedance control of flexible joint robots. In *Robotics and Automation, Proceedings. IEEE International Conference on*, volume 3, pages 3101–3107, Sept 2003. doi: 10.1109/ROBOT.2003.1242067.

C. Ott, A. Kugi, and G. Hirzinger. On the passivity-based impedance control of flexible joint robots. *Robotics, IEEE Transactions on*, 24(2):416–429, 2008. ISSN 1552-3098. doi: 10.1109/TRO.2008.915438.

C. Ott, M. A. Roa, and G. Hirzinger. Posture and balance control for biped robots based on contact force optimization. In *Humanoid Robots (Humanoids), 11th IEEE-RAS International Conference on*, pages 26–33, Oct 2011. doi: 10.1109/Humanoids.2011.6100882.

N. Paine and L. Sentis. A new prismatic series elastic actuator with compact size and high performance. In *Robotics and Biomimetics (ROBIO), IEEE International Conference on*, pages 1759–1766, Dec 2012. doi: 10.1109/ROBIO.2012.6491222.

N. Paine and L. Sentis. A closed-form solution for selecting maximum critically damped actuator impedance parameters. *Journal of Dynamic Systems, Measurement, and Control (In Press)*, 2014.

N. Paine, J. S. Mehling, J. Holley, N. Radford, G. Johnson, C. Fok, and L. Sentis. Actuator Control for the NASA-JSC Valkyrie Humanoid Robot: A Decoupled Dynamics Approach for Torque Control of Series Elastic Robots. *Journal of Field Robotics (In Press)*, 2014a.

N. Paine, S. Oh, and L. Sentis. Design and control considerations for high-performance series elastic actuators. *Mechatronics, IEEE/ASME Transactions on*, 19(3):1080–1091, June 2014b. ISSN 1083-4435. doi: 10.1109/TMECH.2013.2270435.

D. Paluska and H. Herr. Series elasticity and actuator power output. In *Robotics and Automation, ICRA Proceedings IEEE International Conference on*, pages 1830 –1833, May 2006. doi: 10.1109/ROBOT.2006.1641972.

A. Parmiggiani, G. Metta, and N. Tsagarakis. The mechatronic design of the new legs of the icub robot. In *Humanoid Robots (Humanoids), 12th IEEE-RAS International Conference on*, pages 481–486, Nov 2012. doi: 10.1109/HUMANOIDS.2012.6651563.

J. Pestana, R. Bobin, J. C. Arevalo, and E. Garcia Armada. Characterization of emerging actuators for empowering legged robots. In *CLAWAR*, 2010.

P. P. Pott, H. Graefenstein, J. Fischer, R. Mueller, H. F. Schlaak, and E. Abele. Series elastic actuators for man-machine cooperation. In *Innovative Small Drives and Micro-Motor Systems, 9. GMM/ETG Symposium*, pages 1–5, Sept 2013.

E. Poulin, A. Pomerleau, A. Desbiens, and D. Hodouin. Development and evaluation of an auto-tuning and adaptive PID controller. *Automatica*, 32 (1):71–82, 1996.

G. A. Pratt and M. M. Williamson. Series elastic actuators. In *Intelligent Robots and Systems. 'Human Robot Interaction and Cooperative Robots', Proceedings. IEEE/RSJ International Conference on*, volume 1, pages 399 –406, Aug 1995. doi: 10.1109/IROS.1995.525827.

G. A. Pratt, P. Willisson, C. Bolton, and A. Hofman. Late motor processing in low-impedance robots: impedance control of series-elastic actuators. In *American Control Conference, Proceedings of the*, volume 4, pages 3245 – 3251, July 2004.

J. E. Pratt and B. T. Krupp. Series elastic actuators for legged robots. *Proc. SPIE 5422, Unmanned Ground Vehicle Technology VI*, pages 135–144, 2004. doi: 10.1117/12.548000.

J. E. Pratt and G. A. Pratt. Intuitive control of a planar bipedal walking robot. In *Robotics and Automation, Proceedings. IEEE International Conference on*, volume 3, pages 2014 –2021, May 1998. doi: 10.1109/ROBOT.1998.680611.

J. E. Pratt, P. Dilworth, and G. A. Pratt. Virtual model control of a bipedal walking robot. In *Robotics and Automation, Proceedings. IEEE International Conference on*, volume 1, pages 193–198, Apr 1997. doi: 10.1109/ROBOT.1997.620037.

J. E. Pratt, C.-M. Chew, A. Torres, P. Dilworth, and G. Pratt. Virtual model control: An intuitive approach for bipedal locomotion. *International Journal of Robotics Research*, 20(2):129–143, 2001. doi: 10.1177/02783640122067309.

J. E. Pratt, B. Krupp, and C. Morse. Series elastic actuators for high fidelity force control. *Industrial Robot: An International Journal*, 29(3):234 – 241, 2002.

D. Ragonesi, S. Agrawal, W. Sample, and T. Rahman. Series elastic actuator control of a powered exoskeleton. In *Engineering in Medicine and Biology Society,EMBC, Annual International Conference of the IEEE*, pages 3515 –3518, Sept. 2011. doi: 10.1109/IEMBS.2011.6090583.

M. H. Raibert. *Legged Robots That Balance*. MIT Press, Cambridge, Mass, 1986.

M. H. Raibert, K. Blankespoor, G. Nelson, R. Playter, and T. B. Team. *BigDog, the Rough-Terrain Quadruped Robot*, pages 10822–10825. 2008. URL `http://www.df.unibo.it/divulgazione/attualita/bigdog.pdf`.

M. H. Rashid. *Power Electronics Handbook*. Academic Press, San Diego, California, 2001.

K. J. Åström. Automatic tuning and adaptation for PID controllers - A survey. *Control Eng. Practice*, 1:699–714, 1993.

K. J. Åström, H. Panagopoulos, and T. Hägglund. Design of PI controllers based on non-convex optimization. *Automatica*, 34(5):585–601, May 1998. ISSN 0005-1098. doi: 10.1016/S0005-1098(98)00011-9. URL `http://dx.doi.org/10.1016/S0005-1098(98)00011-9`.

M. J. Reiland, B. Hargrave, R. Platt, M. E. Abdallah, and F. N. Permenter. Architecture for robust force and impedance control of series elastic actuators, Sept. 2013. US Patent 8,525,460.

Remy Motors. Remy HVH250-090-SOM electric motor. 2014. `http://www.remyinc.com/docs/hybrid/REM-03_HVH250_DataSht_master.pdf` Accessed: 2014-06-27.

D. W. Robinson. *Design and Analysis of Series Elasticity in Closed-loop Actuator Force Control*. PhD thesis, Massachusetts Institute of Technology, 2000.

D. W. Robinson, J. E. Pratt, D. J. Paluska, and G. A. Pratt. Series elastic actuator development for a biomimetic walking robot. In *Advanced Intelligent Mechatronics, Proceedings. IEEE/ASME International Conference on*, pages 561 –568, 1999. doi: 10.1109/AIM.1999.803231.

C. Rossa, J. Lozada, and A. Micaelli. Stable haptic interaction using passive and active actuators. In *Robotics and Automation, IEEE International Conference on*, pages 2386–2392, May 2013. doi: 10.1109/ICRA.2013.6630901.

S. Safavi, A. S. Ghafari, and A. Meghdari. Design of an optimum torque actuator for augmenting lower extremity exoskeletons in biomechanical framework. In *Robotics and Biomimetics (ROBIO), IEEE International Conference on*, pages 1979–1983, Dec 2011. doi: 10.1109/ROBIO.2011.6181581.

Y. Sakagami, R. Watanabe, C. Aoyama, S. Matsunaga, N. Higaki, and K. Fujimura. The intelligent asimo: system overview and integration. In *Intelligent Robots and Systems, IEEE/RSJ International Conference on*, volume 3, pages 2478–2483, 2002. doi: 10.1109/IRDS.2002.1041641.

Schaft. Schaft team. 2013. `http://theroboticschallenge.org/teams/schaft` Accessed: 2014-06-27.

A. Schepelmann, M. Taylor, and H. Geyer. Development of a testbed for robotic neuromuscular controllers. In *Proceedings of Robotics: Science and Systems*, 2012.

C. Semini. *HyQ Design and Development of a Hydraulically Actuated Quadruped Robot*. PhD thesis, Italian Institute of Technology and University of Genoa, 2010.

C. Semini, N. G. Tsagarakis, E. Guglielmino, and D. G. Caldwell. Design and experimental evaluation of the hydraulically actuated prototype leg of the HyQ robot. In *Intelligent Robots and Systems (IROS), IEEE/RSJ International Conference on*, pages 3640 –3645, Oct. 2010. doi: 10.1109/IROS.2010.5651548.

J. W. Sensinger and J. H. Lipsey. Cycloid vs. harmonic drives for use in high ratio, single stage robotic transmissions. In *Robotics and Automation (ICRA), IEEE International Conference on*, pages 4130–4135, May 2012. doi: 10.1109/ICRA.2012.6224739.

J. W. Sensinger and R. F. Weir. Unconstrained impedance control using a compact series elastic actuator. In *Mechatronic and Embedded Systems and Applications, Proceedings of the 2nd IEEE/ASME International Conference on*, pages 1 –6, Aug. 2006. doi: 10.1109/MESA.2006.296972.

L. Sentis and B. Fernandez. Perturbation theory to plan dynamic locomotion in very rough terrains. In *Intelligent Robots and Systems (IROS), IEEE/RSJ International Conference on*, pages 2267–2273, Sept 2011. doi: 10.1109/IROS.2011.6094801.

L. Sentis, J. Petersen, and R. Philippsen. Implementation and stability analysis of prioritized whole-body compliant controllers on a wheeled humanoid robot in uneven terrains. *Autonomous Robots*, 35(4):301–319, 2013. ISSN 0929-5593. URL `http://dx.doi.org/10.1007/s10514-013-9358-8`.

S. Seok, A. Wang, D. Otten, and S. Kim. Actuator design for high force proprioceptive control in fast legged locomotion. In *Intelligent Robots and Systems (IROS), IEEE/RSJ International Conference on*, pages 1970–1975, Oct 2012. doi: 10.1109/IROS.2012.6386252.

S. Seok, A. Wang, M. Y. Chuah, D. Otten, J. Lang, and S. Kim. Design principles for highly efficient quadrupeds and implementation on the mit cheetah

robot. In *Robotics and Automation (ICRA), IEEE International Conference on*, pages 3307–3312, May 2013. doi: 10.1109/ICRA.2013.6631038.

F. Sergi, D. Accoto, G. Carpino, N. L. Tagliamonte, and E. Guglielmelli. Design and characterization of a compact rotary series elastic actuator for knee assistance during overground walking. In *Biomedical Robotics and Biomechatronics (BioRob), 4th IEEE RAS EMBS International Conference on*, pages 1931–1936, June 2012. doi: 10.1109/BioRob.2012.6290271.

M. Slovich, N. Paine, K. Kemper, A. Metger, A. Edinger, J. Weber, and L. Sentis. Building HUME: A bipedal robot for human-centered hyper-agility. In *Dynamic Walking*, May 2012.

D. Sourlas, J. Choi, and V. Manousiouthakis. Best achievable control system performance: The saturation paradox. In *Decision and Control, Proceedings of the 33rd IEEE Conference on*, volume 4, pages 3816–3818, Dec 1994. doi: 10.1109/CDC.1994.411754.

M. Spong. Modeling and control of elastic joint robots. *J. Dyn. Sys., Meas., Control.*, 109(4):310–318, 1987.

P. Suchomski. Robust PI and PID controller design in delta domain. *Control Theory and Applications, IEE Proceedings* -, 148(5):350–354, Sep 2001. ISSN 1350-2379. doi: 10.1049/ip-cta:20010627.

M. D. Taylor. A compact series elastic actuator for bipedal robots with human-like dynamic performance. Masters thesis, Carnegie Mellon University, 2011.

I. Thorson and D. Caldwell. A nonlinear series elastic actuator for highly dynamic motions. In *Intelligent Robots and Systems (IROS), IEEE/RSJ International Conference on*, pages 390 –394, Sept. 2011. doi: 10.1109/IROS.2011.6094913.

O. Tokatli and V. Patoglu. Series elastic actuation for force controlled micromanipulation. In *Mechatronics (ICM), IEEE International Conference on*, pages 421–426, April 2011. doi: 10.1109/ICMECH.2011.5971323.

E. Torres-Jara and J. Banks. A simple and scalable force actuator. In *International Simposium of Robotics*, March 2004.

N. G. Tsagarakis, M. Laffranchi, B. Vanderborght, and D. G. Caldwell. A compact soft actuator unit for small scale human friendly robots. In *Robotics and Automation, ICRA. IEEE International Conference on*, pages 4356–4362, May 2009. doi: 10.1109/ROBOT.2009.5152496.

J. Urata, T. Hirose, Y. Namiki, Y. Nakanishi, I. Mizuuchi, and M. Inaba. Thermal control of electrical motors for high-power humanoid robots. In *Intelligent Robots and Systems, IROS. IEEE/RSJ International Conference on*, pages 2047–2052, Sept 2008. doi: 10.1109/IROS.2008.4651110.

J. Urata, Y. Nakanishi, K. Okada, and M. Inaba. Design of high torque and high speed leg module for high power humanoid. In *Intelligent Robots and Systems (IROS), IEEE/RSJ International Conference on*, pages 4497–4502, Oct 2010. doi: 10.1109/IROS.2010.5649683.

H. Vallery, R. Ekkelenkamp, H. van der Kooij, and M. Buss. Passive and accurate torque control of series elastic actuators. In *Intelligent Robots and Systems, IROS IEEE/RSJ International Conference on*, pages 3534 –3538, Nov. 2007. doi: 10.1109/IROS.2007.4399172.

J. F. Veneman, R. Ekkelenkamp, R. Kruidhof, F. van der Helm, and H. van der Kooij. Design of a series elastic- and bowden cable-based actuation system for use as torque-actuator in exoskeleton-type training. In *Rehabilitation Robotics, ICORR. 9th International Conference on*, pages 496–499, June 2005. doi: 10.1109/ICORR.2005.1501150.

D. Villegas, M. Van Damme, B. Vanderborght, P. Beyl, and D. Lefeber. Thirdgeneration pleated pneumatic artificial muscles for robotic applications: Development and comparison with mckibben muscle. *Advanced Robotics*, 26(11-12):1205–1227, 2012. doi: 10.1080/01691864.2012.689722. URL `http://dx.doi.org/10.1080/01691864.2012.689722`.

M. Vukobratović and B. Borovac. Zero-moment point - thirty five years of its life. *International Journal of Humanoid Robotics*, 1(1):157–173, 2004.

C. Wang and D. Li. Decentralized PID controllers based on probabilistic robustness. *Journal of Dynamic Systems, Measurement, and Control*, 133 (6):061015, Nov 2011. ISSN 0022-0434. doi: 10.1115/1.4004781. URL `http://dx.doi.org/10.1115/1.4004781`.

S. Wang, C. Meijneke, and H. van der Kooij. Modeling, design, and optimization of mindwalker series elastic joint. In *Rehabilitation Robotics*

(ICORR), IEEE International Conference on, pages 1–8, June 2013. doi: 10.1109/ICORR.2013.6650381.

D. W. Weir, J. E. Colgate, and M. A. Peshkin. Measuring and increasing Z-width with active electrical damping. In *Haptic interfaces for virtual environment and teleoperator systems, Symposium on*, pages 169–175, March 2008. doi: 10.1109/HAPTICS.2008.4479938.

B. W. Williams. *Power Electronics: Devices, Drivers, Applications, and Passive Components*. Mcgraw-Hill, New York City, New York, September 1992.

T. Wimbock, C. Ott, and G. Hirzinger. Impedance behaviors for two-handed manipulation: Design and experiments. In *Robotics and Automation, IEEE International Conference on*, pages 4182–4189, April 2007. doi: 10.1109/ROBOT.2007.364122.

G. Wyeth. Control issues for velocity sourced series elastic actuators. In *Australian Conference on Robotics and Automation*, December 2006.

O. Yaniv and M. Nagurka. Design of PID controllers satisfying gain margin and sensitivity constraints on a set of plants. *Automatica*, 40(1):111–116, 2004.

T.-J. Yeh, M.-J. Wu, T.-J. Lu, F.-K. Wu, and C.-R. Huang. Control of mck-ibben pneumatic muscles for a power-assist, lower-limb orthosis. *Mechatronics*, 20(6):686–697, 2010.

J. Zhu, Q. Wang, and L. Wang. Pantoe 1: Biomechanical design of powered ankle-foot prosthesis with compliant joints and segmented foot. In *Advanced Intelligent Mechatronics (AIM), IEEE/ASME International Conference on*, pages 31–36, July 2010. doi: 10.1109/AIM.2010.5695879.

J. G. Ziegler and N. B. Nichols. Optimum settings for automatic controllers. *Trans. ASME*, 140(3):759–768, 1942.

A. Zoss, H. Kazerooni, and A. Chu. Biomechanical design of the berkeley lower extremity exoskeleton (BLEEX). *Mechatronics, IEEE/ASME Transactions on*, 11(2):128 –138, April 2006. ISSN 1083-4435. doi: 10.1109/TMECH.2006.871087.

Vita

Nicholas Arden Paine is a native of Austin, Tx. He attended elementary, middle, and high school in Dripping Springs, Tx. In the fall of 2004 he attended the Electrical and Computer Engineering program at the University of Texas at Austin, and received the Bachelor's and Master's degrees in Electrical Engineering in 2008 and 2010, respectively. From 2008 to 2010 he was employed as a graduate research assistant in the Pharos Robotics Lab at U.T. Austin. From 2010 to 2014 he worked in the Human Centered Robotics Lab. Upon graduating with a Doctoral degree, he plans to pursue work designing high-performance mobile robots in industry.

Contact email: n.a.paine@gmail.com

This dissertation was typeset with LaTeX[†] by the author.

[†]LaTeX is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's TeX Program.