Series Elasticity in Linearly Actuated Humanoids

Viktor Leonidovich Orekhov

Dissertation submitted to the faculty of the Virginia Polytechnic Institute and State University in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
In
Mechanical Engineering

Dennis W. Hong, Chair
Daniel M. Dudek
Brian Y. Lattimer
Alexander Leonessa
Robert H. Sturges

December 8, 2014
Blacksburg, VA

Keywords: Series Elastic Actuators, Compliant Actuators, Configurable Compliance, Actuator Model, Humanoid Robots

Series Elasticity in Linearly Actuated Humanoids

Viktor Leonidovich Orekhov

ABSTRACT

Recent advancements in actuator technologies, computation, and control have led to major leaps in capability and have brought humanoids ever closer to being feasible solutions for real-world applications. As the capabilities of humanoids increase, they will be called on to operate in unstructured real world environments. This realization has driven researchers to develop more dynamic, robust, and adaptable robots.

Compared to state-of-the-art robots, biological systems demonstrate remarkably better efficiency, agility, adaptability, and robustness. Many recent studies suggest that a core principle behind these advantages is compliance, yet there are very few compliant humanoids that have demonstrated successful walking.

The work presented in this dissertation is based on several years of developing novel actuators for two full-scale linearly actuated compliant humanoid robots, SAFFiR and THOR. Both are state-of-the-art robots intended to operate in the extremely challenging real world scenarios of shipboard firefighting and disaster response.

The design, modeling, and control of actuators in robotics application is critical because the rest of the robot is often designed around the actuators. This dissertation seeks to address two goals: 1) advancing the design of compliant linear actuators that are well suited for humanoid applications, and 2) developing a better understanding of how to design and model compliant linear actuators for use in humanoids.

Beyond just applications for compliant humanoids, this research tackles many of the same design and application challenges as biomechanics research so it has many potential applications in prosthetics, exoskeletons, and rehabilitation devices.

Acknowledgements

The work represented in this dissertation would not have been possible without the encouragement and support of so many people.

Most importantly, I want to thank my parents, Leonid and Yelena, whose countless sacrifices and relentless work ethic have made it possible for me to pursue this degree. I'm thankful for my older brother, Vitaliy, who led the way into engineering and whose coattails I've been riding from the very beginning. To my younger siblings, Galina and Andrew, thanks for your consistent support and interest in my work, it's a bigger deal than you probably realize.

I want to thank my committee members, Brian Lattimer, Alexander Leonessa, Robert Sturges, and Daniel Dudek for your time and advice towards improving my research efforts over the years. I especially want to thank my advisor, Dennis Hong, for your many years of support and for your contagious passion for robotics. It's truly been a privilege to "work" in the creative and inspiring environment you cultivated in RoMeLa.

To all of my labmates in RoMeLa and TREC, it has been an honor and distinct privilege to work with each of you. I doubt that I'll ever find another group that is as talented and enjoyable to work with. There are too many to list everybody by name, but the core group during the early development of SAFFiR and THOR deserves special recognition for putting up with me the most: Derek Lahr, Mike Hopkins, Bryce Lee, Steve Ressler, Coleman Knabe, Jake Webb, and Jack Newton. Robots are hard. But even the setbacks and late nights are tolerable when you work among friends. I'm proud of what we were able to accomplish together.

Finally, I'd like to thank the Graduates and Professionals (GAP) group and the entire Northstar Church family for your friendship, encouragement, perspective, and support. You have been a family and a home away from home. It has been a blessing to serve alongside such a devoted group of believers.

Of making many books there is no end,
and much study is a weariness of the flesh.
The end of the matter; all has been heard.
Fear God and keep his commandments,
for this is the whole duty of mankind.
– Ecclesiastes 12:12-13

Table of Contents

1 Introduction 1
1-1 SAFFiR | Shipboard Autonomous Fire Fighting Robot 1
1-2 THOR | Tactical Hazardous Operations Robot 2
1-3 SAFFiR & THOR Design Approach 3
1-3-1 Why Humanoids 3
1-3-2 Why Linear Actuators 3
1-3-3 Why Compliance 4
1-4 State-of-the-Art in Humanoids 4
1-5 Problem Statement 7
1-6 Contributions 7
1-7 Outline of Dissertation 8
1-8 Attribution 8
2 Configurable Compliance for Linear Series Elastic Actuators 10
2-1 Abstract 10
2-2 Introduction 10
2-2-1 Series Elastic Actuators 11
2-2-2 Variable Compliance 12
2-3 Configurable Compliance 13
2-3-1 SAFFiR Linear SEA 14
2-3-2 THOR Linear SEA 15
2-3-3 THOR Linear-Hoekens SEA 16
2-3-4 Cantilevered Beam Benefits 16
2-4 SAFFiR Configurable Compliance - End Loading 17
2-4-1 Cantilever Beam Material Selection 19
2-4-2 Cantilevered Beam Design 20
2-4-3 SAFFiR Stiffness Tuning 23
2-5 THOR Configurable Compliance - Moment Loading 23
2-5-1 Cantilevered Beam Design 25
2-6 Discussion 27
3 An Unlumped Model for Linear Series Elastic Actuators with Ball Screw Drives 29
3-1 Abstract 29
3-2 Introduction 29

3-2-1 Related Work on Moving Output and Unlumped Models 30
3-2-2 Model Simplicity vs. Fidelity 31
3-2-3 Ball Screw Driven Linear SEAs 31
3-3 Unlumped Rack & Pinion Model 32
3-3-1 Changing Ground 32
3-3-2 Lumped Mass & Inertia 33
3-3-3 Unlumped Model 34
3-3-4 High Impedance Model Comparison 35
3-3-5 Initial Observations 36
3-4 Results 36
3-4-1 System Identification – High Impedance Test Case 37
3-4-2 Fitting a Model to the $F1/Fm$ Response 38
3-4-3 Comparing the $F2/Fm$ and $F1/F2$ Responses 39
3-4-4 Intuitive Interpretations 40
3-4-5 Moving Output Results 40
3-5 Discussion 42
3-6 Future Work 42
3-7 Acknowledgment 43
4 Design, Modeling, and Stiffness Selection of Linear Series Elastic Actuators 44
4-1 Abstract 44
4-2 Introduction 44
4-2-1 Depicting Screw-Type Actuators 45
4-2-2 Linear Series Elastic Actuators 46
4-2-3 Stiffness Selection 48
4-2-4 Modeling Series Elastic Actuators 48
4-2-5 Joint Torque & Intermediate Inertias 49
4-2-6 Paper Outline 50
4-3 Spring Location in Linear SEAs 50
4-4 Model Derivations 51
4-4-1 Equations of Motion | Sprung Ball Nut 52
4-4-2 Equations of Motion | Sprung Motor Housing 52
4-4-3 Simulink Models 53
4-4-4 High Impedance Test Case 53
4-4-5 Discussion | Actuator Models 55

4-5 Actuator Dynamics 56
4-5-1 High Impedance Test Case 56
4-5-2 Moving Output Test Case 60
4-5-3 Stiffness Criteria | Pure Force Source Approximation 61
4-6 Controlled Performance 61
4-6-1 Ideal Compensator 61
4-6-2 Effect of Current Limit 64
4-6-3 Effect of Input Amplitude 65
4-6-4 Design Implications | Force Bandwidth 66
4-7 Spring Location & Stiffness Selection 67
4-8 Conclusions 68
4-8-1 Model Derivations 69
4-8-2 Actuator Dynamics 69
4-8-3 Controlled Performance 69
4-8-4 Future Work 70
5 Conclusions 71
5-1 SAFFiR & THOR Results 71
5-2 Future Work 73
References 75

List of Figures

Figure 1-1. (left) Potential shipboard fire fighting scenario, used with permission of B. Lattimer, (right) Picture of the SAFFiR Prototype. 2
Figure 1-2. (left) Potential disaster response scenario, image courtesy of DARPA, (right) Picture of the THOR robot, image used with permission of J. Holler. 2
Figure 2-1. Linear Series Elastic Actuator from [15], used with permission of J. Pratt. 11
Figure 2-2. SAFFiR lower body actuator forces as a function of time for the right hip, knee, and ankle during a walking cycle. Positive forces represent compression, negative forces represent tension. SS stands for single support, DS stands for double support. Image used with permission of D. Lahr. 13
Figure 2-3. SAFFiR Linear SEA with Configurable Compliance. 14
Figure 2-4. THOR Linear SEA with Configurable Compliance, used with permission of J. Holler. 15
Figure 2-5. Schematic of THOR Linear-Hoekens SEA with Configurable Compliance, used with permission of C. Knabe. 16
Figure 2-6. Mechanical advantage profile of THOR Linear-Hoekens SEA over a 160 degree range of motion, used with permission of C. Knabe. 16
Figure 2-7. SAFFiR Configurable Compliance design. 18
Figure 2-8. Outer clamp and inner pivots (left), cross section view of the movable pivot clamp (right). 18
Figure 2-9. SAFFiR rigid member used in place of the cantilevered beam for rigid actuators. 19
Figure 2-10. Configurable Compliance in the hip joint of SAFFiR. 21
Figure 2-11. Cantilevered beam loading conditions. 21
Figure 2-12. Experimental load vs. deflection of the SAFFiR Configurable Compliance design. 22
Figure 2-13. Stiffness vs. movable pivot position of the SAFFiR Configurable Compliance design. 23
Figure 2-14. Rendering and schematic of the Configurable Compliance design for the THOR Series Elastic Actuators, used with permission of J. Holler. 24
Figure 2-15. Cross section of the THOR Configurable Compliance design, used with permission of C. Knabe. 24
Figure 2-16. Exploded view of the THOR Configurable Compliance mounting and assembly, used with permission of C. Knabe. 25
Figure 2-17. Schematic for a cantilevered beam under moment loading, used with permission of C. Knabe. 26
Figure 2-18. Simulated effective stiffness of cantilevered beam under moment loading. 26
Figure 2-19. Experimental stiffness experiments for moment loading Configurable Compliance. Displacement represents the total actuator length change, for the given load case. 27
Figure 2-20. Available stiffness settings for the two Configurable Compliance designs. 28
Figure 3-1. Early lumped models for SEAs. Fm is the motor force, mk is the lumped sprung mass, bm is the lumped damping, and k is the stiffness of the physical spring placed in series. 30
Figure 3-2. Schematic of the THOR Linear SEA used in the lower body of THOR. Two sets of parallel actuators power the hip roll/yaw and ankle pitch/roll DOF. A modified version of this design with an inverted Hoekens linkage output powers the hip pitch and knee pitch [34]. Used with permission of J. Holler 31
Figure 3-3. Two-link moving output models for a legged robot with changing ground contacts. The general model (a) shows two moving links. The model can be simplified for the stance phase (b, c, and d) or the swing phase (e, f, and g). 32
Figure 3-4. Lumped and unlumped models for linear SEAs showing the moving output test case. 34
Figure 3-5. Lumped and unlumped models for linear SEAs showing the high impedance test case. 35
Figure 3-6. THOR Linear SEA in the high impedance test case with a load cell at either end. 37
Figure 3-7. Open loop frequency response from Simulink simulations and experimental system identification for F1/Fm (a), F2/Fm (b), and F1/F2 (c). 38
Figure 3-8. THOR Linear SEA in the moving output test case with a 4 kg mass. 41

Figure 3-9. Open loop frequency response from Simulink simulations and experimental system identification for F1/Fm with a moving output. 41
Figure 3-10. Proposed unlumped models for (a) an alternate linear SEA design, and (b) for a geared rotary SEA design. 43
Figure 4-1. Lumped models for SEAs. Fm is the motor force, mk is the lumped sprung mass, bm is the lumped damping, and k is the stiffness of the physical spring placed in series. 45
Figure 4-2. Ball screw driven linear actuator showing a belt reduction to ball screw transmission (a) and the equivalent rack & pinion representation (b). 46
Figure 4-3. Linear Series Elastic Actuator from [15], used with permission of J. Pratt. 47
Figure 4-4. Schematic of the THOR Linear SEA used in the lower body of THOR, used with permission of J. Holler. 47
Figure 4-5. The effect of changing ground contacts in the swing phase (a, b, and c), and stance phase (d, e, and f). 49
Figure 4-6. Schematic and models of the two most common spring locations for linear Series Elastic Actuators, the Sprung Ball Nut (a, b) and Sprung Motor Housing (c, d). 51
Figure 4-7. Free body diagrams for a sprung ball nut linear SEA actuator model. 52
Figure 4-8. Free body diagrams for a sprung motor housing linear SEA actuator model. 52
Figure 4-9. High impedance test case for the two spring configurations. 54
Figure 4-10. Open loop frequency response for the high impedance test case from Simulink simulations for F1/Fm, F2/Fm, and Fsprung/Funsprung. 57
Figure 4-11. Contour plot showing the relationship between spring stiffness, actuator sprung mass, and the maximum bandwidth for which the pure force source assumption is valid. 59
Figure 4-12. Open loop frequency response of the moving output test case from Simulink simulations for F1/Fm, F2/Fm, and Fsprung/Funsprung. 60
Figure 4-13. Control diagram of an ideal inverse plant compensator with motor saturation. 62
Figure 4-14. Controlled performance of an ideal compensator with motor saturation. The input force amplitude is 200 N and the current limit is 10 A. 62
Figure 4-15. Signal clipping due to saturation of a control signal. 64
Figure 4-16. Effect of current limit on the controlled performance of an ideal compensator with motor saturation. The input force amplitude is 200 N and the current limit is varied from 2.5 A to 40 A. 65
Figure 4-17. Effect of input force amplitude on the controlled performance of an ideal compensator with motor saturation. The current limit is 10 A while the input force amplitude is varied from 50 N to 800 N. 66
Figure 5-1. SAFFiR Prototype robot demonstrating stable walking on gravel, grass, and sand. Images used with permission of D. Lahr. 72
Figure 5-2. THOR robot demonstrating stable walking on concrete, gravel, and grass. Images used with permission of M. Hopkins. 73

List of Tables

Table 1: Comparison of State of the Art Humanoids. 6
Table 2: Comparison of SAFFiR and THOR SEA designs. 15
Table 3: Young's Modulus and Yield Strength of Common Compliant Materials. 20
Table 4: Comparison of the two Configurable Compliance approaches. 28
Table 5: Known and Extracted model Variables 39
Table 6: Known and Extracted model Variables 53
Table 7: Actuator Model Transfer Functions 56
Table 8: Comparison of Different Configurations of the THOR-Linear SAE ($Imax$ = 10 A, AFd = 200 N) 67

1 Introduction

Recent advancements in actuator technologies, computation, and control have led to major leaps in capability and have brought humanoids ever closer to being feasible solutions for real-world applications. As the capabilities of humanoids increase, they will be called on to operate in unstructured real world environments. This realization has driven researchers to develop more dynamic, robust, and adaptable robots.

Out of these efforts have emerged new actuator technologies such as Series Elastic Actuators as well as new control approaches such as force control and momentum control walking. One of the common threads in recent work has been an interest in compliant actuators. This interest is largely motivated by a growing understanding of the role of compliance in animal locomotion which includes potential benefits such as impact absorption, low impedance actuation, and energy storage.

The work presented below is based on several years of developing novel actuators for two full-scale compliant humanoid robots, SAFFiR and THOR. Both are state-of-the-art robots intended to operate in the extremely challenging real world scenarios of shipboard firefighting and disaster response. This dissertation presents some of the development aimed at gaining a better understanding of compliant actuation and its role in humanoid robots. It specifically address linear Series Elastic Actuators as they are used in humanoids, but much of the analysis and insights are applicable to other implementations (rotary, cable drive), actuator technologies (hydraulic, pneumatic), and applications (quadrupeds, industrial actuators).

1-1 SAFFiR | Shipboard Autonomous Fire Fighting Robot

The SAFFiR (Shipboard Autonomous Fire Fighting Robot) Prototype is a full-scale lower body biped that was developed as part of the SAFFiR project for the US Navy with the application of shipboard fire fighting [1]. The SAFFiR Prototype has 12 degrees of freedom (DOF): 3 DOF in each hip (roll, pitch, yaw); 1 DOF in each knee (pitch); and 2 DOF in each ankle (roll, pitch). Custom linear Series Elastic Actuators were developed specifically for this application and are discussed in more detail in Section 2-3-1. SAFFiR was designed with the intent to investigate compliant walking and the role of compliance at different joints. One of the design challenges, which this dissertation will address, was selecting stiffness values and how to physically implement compliance for linear actuators. The current status, as shown in Figure 1-1 (right), is a fully functioning and walking lower body capable of walking across level ground, strewn plywood, thick turf grass, gravel, and sand.

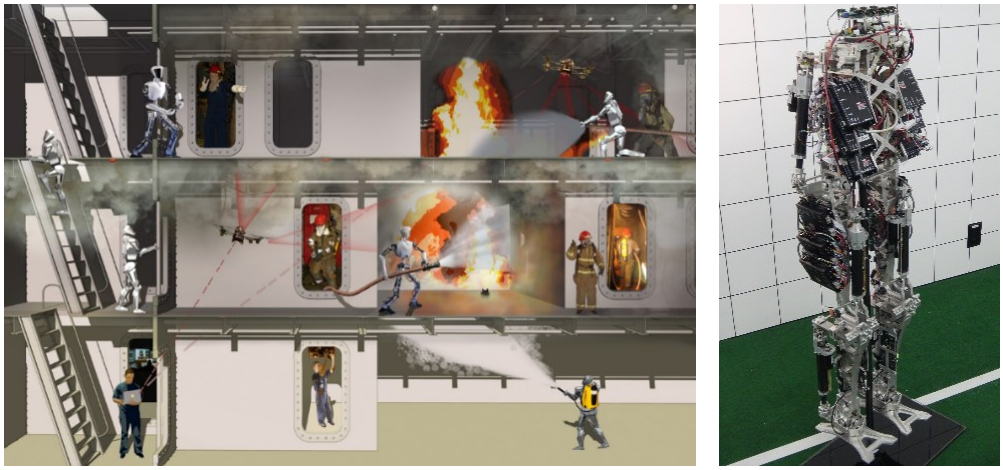

Figure 1-1. (left) Potential shipboard fire fighting scenario, used with permission of B. Lattimer, (right) Picture of the SAFFiR Prototype.

1-2 THOR | Tactical Hazardous Operations Robot

THOR (Tactical Hazardous Operations Robot) is a full-scale humanoid robot which was developed for the DARPA Robotics Challenge (DRC) as a funded Track A entry. The challenge requires participating robots to be capable of tasks such as driving a vehicle, climbing ladders, using hand tools, turning valves, walking through closed doors, clearing debris, attaching a water hose, and walking over rough terrain. THOR is 1.7 m tall, weighs 60kg, and has 34 total DOF; 6 DOF legs, 7 DOF arms, a 2 DOF waist, 2 DOF neck, and 2 DOF hands. The linear Series Elastic Actuators were completely redesigned for the more demanding DRC Tasks including both a linear version and a linear-Hoekens linkage version which are discussed in more detail in Sections 2-3-2 and 2-3-3. Team THOR participated in the DRC Trials event in December of 2013, and was selected as one of the nine funded Track A teams to continue as DRC Finalists and will compete in the DRC Finals in 2015. The current status of THOR, as shown in Figure 1-2 (right), is a fully functioning robot capable of compliant balancing and walking [2].

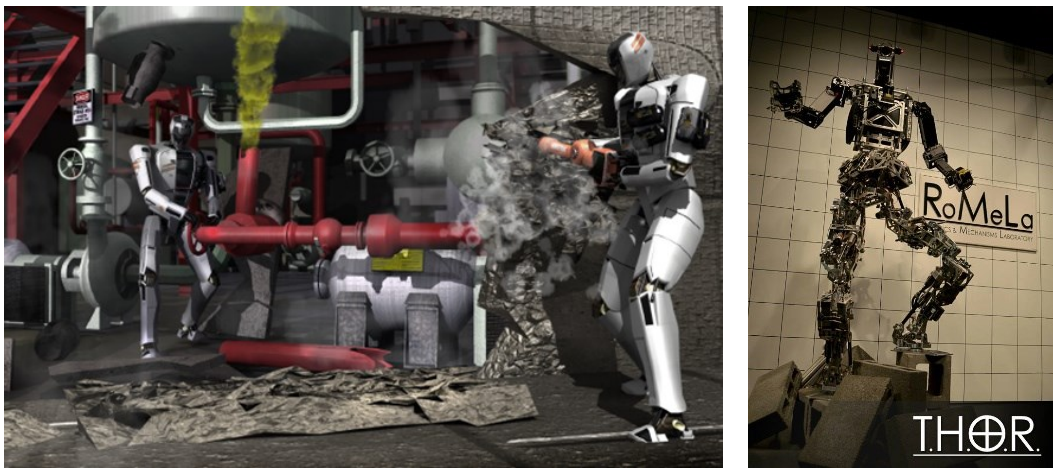

Figure 1-2. (left) Potential disaster response scenario, image courtesy of DARPA, (right) Picture of the THOR robot, image used with permission of J. Holler.

1-3 SAFFiR & THOR Design Approach

The applications of SAFFiR and THOR place considerable demands on their design and performance. Very few robots are capable or even intended to operate in complex, unstructured, or non-static environments. To enable their state-of-the-art performance, the design approach for both SAFFiR and THOR centered on three common themes; humanoid form factor, linear actuation, and compliant walking. The motivation for each of the three themes is given in the following sections.

1-3-1 Why Humanoids

Science Fiction: History has shown that humanity is and undoubtedly always will be fascinated with humanoid robots. Humanoid automata were mentioned as early as the 3rd Century in Chinese texts, Leonardo da Vinci had sketches for a mechanical knight around 1495, and it is no surprise that the term "robot" was first introduced in a 1920 Czech science fiction play referring to artificially created humanoids [3].

Environment: Many of the applications for robots are in environments designed and built for humans to operate. Existing homes, factories, and especially military vessels are intentionally designed around the human form factor. Many of the desired tasks for robots such as mobile manipulation are also particularly well suited for humanoid robots. A human-like form and function means that there would be little need for adaptation of the environment or processes.

Safety: By having a human-like form and function, it is safer for humanoid robots to operate alongside humans because humans would have an intuitive understanding of how the robot is expected to move and operate.

Biomechanics: The human body is a very complex system, the function of which is still not fully understood. Humanoid robots can be treated as simplified models of the very complex human body. As such, they can serve as effective stand-ins for humans in locomotion research. The information flow between biomechanics research and robotics research has historically gone both ways and benefited both fields. Robotics researchers look to the insights gleaned from biomechanics research to inform their designs, which if successful, serve to validate the biomechanics research. Humanoid research can, in turn, provide unexpected biomechanics insights [4].

Rehabilitation & Prosthetics: Humanoid research tackles many of the same design and application challenges as rehabilitation and prosthetics research. Therefore, there are many potential applications of the new technology towards benefiting elderly, stroke, and amputee patients [5].

1-3-2 Why Linear Actuators

The majority of robots, especially electrically powered humanoid robots, use rotary actuators in serial chains. Electric motors have traditionally been available in a rotary form factor, and it is kinematically simpler to design a serial arrangement for the robot's limbs. By comparison, linear actuators can be limited by their packaging constraints and a limited range of motion. However, by carefully designing linear actuators they can provide some significant benefits.

Linear actuators typically act on a lever arm such that the mechanical advantage varies throughout the range of motion. While normally considered an inconvenience, the varying mechanical advantage can be exploited to provide peak force where it is needed within the range of motion. Linear actuators are also conducive to parallel actuation arrangements in which two or more actuators control two or more degrees of freedom simultaneously. By doing so, multiple actuators can be recruited for high power tasks when necessary. Finally, linear actuators can be packaged externally around a light weight yet rigid internal structure, much like the skeletal structure of humans.

1-3-3 Why Compliance

Compared to state-of-the-art robots, biological systems demonstrate remarkably better efficiency, agility, adaptability, and robustness. Many recent studies suggest that a core principle behind these advantages is compliance. For example, Alexander proposed that legged animals make use of compliance to absorb foot impacts, to bounce like pogo stick-like springs during walking and running, and to serve as return springs for reversing the direction of swinging limbs [6].

Numerous other studies have investigated compliance in a variety of species and in behaviors even beyond locomotion [7]. In [8], the authors present a compilation of studies that show how compliance serves diverse roles in metabolic efficiency, muscle power amplification and attenuation, and mechanical feedback for stability. Robots could benefit from many of these advantages, especially in applications where moving naturally, absorbing impact loads, storing energy, and working safely around humans are priorities. Robots with the ability to adjust stiffness could be especially advantageous by being able to adapt to different loading conditions and environments [9].

1-4 State-of-the-Art in Humanoids

Given the vast variety of humanoid robots, it would be beneficial to identify where SAFFiR and THOR stand in the field. The intention below is to provide an overview of the humanoids research field, not an exhaustive review. Compliant actuators are widely used in bipeds, prosthetics, and exoskeleton applications. The focus in this section is on current state-of-the-art humanoids that are full-scale, compliant, linearly actuated, and capable of 3D walking.

Advanced: The three most advanced humanoid robots, based on overall capability and performance, are ATLAS from Boston Dynamics (recently acquired by Google) [10], [11], the SCHAFT robot which is based on the successful line of HRP robots (also recently acquired by Google) [12], and the Honda ASIMO robot [13]. Each of these have demonstrated impressive dynamic walking capability over uneven terrain and, in the case of HRP and ASIMO, even running.

Compliant: There are a growing number of compliant humanoid robots, but few are full-scale and capable of 3D walking. The three most successful examples are the cable driven series elastic robots Flame and Tulip from TU Delft [14], the linear series elastic robot M2V2 from Yobotics/IHMC [15], and the rotary series elastic robot COMAN from IIT [16]. Of these, Flame and Tulip have compliant actuators only in the ankle pitch, knee pitch, and hip pitch DOF, M2V2

uses identical SEAs at all 12 DOF, and COMAN only uses SEAs at the ankle pitch and knee pitch. The developers of the DLR Biped use a compliant walking approach which, although it does not have a physical spring, takes advantage of the inherent compliance in the harmonic drives [17]. The DLR group is also actively developing new actuator designs which have variable compliance and variable damping. The most recently developed compliant humanoid is Valkyrie, NASA JSC's entry into the DARPA Robotics Challenge which uses Series Elastic Actuators [18], [19].

Linear & Parallel: Already mentioned above, M2V2 uses linear SEAs for all 12 DOF, with the ankles being parallelly actuated. HRP-4, one of the robots in the HRP/SCHAFT line, uses a linear actuator only for the ankle pitch. Valkyrie uses 2 linear SEAs for a parallelly actuated ankle. Not yet mentioned are LOLA, which uses linear actuators for a 1 DOF knee and a 2 DOF parallelly actuated ankle [20], and JOHNNIE which uses linear actuators for a 2 DOF ankle [21]. All of these robots have demonstrated successful 3D walking.

Pneumatic: There are also several pneumatically powered robots which are inherently linearly actuated and inherently compliant. Lucy serves as a good example and has been the most successfully walking pneumatic robot to date [22].

2D Walkers: While this discussion has specifically focused on 3D walkers, it is important to note that there are several groups doing related research with full-scale 2D walkers. Of these, the best examples are Mabel and Atrias [23]. Both have demonstrated very impressive and robust compliant walking and running [24].

Within this field, SAFFiR and THOR are two of only seven successfully walking electrically powered compliant humanoids. They are the only robots with parallelly actuated hips, and with the exception of M2V2 and Valkyrie, THOR is the only other humanoid with a fully compliant lower body.

Table 1: Comparison of State of the Art Humanoids.

Robot	Walking Capability	Compliance (Lower Body)	Linear Actuators	Parallel Actuation
Petman/ATLAS	Dynamic, Rough Terrain	Hydraulic	Hydraulic	4 DOF
HRP/SCHAFT	Dynamic, Running	---	2 DOF	---
ASIMO	Dynamic, Running	---	---	---
M2V2	Compliant, Level Ground	12 DOF	12 DOF	4 DOF
Flame/Tulip	Compliant, Level Ground	Cable Drives	---	---
COMAN	Compliant, Level Ground	4 DOF	---	---
DLR Biped	Compliant, Level Ground	12 DOF*	---	---
Valkyrie	Compliant, Level Ground	12 DOF	4 DOF	4 DOF
LOLA	Level Ground	---	6 DOF	4 DOF
JOHNNIE	Level Ground	---	4 DOF	4 DOF
Lucy	Pneumatic, 2D Biped	6 DOF	6 DOF	Antagonistic
Mabel/Atrias	2D Biped w/ Boom	Compliant Leg	---	---
SAFFiR	Level Ground, Gravel, Grass, Sand	10 DOF	10 DOF	8 DOF
THOR	Level Ground, Gravel, Grass,	12 DOF	12 DOF	8 DOF

1-5 Problem Statement

The design, modeling, and control of actuators in robotics application is critical because the rest of the robot is often designed around the actuators. As the field of compliant humanoids advances, there will be a continuing need for higher performance actuators and for more descriptive dynamic models. The work below addresses the design and modeling of linear Series Elastic Actuators as they are used in humanoid robots.

Compliance Design: While there are many existing designs, there is still a need for better compliant actuator designs which can be used in high performance humanoid applications. The wide variety of compliant actuators can be summarized into two main categories; Variable Compliance Actuators (VCA) which can actively adjust the physical stiffness of the spring element, and Series Elastic Actuators (SEA) which use a spring with a set stiffness. VCAs require an extra actuator as well as added complexity, size, and weight. SEAs, on the other hand, retain their simplicity but have a fixed stiffness which requires the actuator to be disassembled in order to physically exchange the spring for one of a different stiffness. Each of the existing designs has significant tradeoffs, so there is a clear need and opportunity for more and better designs.

Actuator Model: Compliant actuators are frequently modeled using an oversimplified lumped mass model which has remained mostly unchanged in almost two decades. As a result, most of the literature oversimplifies the actuator dynamics, only considers the forces acting on one side of the actuator, and overemphasizes the importance of spring location. The conventional lumped model served well for early development but more descriptive models are now needed to more faithfully represent the true actuator dynamics and to better compare different actuator designs.

1-6 Contributions

The major contributions of this work are:

1. The presentation of a novel approach to implementing compliance in linear Series Elastic Actuators. The approach is called Configurable Compliance and uses a structural cantilevered beam with a movable/removable pivot to enable a range of stiffnesses without the complexity of existing designs.

2. The Configurable Compliance design for the SAFFiR-Linear actuator which uses a perpendicular beam orientation and offers a wide range of stiffnesses via a movable pivot.

3. The Configurable Compliance design for the THOR-Linear and THOR-Hoekens actuators which uses a parallel beam orientation with three discrete stiffness settings and high load capacity via a pure moment loading.

4. The presentation of a new unlumped actuator model for linear Series Elastic Actuators which is more descriptive and better represents the actuator dynamics. The unlumped model uses a rack & pinion representation which is shown to be significantly more accurate than the conventional lumped model.

5. The development of design implications for linear Series Elastic Actuators using the new unlumped rack & pinion model. The analysis reveals how spring location, spring stiffness, sprung mass, motor saturation, and input amplitude impact the actuator dynamics and controlled force bandwidth.

6. Contributed to the design, fabrication, and integration of two successful, full-scale, state-of-the-art, compliant walking humanoids: SAFFiR and THOR.

1-7 Outline of Dissertation

The remainder of this dissertation is organized such that Chapters 2-4 are each intended to serve as a standalone manuscript with Chapter 5 serving as a conclusion.

Chapter 2 presents the motivation and design approach for the compliant actuators developed for SAFFiR and THOR. Both platforms use Series Elastic Actuators with a novel spring implementation called Configurable Compliance, which allows for the stiffness of the physical spring to be manually adjusted without replacing any components. Two different implementations are presented, one of which has a broad range of available stiffness settings while the other only has three stiffness settings but is capable of much higher loads. Both approaches offer numerous advantages over other designs and are novel in the way the physical springs are implemented as well as how the stiffness settings are adjusted.

Chapter 3 introduces a new model for linear Series Elastic Actuators in which the rotational mass of the motor rotor and transmission is decoupled from the linear translational mass of the motor housing. The model uses an intuitive rack & pinion representation of the ball screw transmission and allows for both sides of the actuator to have moving outputs. While this complicates the conventional model for SEAs, the unlumped model is shown to be more accurate in representing the true actuator dynamics.

Chapter 4 uses the new unlumped model to compare the two most common design approaches for linear SEAs which differ primarily in the location of the spring element. The two design approaches are analyzed both analytically and in simulation to investigate the design implications of spring location, spring stiffness, sprung mass, motor saturation, and input amplitude.

Chapter 5 provides a summary of the conclusions, walking results for the SAFFiR and THOR platforms, and ideas for future work.

1-8 Attribution

Chapter 2 of this dissertation is a journal manuscript in preparation for submission.

Chapter 3 of this dissertation has been submitted and is under review as a paper in the IEEE International Conference on Robotics and Automation to be held in May 25-30, 2015 in Seattle, WA. The paper is titled "*An Unlumped Model for Linear Series Elastic Actuators with Ball Screw Drives*" and has three additional co-authors:

Coleman S. Knabe contributed to the early development of the conceptual unlumped model, specifically the use of a rack & pinion representation for screw-based transmissions.

Michael A. Hopkins contributed by developing the system identification code used to generate the frequency response plots and provided valuable insights and feedback for the early drafts of the paper.

Dennis W. Hong contributed by providing valuable insights and feedback for the early drafts of the paper.

Chapter 4 of this dissertation is a journal manuscript in preparation for submission.

2 Configurable Compliance for Linear Series Elastic Actuators

2-1 Abstract

Variable compliance has been a growing topic of interest in legged robotics due to recent studies showing that animals adjust their leg and joint stiffness to adjust their natural dynamics and to accommodate changes in their environment. However, existing designs add significant weight, size, and complexity. Series Elastic Actuators (SEAs), on the other hand, are designed with a set stiffness usually tuned for actuator performance. We propose a new concept for implementing a physical spring in series with a linear SEA using a cantilevered spring. A movable pivot is used to adjust the stiffness by changing the effective length of the cantilever. While the proposed design does not allow for variable compliance, it does retain many of the benefits of passive spring elements such as absorbing impacts, storing energy, and enabling force control. The primary advantage of the new approach is the ability to adjust the stiffness of each joint individually without the increased weight and complexity of variable stiffness designs. This paper introduces the motivation for Configurable Compliance, describes two different implementations of the concept, explains the design methods, and presents experimental data from completed prototypes.

2-2 Introduction

The role of compliance in nature continues to be a growing research field as researchers learn more about how animals and humans utilize the natural compliance of muscles and tendons. Compared to state-of-the-art robots, biological systems demonstrate remarkably better efficiency, agility, adaptability, and robustness. Many recent studies suggest that a core principle behind these advantages is compliance.

For example, Alexander proposed that legged animals make use of compliance to absorb foot impacts, to bounce like pogo stick-like springs during walking and running, and to serve as return springs for reversing the direction of swinging limbs [6]. Animals have been shown to tune their tendon stiffnesses over a variety of time scales; in response to long term exercise, during periods of heavy exercise, and even during a single muscle contraction. In [8], the authors present a compilation of studies that show how compliance serves diverse roles in metabolic efficiency, muscle power amplification and attenuation, and mechanical feedback for stability.

Unlike animals, legged robots have fixed and uniform stiffnesses at each joint and underperform biological systems in power density, efficiency, adaptability, and robustness. Robots could benefit from many of these advantages, especially in applications where moving naturally, absorbing

impact loads, storing energy, and working safely around humans are priorities. Robots with the ability to adjust stiffness could be especially advantageous by being able to adapt to different loading conditions and environments [9]. One approach to reducing this deficiency is improving the design and performance of robotic actuators.

Within the combined fields of humanoid robotics, prosthetics, wearable robotics, and rehabilitation robotics, researchers are actively working on developing more and better compliant actuator designs. Each of these fields stands to benefit from the shock tolerance, energy storage, safety, and efficiency improvements promised by compliant actuation. Each of these fields, however, places a unique set of criteria on the actuator performance and packaging. Given the variety of applications, it is no surprise that the design approaches vary greatly in their power sources (electric, hydraulic, pneumatic), spring arrangements (series, parallel, antagonistic), and form factors (rotary, linear). Despite the variety, no one approach to compliant actuation has yet emerged as a leading contender. Each of the current designs has significant tradeoffs, so there is a persistent need and opportunity for more and better designs.

Both SAFFiR and THOR are electrically powered and use linear actuators. Within this subset of applications (electric, linear), the most successful design approach to date has been Series Elastic Actuators (SEAs) [25]. There has also been a growing interest within the robotics community in Variable Compliance Actuators. Both of these approaches show promise for high performance applications but have disadvantages. Configurable Compliance, which we introduce in this paper, serves as a compromise between the two by enabling the adjustment of physical spring stiffness without a significant increase in complexity, size, or weight [26].

2-2-1 Series Elastic Actuators

An example of an actuator with fixed stiffness springs is the Series Elastic Actuator, shown in Figure 2-1, in which a set of die springs are placed between the ball screw transmission and the actuator output. The springs are preloaded to half of their capacity which essentially doubles their effective stiffness. A high resolution position sensor measures the spring displacement as a means to measure force, essentially using the spring as a sensor for force control [25]. The fixed stiffness of SEAs is typically designed for actuator performance, tuned for low impedance and high bandwidth force control. Because the spring is treated as a component of the actuator, the same spring stiffness is used at every DOF.

Figure 2-1. Linear Series Elastic Actuator from [15], used with permission of J. Pratt.

The Configurable Compliance design which is the subject of this paper is an alternative implementation of Series Elastic Actuators. Much of the same approach is used for the actuator design with the exception of the physical spring element. Instead of a set of die springs, a cantilevered titanium beam is used.

2-2-2 Variable Compliance

Several biomechanics studies have found that animals actively adjust the stiffness of their joints and limbs to adjust the natural dynamics of the system in response to disturbances or changes in their environment. For example, Ferris et al. have shown that humans adjust the stiffness of their legs to accommodate changes in surface stiffness during hopping and running [27], [28]. Because of this, a number of research groups have been working to develop novel approaches for implementing variable compliance. Unfortunately, existing designs add significant weight, size, and complexity.

There are many different approaches to variable compliance designs [9], [29]. As is often the case in robotics, the particular implementation of a design depends in large part on the specific needs of the application. Van Ham et al. provide an overview of the main design approaches and provides multiple examples of each [30]. Vanderborght et al. provide a similar review of variable impedance actuators, which are capable of varying both stiffness and damping [31].

It should be noted that different research groups seem to use different naming conventions to describe what is essentially the same concept. The confusion arises from the complementary nature of the terms compliance and stiffness as well as the interchangeability of the terms adjustable, variable, and controllable. We use the term variable compliance, by which we mean a mechanism or method which uses a passive physical spring element, the stiffness of which can be adjusted during operation by some mechanical means.

The most common drawback of variable compliance designs is the use of two motors. In some cases the motors are arranged antagonistically, using both motors to adjust the preload and the equilibrium position simultaneously. In other cases one motor is used to adjust the stiffness while the other motor controls the equilibrium position of the joint. Either way, the addition of a second motor has a significant impact on weight. There are also typically gear trains, mechanisms, or additional moving parts used for adjusting the stiffness of the spring element. These additional components not only add to the weight but also increase the actuator's complexity, part count, and volume. Unfortunately, these drawbacks are especially a challenge for humanoids which are high degree of freedom (DOF) systems. Having many DOF means that any additional weight, complexity, or size gets multiplied by the total number of actuators; leading to a large increase in the overall weight.

It is important to note that this discussion is not commenting on the validity or usefulness of any existing designs. Each design has its advantages and is undoubtedly well suited for the designer's intended application. However, in the case of full-scale humanoids there are stringent constraints on the weight, complexity, and size of components. Unfortunately, many of the existing variable compliance designs feature one or more of these drawbacks.

2-3 Configurable Compliance

Variable Compliance Actuators require an extra actuator as well as added complexity, size, and weight. SEAs, on the other hand, retain their simplicity but have a fixed stiffness which requires the actuator to be disassembled in order to physically exchange the spring for one of a different stiffness. In contrast to this "one-stiffness-fits-all" approach, we believe there is a strong case for using different stiffness settings throughout a robot. Distal joints, for example, are more likely to experience impact loads (lower stiffness desired) while proximal joints experience larger inertial loads (higher stiffness desired). Furthermore, each joint has different power requirements, mechanical advantage, velocity profiles, and distinct contributions to the overall locomotion behavior.

Figure 2-2 shows the experimentally measured forces of each actuator in the SAFFiR right leg during a typical walking gait. The yaw DOF is not represented due to the lack of a load cell on the yaw actuator which is a rigid actuator operated in position mode. As shown by the plots, the ankle actuators see half the load of the hip actuators but operate in both tension and compression. The knee actuator, on the other hand, operates almost exclusively in compression. Each joint has a unique force-velocity profile.

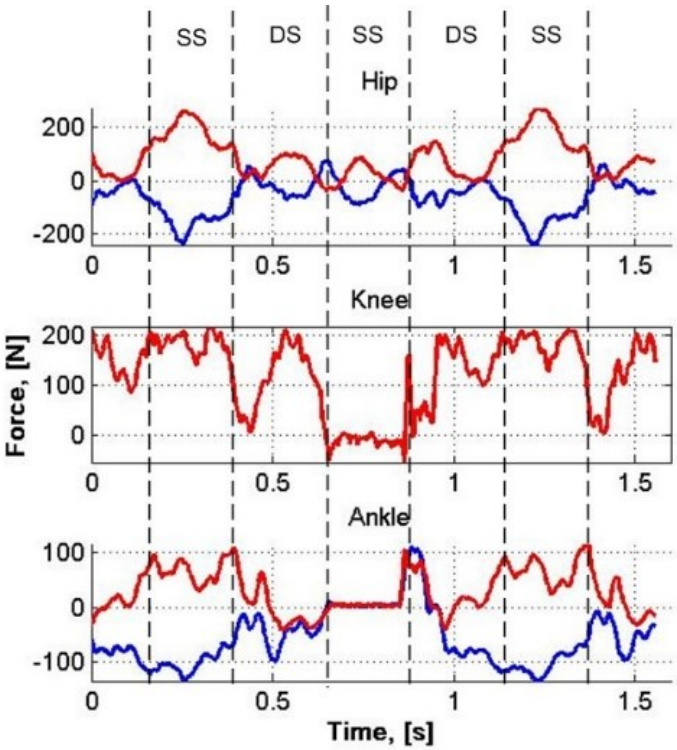

Figure 2-2. SAFFiR lower body actuator forces as a function of time for the right hip, knee, and ankle during a walking cycle. Positive forces represent compression, negative forces represent tension. SS stands for single support, DS stands for double support. Image used with permission of D. Lahr.

As a way to investigate nonuniform distributed stiffness, Configurable Compliance is an approach in which a passive physical spring has a fixed stiffness during operation but can be adjusted manually to achieve different stiffness values at any of the DOF.

In contrast to variable compliance, Configurable Compliance is not controllable and loses the benefit of modifying the natural dynamics on-the-fly. Nevertheless, it does retain many of the benefits of passive spring elements such as absorbing impacts, storing energy, and enabling force control. The primary advantage of the design is the ability to adjust the stiffness of each joint individually without the increased weight and complexity of existing variable compliance designs. In effect, an SEA with Configurable Compliance performs as a conventional SEA during operation, but gives the user the ability to readily change the spring stiffness in between runs. To use the terminology in [30], Configurable Compliance acts as an equilibrium controlled stiffness under operation but has a structure controlled stiffness which can be adjusted manually.

Two different implementations of Configurable Compliance have been developed and implemented into three different actuator designs. Both Configurable Compliance designs use a titanium cantilevered beam as the elastic element but the beam is placed under very different loading conditions in each arrangement. The SAFFiR actuator uses an end loading arrangement in which the beam is mostly loaded in bending by a concentrated end load from the actuator. Two different SEA designs are used in THOR; both of which use a cantilevered beam in almost pure moment loading.

2-3-1 SAFFiR Linear SEA

The overall design of the SAFFiR Series Elastic Actuator is shown in Figure 2-3. A custom lightweight linear actuator is coupled with a cantilevered titanium beam which serves as the series spring. The actuator is powered by a 100 Watt Maxon EC 4-pole brushless DC motor running at 48 volts which drives a belt reduction and then a ball screw reduction. The actuator is attached to the robot on either end using u-joints. The actuator generates 1000 N, travels at 0.35 m/s, has a stroke of 110 mm, and weighs 0.816 kg including the compliant titanium beam.

The cantilevered beam is positioned perpendicular to the actuator such that the beam deflection is in line with the primary axis of the linear actuator. The u-joints at either end turn the actuator into a two force member and are used to restrict the relative rotation of the ball screw and the ball nut. The upper u-joint has a split-trunnion design, which serves to house a low-profile load cell for force feedback. A more detailed description of the SAFFiR linear SEA can be found in [32].

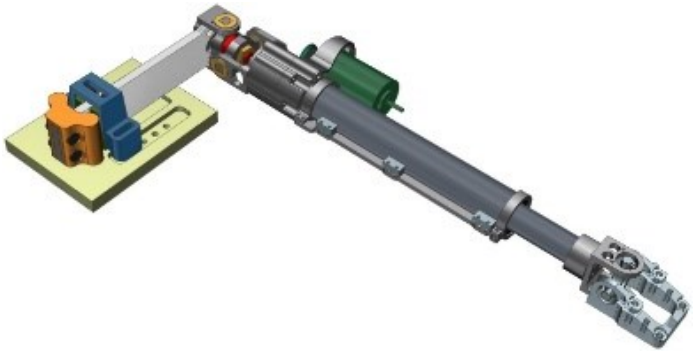

Figure 2-3. SAFFiR Linear SEA with Configurable Compliance.

2-3-2 THOR Linear SEA

The THOR Linear SEA is a redesign and improvement of the SAFFiR version. The key changes are a larger belt reduction ratio (3:1 vs. 2.5:1), a smaller ball screw lead (2 mm vs. 3.175 mm), and a doubling of the peak force from 1000 N to 2225 N. Table 2 provides a comparison of the two actuators and a more detailed description of the THOR SEA design can be found in [33].

Table 2: Comparison of SAFFiR and THOR SEA designs.

	SAFFiR	THOR
Weight (Actuator Only) [kg]	0.653	0.726
Weight (Full SEA) [kg]	0.816	0.938
Maximum Speed [m/s]	0.35	0.19
Continuous Force [N]	300	685
Maximum Force [N]	1000	2225
Spring Constant [kN/m]	145– 512	372 or 655

A major design change in the THOR-Linear actuator is the rearrangement of the compliant beam. As shown in Figure 2-4 the beam is positioned parallel to the primary axis of the actuator and acted on through a lever arm. One of the advantages for this configuration is a better packaging profile compared to the SAFFiR design which had the compliant beam perpendicular to the primary axis. Another advantage is that the beam is loaded with an almost pure moment loading instead of the end loading condition on SAFFiR. The moment loading arrangement allows the beam to store energy uniformly across the entire length of the beam, resulting in a greater overall energy storage capacity. This proved to be critical since the peak force of the THOR actuators doubled that of SAFFiR.

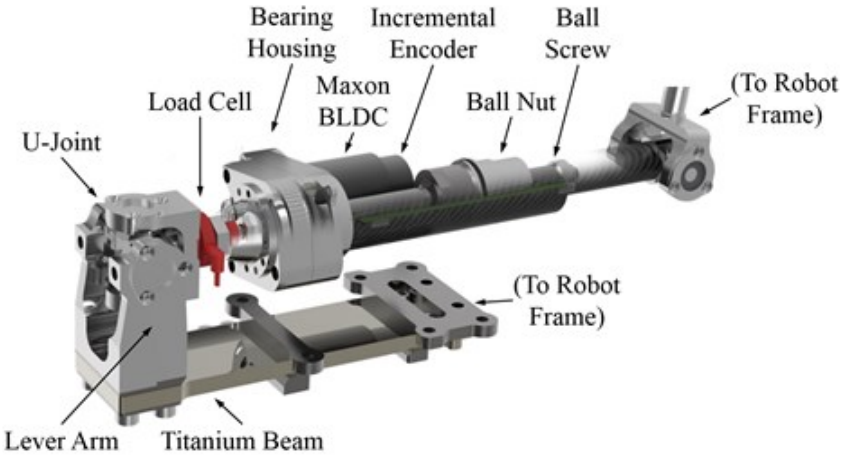

Figure 2-4. THOR Linear SEA with Configurable Compliance, used with permission of J. Holler.

2-3-3 THOR Linear-Hoekens SEA

The hip pitch and knee pitch actuators in THOR use a modified version of the THOR Linear SEA in which the output drives an inverted Hoekens Linkage, shown in Figure 2-5. The other eight actuators in the THOR lower body use a conventional lever arm but the large range of motion requirements of the hip pitch and knee pitch DOF required a different approach. A Hoekens four-bar linkage is conventionally used to transfer a rotary input into a straight-line linear output. By using a novel inversion of the Hoekens linkage, these actuators perform the opposite function; converting the linear output of the SEA to drive the rotation of the joint.

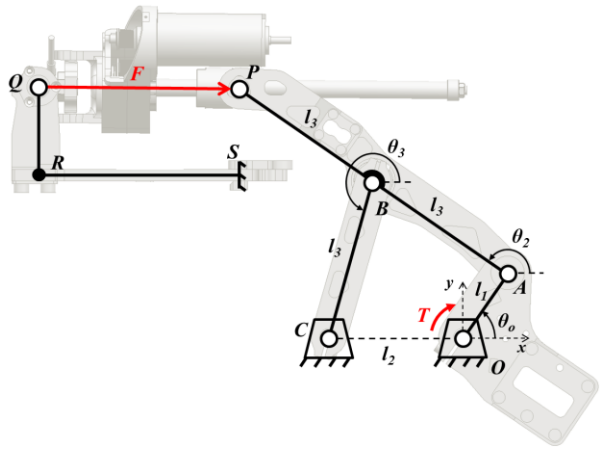

Figure 2-5. Schematic of THOR Linear-Hoekens SEA with Configurable Compliance, used with permission of C. Knabe.

One of the benefits of the Hoekens linkage is that it can be designed to achieve an almost constant mechanical advantage over a large range. Figure 2-6 shows the mechanical advantage only varies by 5% over a 160 degree joint range of motion. A more detailed description of the THOR Linear-Hoekens SEA can be found in [34].

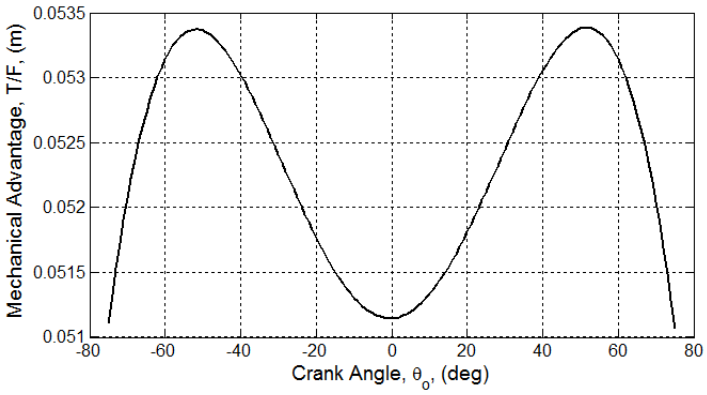

Figure 2-6. Mechanical advantage profile of THOR Linear-Hoekens SEA over a 160 degree range of motion, used with permission of C. Knabe.

2-3-4 Cantilevered Beam Benefits

Cantilevered beams are often used in variable stiffness applications due their simple geometry, straightforward design, and easily adjustable stiffness. The stiffness is usually adjusted by varying

the length of the cantilever, but can also be achieved through other means. In [35] and [36], a laminated beam approach is proposed to implement tunable bending stiffness by adjusting the compressive forces acting on the layers either through electrostatic or pneumatic forces. Of the existing designs that use cantilevers, most are rotary actuators[37]–[40]. The AVSER design in [41], similar to the designs in this paper, uses a cantilevered beam and ball screw transmission but the ball screw drives a cable output resulting in a cable-driven actuator. Other notable examples include a 2 DOF finger that uses leaf springs in [42], and a 2 DOF robotic arm that uses two rotary SEAs to drive a five-bar linkage [43].

To our knowledge, the two designs presented in this paper are the first time a cantilevered beam has been used in a linearly actuated humanoid. When compared to the conventional approach of using die springs, a cantilevered beam offers several advantages:

- The cantilevered beam serves the dual purposes of providing compliance and serving as an attachment point for the actuator, allowing the beam to serve as a structural component and leading to a reduction in overall weight.
- Using a beam moves the spring element away from the main body of the actuator resulting in a shorter overall length compared to methods which place the spring elements in line with the actuator.
- Identical components are used at every joint while still accommodating a wide range of stiffness settings, resulting in a very modular design.
- The stiffness setting can be adjusted independent of the actuator, making the actuators easily interchangeable and allowing the stiffness of any joint to be adjusted on a fully assembled robot.
- Cantilevers are bidirectional, unlike die springs, which can only be loaded in compression.
- The cantilever beam geometry can be readily designed and fabricated on conventional shop equipment. Selecting the strength and stiffness range of the beam is only a matter of selecting the beam material and geometry (length, width, thickness).
- The stiffness setting can be easily adjusted by changing the effective beam length via a movable or removable pivot. The full length of the beam provides a wide range of possible stiffnesses at relatively high resolution (length is an easy variable to measure).

2-4 SAFFiR Configurable Compliance - End Loading

The SAFFiR-Linear Configurable Compliance design, shown in Figure 2-7, is composed of a titanium cantilever beam that is fixed to a base plate at one end, has a movable pivot, and is attached to the upper u-joint of an actuator at the other end. A load cell is placed within the split trunion of the upper u-joint. We used a load cell to directly provide force feedback, instead of the usual approach of measuring spring deflection, so that the force measurements are not dependent on knowing the exact stiffness setting.

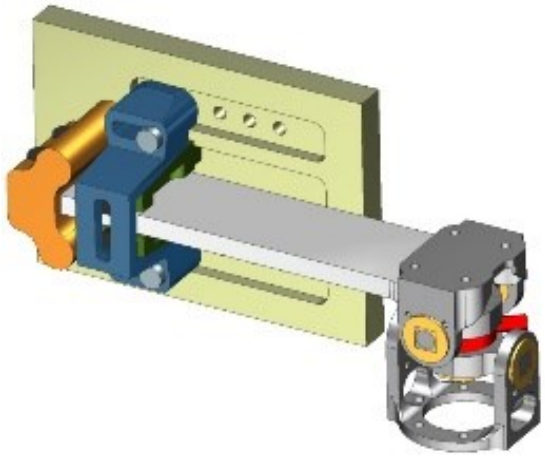

Figure 2-7. SAFFiR Configurable Compliance design.

The fixed end of the cantilever interfaces with the robot using a fixed clamp, shown in orange, that fits into a matched recess in a base pattern. The base pattern can be machined into any rigid support structure or link on the robot and includes other features that interface with the movable pivot, shown in Figure 2-8. It is important for the pivot to be rigidly fixed to both the robot and the cantilever. We accomplish this by using a tapered outer clamp, shown in blue, and two tapered inner pivots, shown in green.

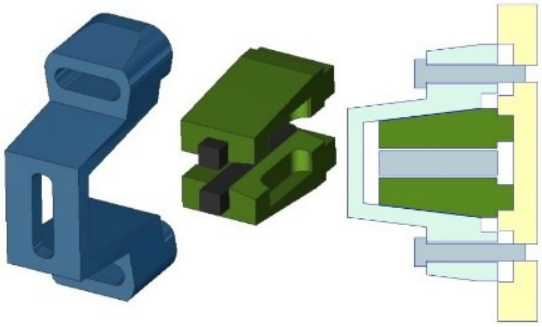

Figure 2-8. Outer clamp and inner pivots (left), cross section view of the movable pivot clamp (right).

The taper provides a way to tighten the pivots firmly against the cantilever while also clamping the pivots into the base pattern. The cross section view illustrates the taper and how the clamp interfaces with the base pattern. The clamp is designed with a 7 degree taper and an axial travel of 4mm which translates into a 1mm travel in the direction of the clamping force. The base pattern's two outer slots are used to keep the outer clamp from flexing outward as it is tightened. The inner slot of the base pattern serves a similar purpose, keeping the inner pivots centered and providing a rigid attachment to the base.

Stiffness adjustment is made possible by the threaded holes in the base pattern and the horizontal slots in the outer clamp. The threaded holes are placed 10 mm apart while the slot allows the outer clamp to slide 11 mm horizontally. The slot travel overlaps the bolt holes, allowing the movable pivot to be positioned at any length along the cantilever by a combination of sliding the clamp within the slot or changing the threaded holes being used to bolt the pivot together.

The inner pivots are composed of aluminum which would not serve well as a bearing surface. Therefore, two flat steel inserts are used to take the load of the cantilever beam. Note that the forward facing side of the inner pivots has an angled cutout in the aluminum to ensure that the cantilever beam only bends about the steel insert.

If it becomes necessary to completely lock out the compliance, the titanium cantilever can be replaced with an aluminum rigid member shown in Figure 2-9. The rigid member uses the existing mounting holes and slots in the base pattern.

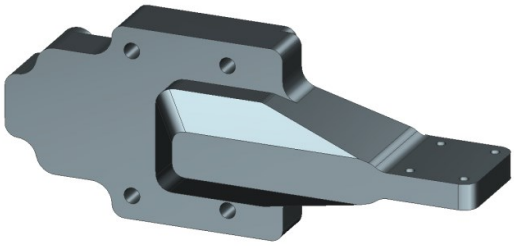

Figure 2-9. SAFFiR rigid member used in place of the cantilevered beam for rigid actuators.

2-4-1 Cantilever Beam Material Selection

The most critical variables in the Configurable Compliance design are the material and the geometry (length, width, thickness) of the cantilevered beam. Considerations for material selection include availability, machinability, creep, thermal properties, electrical conductivity, and many others. However, for compliant mechanisms, an especially important characteristic is the ratio of yield strength to Young's modulus. This ratio provides a measure of how much the material will bend before yielding.

For a beam with a uniform rectangular cross section undergoing bending due to an end load, the max deflection before failure is given by [44]:

$$\delta_{max} = \frac{2}{3}\frac{S_y}{E}\frac{L^2}{h}, \tag{2-1}$$

where S_y is the yield strength, E is the Young's modulus, L is the cantilever length, and h is the cantilever thickness. Note that the max deflection is independent of the beam width. The width is still important, however, as it determines the max load, and thus the stiffness.

$$\sigma_{max} = \frac{6FL}{bh^2}, \tag{2-2}$$

where F is the load and b is the cantilever width.

Equation 2-1 shows that the maximum deflection is dependent on the geometry and on the ratio of yield strength to Young's modulus. For a given geometry, the material with a higher ratio will result in a larger deflection before failure. Table 3, adapted from [44], lists typical values for materials commonly used in compliant mechanisms. Note that the values, especially the yield strength, for a particular material depend heavily on the specific alloy, heat treatment, or epoxy used.

Table 3: Young's Modulus and Yield Strength of Common Compliant Materials.

Material	E (GPa)	S_y (MPa)	(S_y/E) x1000
Steel (1010 hot rolled)	207	179	0.87
Steel (4140 Q&T@400)	207	1641	7.9
Aluminum (7075 Heat Treated)	71.7	503	7.0
Titanium (Ti-13 Heat Treated)	114	1170	10
Polyethylene (HDPE)	1.4	28	20
E-glass (73 vol %)	56	1640	29

While the S_y/E ratio is important, the material's other properties must also be considered. E-glass fiberglass has a very good ratio and yield strength, but it is susceptible to changes in temperature, which would lead to undesirable changes in stiffness. The actual yield strength and modulus values should also be considered. HDPE has a large S_y/E ratio but both of the individual values are small. For a given length and thickness, an HDPE spring would have to be very wide to achieve the same stiffness as a material with the same ratio but higher E and S_y values. Aluminum and 1010 Steel can similarly be eliminated from consideration due to their low yield strengths.

Selecting between steel and titanium is a more difficult decision. Titanium has a higher S_y/E ratio but the high grade alloys are more difficult to find and more difficult to machine. A lower S_y would mean that the spring would have to be larger, but titanium is lighter than steel so the spring ends up having almost the same weight. The titanium that we sourced is Ti-6Al-4V ASTM Grade 5, with a Young's modulus of 115 GPa and a yield strength of 828 MPa, leading to a S_y/E ratio of 0.0072. Titanium, despite the slightly lower ratio, was selected due to its much better corrosion resistance.

2-4-2 Cantilevered Beam Design

With a material in mind, the cantilever geometry can be readily designed using Equations 2-1 and 2-2 to achieve both a desired stiffness and a desired maximum load. For these calculations, the Configurable Compliance is modeled as a fixed end cantilever, the length of which varies, defined as the distance from the movable pivot to the centerline of the linear actuator. Space constraints limited the adjustment range of the movable pivot, resulting in a possible cantilever length range of 65 mm at the shortest setting to 110 mm at the longest setting. The final dimensions of our design are a total length of 125 mm, a beam thickness of 6 mm, and beam width of 30 mm. The resulting beam has a deflection safety factor of 1.46, a load safety factor of 1.35, and a theoretical stiffness range of 145-512 kN/m. Figure 2-10 shows the Configurable Compliance design assembled on SAFFiR.

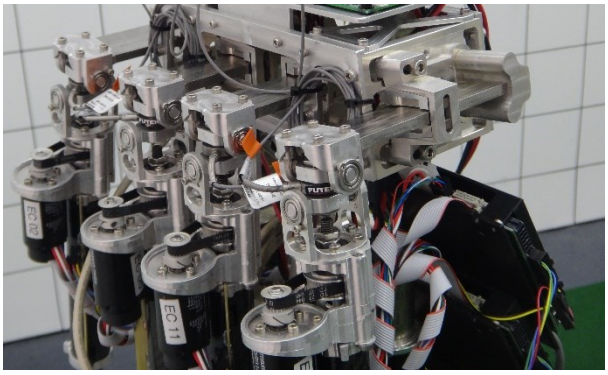

Figure 2-10. Configurable Compliance in the hip joint of SAFFiR.

In designing the beam we modeled it as a cantilever with a fixed end and a concentrated load. However, the fixed end condition may not be the most accurate way to model the movable pivot. Two alternative models were considered: 1) pinned-pinned beam with an overhanging load and 2) fixed-pinned beam with an overhanging load, both shown in Figure 2-11. Which model is most appropriate would depend on how well the fixed clamp behaves as a fixed constraint and how well the movable pivot behaves as a pin constraint.

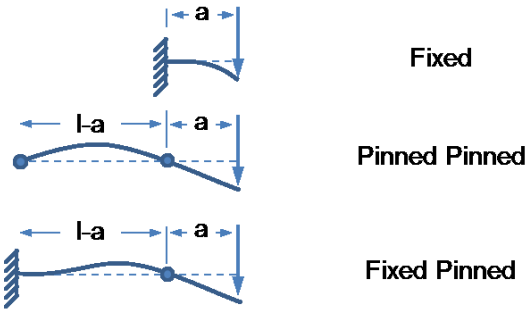

Figure 2-11. Cantilevered beam loading conditions.

The deflection equations for each loading condition are given as:

$$\delta_{\text{fixed}} = -\frac{\text{Fa}^3}{3\text{EI}}, \tag{2-3}$$

$$\delta_{\text{pinpin}} = -\frac{\text{Fla}^2}{3\text{EI}}, \tag{2-4}$$

$$\delta_{\text{fixedpin}} = -\frac{\text{F}}{12\text{EI}}(\text{a}^3 + 3\text{la}^2), \tag{2-5}$$

where δ is the deflection at the end of the beam, F is the load, l is the total length (125mm), and a is the distance from the pivot to the centerline of the actuator. Equations 2-3 and 2-4 can be found in beam formula tables [45] while Equation 2-5 can be derived by the superposition of a fixed cantilever with a concentrated end load and a fixed cantilever with a concentrated intermediate load.

Of the three models, the fixed end case in Equation 2-3 is the most conservative model so it was used in the design. However, in predicting the actual compliance as a function of length, it is

important to experimentally find a model that best fits the behavior of the beam. Experimental data was collected for a range of stiffness settings and applied loads. The data was used to calculate the stiffness at each setting and then to formulate a model for the overall Configurable Compliance behavior.

An actuator with Configurable Compliance was placed on a test stand with the ability to apply known loads in both tension and compression using a series of calibrated weights, up to 690 N (155 lbf). The actual force applied to the cantilever beam was directly measured with the actuator load cell (Futek LCM200) at a resolution of ±2 N. A dial indicator with a 0.0127 mm (0.0005 in) resolution was used to directly measure the deflection of the spring at the centerline of the linear actuator. The movable pivot was adjusted over five different settings (71 mm, 80 mm, 88 mm, 98 mm, and 108 mm) and at each length a series of loads were applied in both tension and compression (±93 N, ±187 N, ±280 N, ±458 N, ±687 N). The data from these experiments is shown in Figure 2-12. Only the tension loading case data is shown for simplicity.

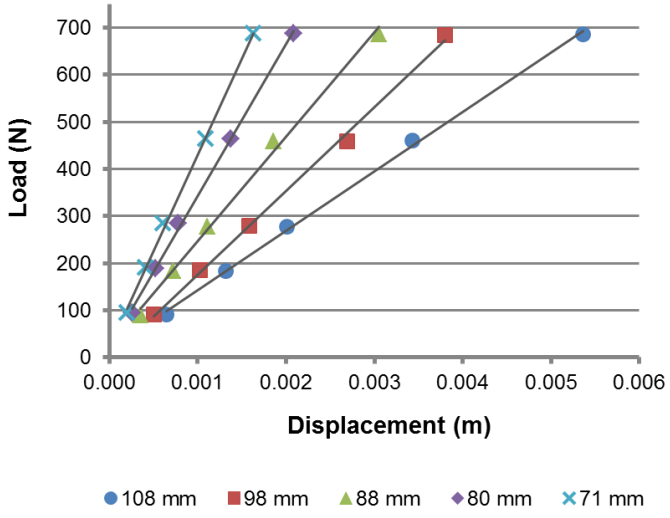

Figure 2-12. Experimental load vs. deflection of the SAFFiR Configurable Compliance design.

To see how the Configurable Compliance varies with length, the stiffness values at each setting were plotted and fit with a power curve as shown in Figure 2-13. Also plotted on the curve are the stiffness curves for each of the three theoretical models shown in Figure 2-11. The equation for the power curve fit of the experimental data is given as:

$$\mathrm{k}_{\mathrm{spring}} = 4.12\mathrm{e}10\, l^{-2.70}\,, \tag{2-6}$$

where l is the equivalent cantilever beam length measured in mm. The data and fit show that the Configurable Compliance design behavior is bounded by the fixed beam and the fixed-pinned beam models. Given this model for stiffness, the position of the movable pivot can be readily adjusted for any desired stiffness within the range.

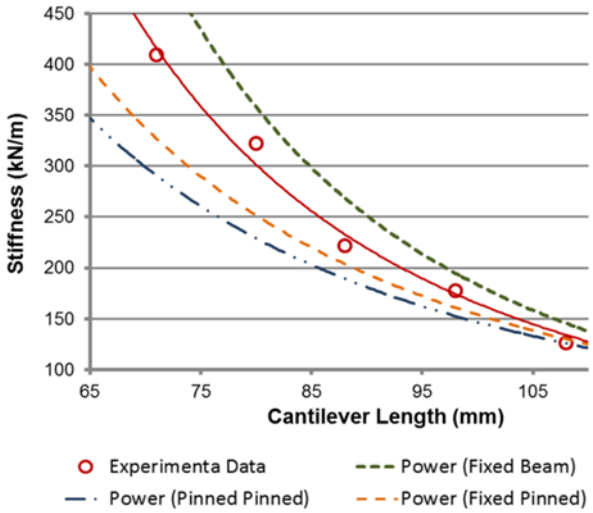

Figure 2-13. Stiffness vs. movable pivot position of the SAFFiR Configurable Compliance design.

2-4-3 SAFFiR Stiffness Tuning

Our approach to compliant walking on SAFFiR was to gradually incorporate the compliance starting with the ankles. The compliant beams in all but the ankle actuators were replaced with rigid members, shown in Figure 2-9. The four ankle actuators retained the titanium beams which were configured to the stiffest setting (65 mm cantilever). A force controlled balancing algorithm was developed and tuned to perform well on level ground. This balancing algorithm served as a baseline and standard by which to compare other stiffness settings.

The stiffness was then tuned by gradually increasing the effective length of the cantilevered beam until balancing performance was significantly affected. One of the concerns we had was added delay in the force control response due to the additional compliance. With several other sources of delay in the system such as communication and control loops, additional delay could significantly impact balancing and walking performance. It was experimentally determined that balancing performance was noticeably affected at a cantilever length of 70 mm, corresponding to a stiffness of 430 kN/m and a measured additional delay of 25-50 milliseconds in the force control response to an external disturbance. This stiffness setting was then used in developing a walking algorithm which has now successfully walked over strewn plywood, gravel, and even sand.

2-5 THOR Configurable Compliance - Moment Loading

The THOR Configurable Compliance design took a similar approach but was modified for the more stringent design requirements of the DRC. The new design goals were to double the maximum load and to improve the packaging footprint. The final arrangement and design is shown in Figure 2-14, Figure 2-15, and Figure 2-16.

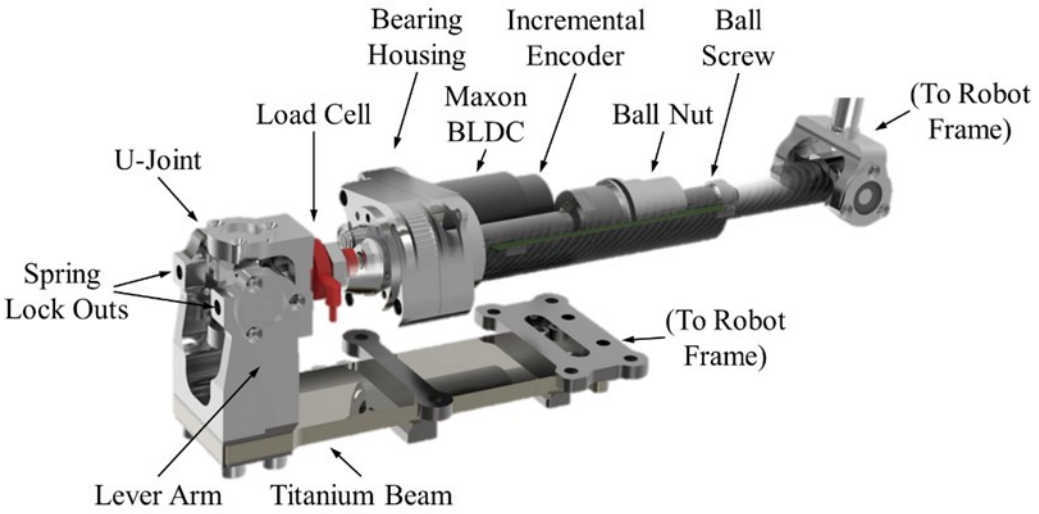

Figure 2-14. Rendering and schematic of the Configurable Compliance design for the THOR Series Elastic Actuators, used with permission of J. Holler.

The actuator acts on the compliant beam through a lever arm, shown in red in the cross section view, which converts an axial load into a moment load on the beam. The beam is arranged parallel to the actuator primary axis and is made of titanium for the same reasons listed in Section 2-4-1. The base clamp, shown in green, secures the beam to the robot link while the removable clamp, shown in blue, is used to configure the stiffness setting between a soft and stiff setting.

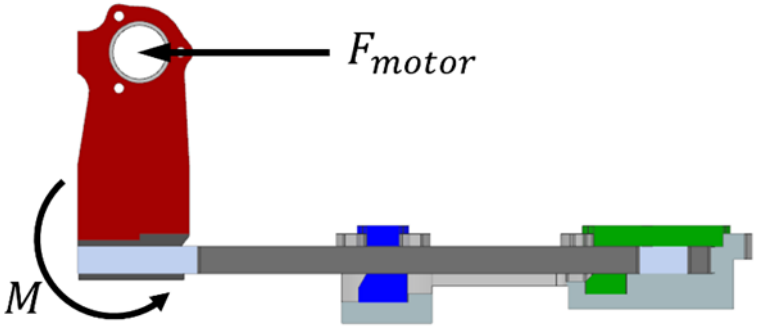

Figure 2-15. Cross section of the THOR Configurable Compliance design, used with permission of C. Knabe.

Only three stiffness settings are available, compliant, stiff, and locked out. The locked out setting is accomplished by rigidly bolting the lever arm into a structural member on the robot using two spring lock out bolts. Loading the beam with a moment loading allows for the entire length of the beam to store energy, enabling a max load of 2225 N at the desired stiffnesses. This arrangement also improves the packaging footprint by allowing a more compact and streamlined actuator.

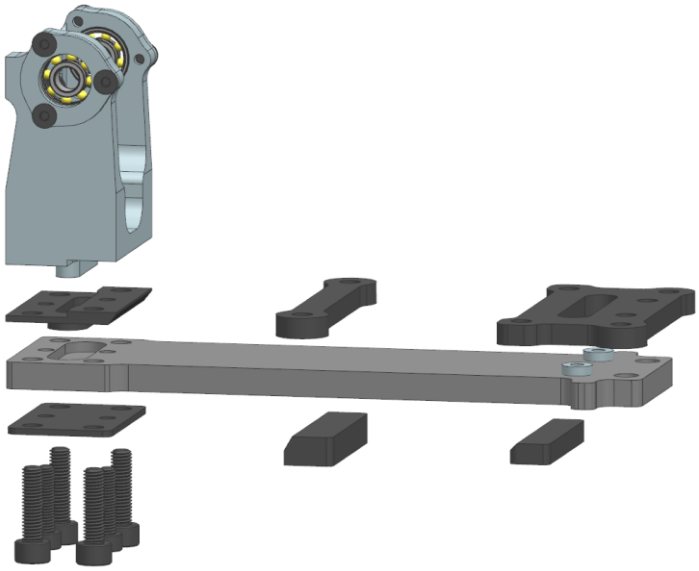

Figure 2-16. Exploded view of the THOR Configurable Compliance mounting and assembly, used with permission of C. Knabe.

2-5-1 Cantilevered Beam Design

One of the compromises made in the THOR Configurable Compliance design was only having three stiffness settings (compliant, stiff, locked out). This decision was based on space limitations and informed by our experience with the SAFFiR platform. Two desired stiffness settings of 380 kN/m and 760 kN/m were selected based on the tuning described in Section 2-4-3. The compliant setting was selected to be 50 kN/m softer in expectation of THOR's more dynamic walking gait. The stiff setting was then selected to be twice the stiffness of SAFFiR's stiffest setting due to THOR's weight and to provide a stiffer setting for initial testing.

Figure 2-17 shows a beam of length l under moment loading M due to a force F acting through lever arm h. The deflection in y of the beam due to a moment is given by:

$$\delta_{max} = -\frac{\mathrm{Ml}^2}{2\mathrm{EI}}. \tag{2-7}$$

And the angle θ at the end of the beam is given by:

$$\theta = \frac{\mathrm{Ml}}{\mathrm{EI}}. \tag{2-8}$$

Note that the actuator attachment point Q deflects both in the x and y directions which can be related to δ and θ through h and l. The effective linear spring stiffness 'felt' at the actuator is here defined by the actuator length change that would cause a given deflection and force in the spring.

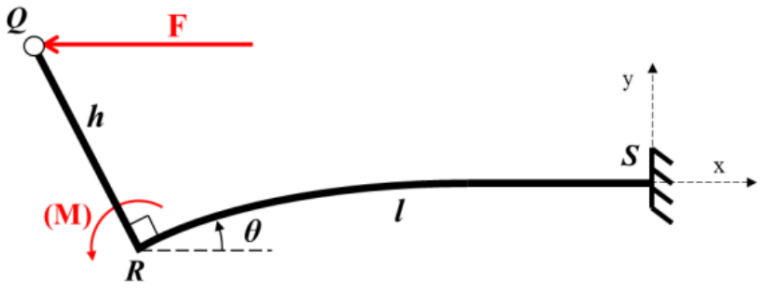

Figure 2-17. Schematic for a cantilevered beam under moment loading, used with permission of C. Knabe.

Using the packaging and loading constraints, as well as the desired stiffness settings, the selected beam dimensions were a width of 38 mm, thickness of 5 mm, and two length settings of 110 mm and 55 mm for the compliant (380 kM/m) and stiff (760 kN/m) settings, respectively. Figure 2-18 shows how the effective stiffness of this arrangement varies with cantilever length.

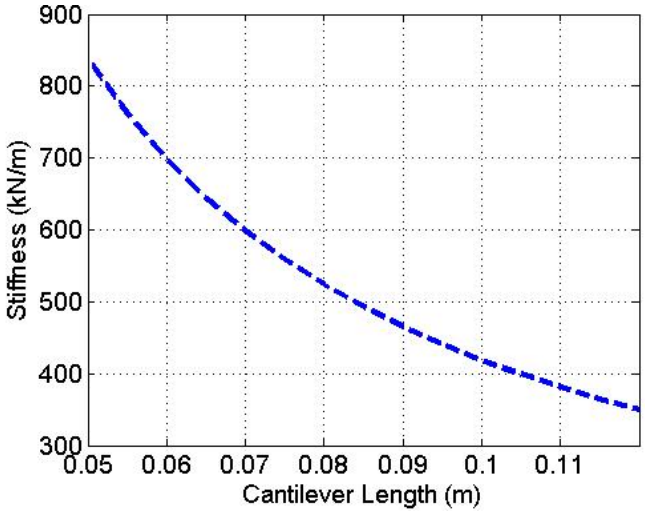

Figure 2-18. Simulated effective stiffness of cantilevered beam under moment loading.

An actuator with this beam design was placed on a test stand with the ability to apply known loads in both tension and compression using a series of calibrated weights, up to 1543 N. The actual force applied to the cantilever beam was directly measured with the actuator load cell (Futek LCM200) at a resolution of ±2 N. A dial indicator with a 0.0127 mm (0.0005 in) resolution was used to directly measure the deflection of the spring in both the x and y directions. A series of loads were applied at the two stiffness settings using calibrated weights to achieve loads of ±51 N, ±104 N, ±155 N, ±207 N, ±255 N, ±259 N, ±511 N, ±767 N, ±1023 N, ±1278 N, and ±1543 N. The data from these experiments is shown in Figure 2-19.

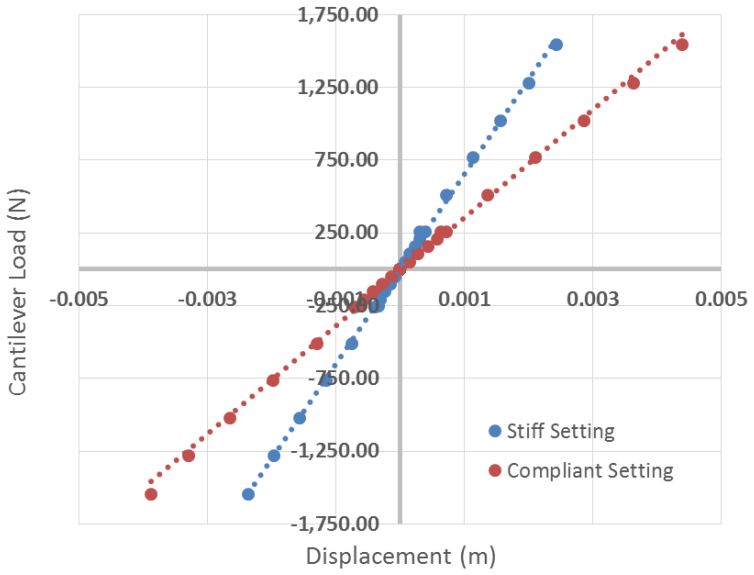

Figure 2-19. Experimental stiffness experiments for moment loading Configurable Compliance. Displacement represents the total actuator length change, for the given load case.

The experimental stiffness of the two settings was found by fitting a linear regression to the data, the slope of which would be the spring stiffness:

$$F = 371643\ \Delta l\ , \tag{2-9}$$

$$F = 655399\ \Delta l\ . \tag{2-10}$$

Rounding to the nearest thousand, the experimental stiffness values are 372 kN/m and 655 kN/m. Note that both of these stiffness values are lower than we originally designed for (380 kN/m and 760 kN/m). We believe this is due to additional deflection happening within the u-joints as well as deflection in the cantilever clamping method.

2-6 Discussion

While both approaches have their tradeoffs, the two different Configurable Compliance designs described in this paper are both well suited their specific applications. Figure 2-20 shows the stiffness range and values covered by the two designs.

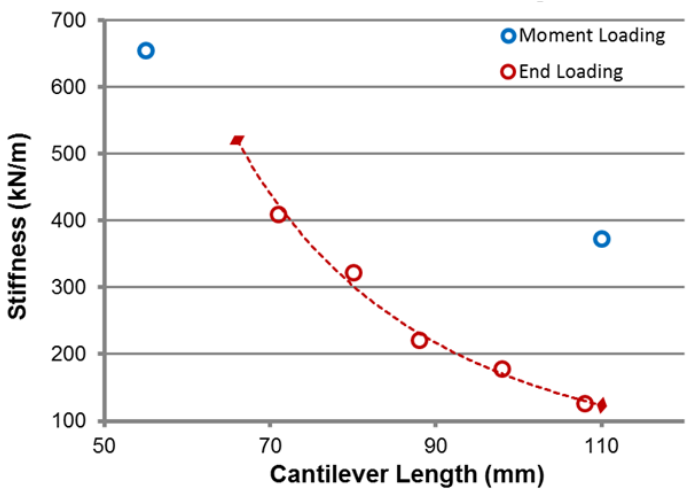

Figure 2-20. Available stiffness settings for the two Configurable Compliance designs.

However, it is important to keep in mind that there are other features which are also important depending on the desired application. Table 4 lists some of the advantages of the two approaches. Despite the differences, the most important feature of both designs is the ability to practically implement a nonuniform distributed stiffness. The compliance is therefore treated as a joint level component instead of simply a component within an actuator. By adjusting the compliance for each joint, the springs can be tuned for each joint's different power requirements, mechanical advantage, velocity profiles, and bandwidth requirements.

Table 4: Comparison of the two Configurable Compliance approaches.

SAFFiR Design | End Loading	THOR Design | Moment Loading
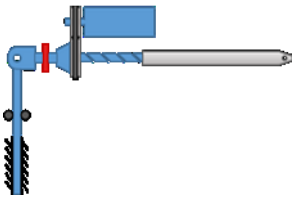	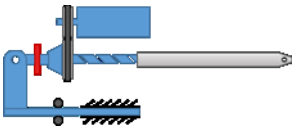
Max Load: 1000 N	Max Load: 2225 N
Perpendicular Arrangement	Parallel Arrangement
Easy access for adjustments	Compact and high energy storage capacity
Large stiffness range with fine resolution adjustability (145-512 kN/m)	Three discrete settings (372 kN/m, 655 kN/m, and locked out)

3 An Unlumped Model for Linear Series Elastic Actuators with Ball Screw Drives

3-1 Abstract

Series elastic actuators are frequently modeled using a conventional lumped mass model which has remained mostly unchanged since their introduction almost two decades ago. The lumped model has served well for early development but more descriptive models are now needed to compare new actuator designs and control approaches. In this paper we argue that the lumped model is an insufficient representation of how Series Elastic Actuators are used in practice. We propose a new unlumped model for linear Series Elastic Actuators which uses a rack & pinion to intuitively depict the mechanics of a ball screw drive. Results from real hardware experiments are compared to simulation results which demonstrate that the new model is significantly more representative of the true actuator dynamics.

3-2 Introduction

Series Elastic Actuators (SEAs) come in a variety of configurations but are generally characterized by a physical spring placed in series with the output. Numerous studies have addressed the various benefits of series elasticity for force control in a variety of robotic applications [25], [46], [47]. When SEAs were first introduced [25], [48], they were depicted as a simple mass-spring-damper, shown in Figure 3-1. The motor rotor and transmission inertias are combined into a single lumped sprung mass, m_k, and assumed to always move together. This same representation, often referred to as the "lumped model" in the literature, was used to analyze the dynamics of several different designs including electric rotary, electric linear, and hydraulic linear SEAs [49]. In the years since, there has been an increasing variety of SEA designs including: rotary actuators [50], [51] and linear actuators [32], [33], [52]; novel design features such as nonlinear stiffness [53], variable damping [54], variable stiffness [55], continuously variable transmissions [56], and clutches[57]; and diverse control approaches [52], [58]–[61].

In spite of the increasingly more complex designs and more sophisticated controllers, the same basic lumped model has continued to be used almost exclusively. With minor variations, the general trend when modeling SEAs is to make the following three assumptions: one side of the actuator is treated as the output with the other side serving as ground, the motor and transmission inertias are combined into a single lumped inertia, and the actuator dynamics are derived with the motor torque as the input and the force on the assumed output side as the output. As a result, most of the literature never considers the forces acting on the "ground" side of the actuator, overemphasizes the importance of spring location, and oversimplifies the actuator dynamics.

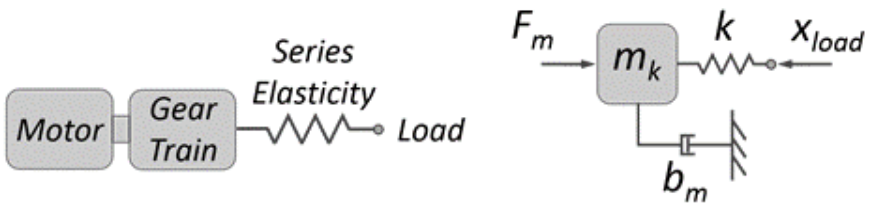

Figure 3-1. Early lumped models for SEAs. F_m is the motor force, m_k is the lumped sprung mass, b_m is the lumped damping, and k is the stiffness of the physical spring placed in series.

We argue that for most robotic applications the selected model should include an actuator acting on two links, each of which can be either fixed or moving. Because of this, the forces acting on both links should be considered and the actuator model should be capable of properly representing these forces. We show that, depending on the mechanical design, the lumped mass simplification can often lead to drastically misleading conclusions about the actuator dynamics. The simple lumped model served well for early development, but we believe more descriptive models are now necessary to better understand the true actuator dynamics and the impact that a given design approach has on performance.

In this paper, we present a new unlumped actuator model for ball screw driven linear Series Elastic Actuators. The model includes a linear SEA acting on two links, each of which are allowed to be fixed or moving. The ball screw drive is modeled with an intuitive rack & pinion representation which retains the important actuator dynamics that are oversimplified in the lumped model. Both models are compared with experimental results which demonstrate that the new unlumped model is a significantly better representation of the true actuator dynamics.

3-2-1 Related Work on Moving Output and Unlumped Models

The majority of the literature on SEAs limits dynamic analysis to the simplified case described above; however, there are a few notable exceptions:

Reaction Forces: In [52], the authors present the UT-SEA, a ball screw driven linear SEA that places the spring element between the motor housing and "ground". Since the spring is also used as the force sensor, the authors investigate the forces acting on the output side as well as the reaction forces acting on the ground side. However, the analysis uses a lumped actuator model, leading to some incorrect conclusions which we address in Section 3-5.

Rack & Pinion Model: In [62], the authors depict the leg of a running robot using a rack & pinion model in order to investigate the energetics of series compliance during impacts. While similar to the model we present below, their analysis is focused on analyzing a cable driven leg as opposed to an SEA and is limited to collisions and energy storage.

Lever Model: In [63], the authors present a lever model to compare what they refer to as distal compliance and proximal compliance. The lever fulcrum is allowed to have a mass which is referred to as being the transmission housing mass. This could be considered an unlumped model; however, a lever is less intuitive and only applicable when considering small motions. A similar lever model is used in [64] except with a massless fulcrum.

3-2-2 Model Simplicity vs. Fidelity

Selecting an appropriate model is a critical step in the development of robotic actuators because it is used in making design decisions, evaluating performance, and designing controllers. An effective model can simplify the analysis and provide intuitive insights for a complex system while an oversimplified model can lead to poor design decisions and misleading analysis. A simpler model is generally preferred as long as it does not neglect any important system dynamics. The challenge is that any model introduces an inherent tradeoff between complexity and fidelity. The key is understanding and retaining the most prominent dynamics while eliminating the elements that can be neglected.

3-2-3 Ball Screw Driven Linear SEAs

Selecting an appropriate model often depends on the specific application. In this paper, we are addressing the modeling of ball screw driven linear Series Elastic Actuators as they are used in legged robots. Most linear SEAs use ball screw drives because they have high efficiencies, low backlash, and are robust to impacts. However, the rotary-to-linear transmission of a ball screw can be difficult to represent in a way that is intuitive. In this paper, we introduce a rack & pinion representation which is an almost direct analogy for the rotary-to-linear behavior of ball screws.

Figure 3-2 shows the linear SEA that is used in THOR, a 1.78 m 34-DOF torque-controlled humanoid [33]. Similar to other designs, we use a 100 W BLDC motor to drive a ball screw with a 0.002 m lead through a 3:1 belt reduction. This configuration provides a maximum speed of 0.19 m/s, a continuous force of 685 N, and a peak force of 2225 N. Universal joints at either end of the actuator simultaneously serve as the anti-rotation mechanism for the ball nut and as connection points attaching the actuator to the robot structure. One of the novel features of this design is that the spring element is implemented as a cantilevered titanium beam. The actuator applies forces through a lever arm to load the beam with an almost pure moment. A common practice with SEAs is to also use the spring as a sensor by measuring the spring deflection to estimate force [48], [49]. Instead of measuring spring deflection, we use a relatively rigid in-line load cell.

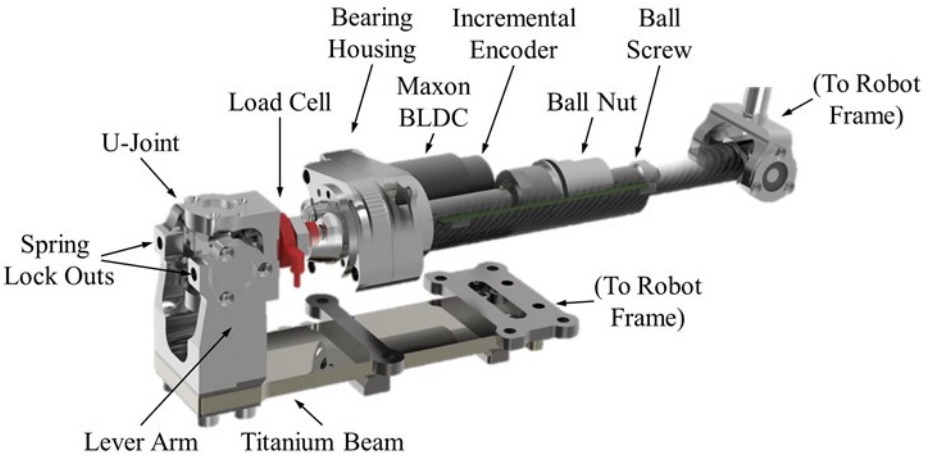

Figure 3-2. Schematic of the THOR Linear SEA used in the lower body of THOR. Two sets of parallel actuators power the hip roll/yaw and ankle pitch/roll DOF. A modified version of this design with an inverted Hoekens linkage output powers the hip pitch and knee pitch [34]. Used with permission of J. Holler

3-3 Unlumped Rack & Pinion Model

In an effort to evaluate the THOR Linear SEAs, we have developed a new unlumped model that addresses the two major limitations of the conventional lumped model: the assumption of a ground link and lumping of the translational and rotational inertias.

3-3-1 Changing Ground

The conventional approach to modeling SEAs is to treat one side as being rigidly attached to ground with the other side serving as the output. The actuator dynamics are then derived with the motor torque as the input and the force on the assumed output link as the output. The reaction forces acting on the "ground" link are typically not considered, with [52] being a rare exception. While this approach works well for evaluating an actuator on a test stand, it is not representative of how the actuators are commonly used in legged robots.

In reality, actuators in legged robots act on two links; each of which are allowed to move subject to joint constraints and have external forces acting on them as shown in Figure 3-3a, where m_1 is the linearized inertia of Link 1, m_2 is the linearized inertia of Link 2, m_k is the lumped sprung mass of the actuator, F_1 and F_2 are the external forces acting on the links either due to internal dynamics or external disturbances, F_m is the linearized motor force, b_m is the damping in the motor and transmission, and k is the spring stiffness. While there is a growing interest in including physical damping in parallel with the spring element [47], [54], we do not include a spring damping term because SEAs are typically designed such that the spring element has negligible friction.

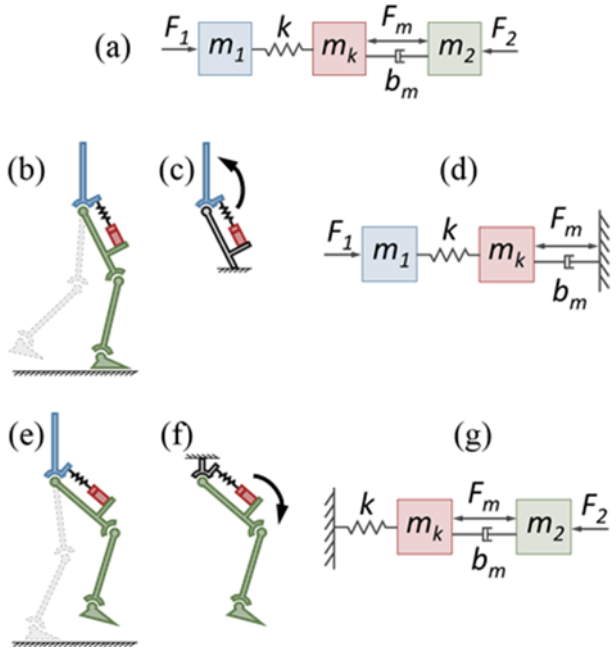

Figure 3-3. Two-link moving output models for a legged robot with changing ground contacts. The general model (a) shows two moving links. The model can be simplified for the stance phase (b, c, and d) or the swing phase (e, f, and g).

As a case study, we can consider the scenario of a hip pitch actuator that acts across the torso and thigh links. In this scenario, the torso inertia can be represented by m_1, the leg inertia can be represented by m_2, and both ends can be thought of as being the fixed end or the moving output depending on the phase of the walking cycle. In the stance phase (Figure 3-3b), it would be a reasonable approximation to treat the thigh as ground (Figure 3-3c), resulting in the actuator model being drawn with the ground on the right side (Figure 3-3d). However, in the swing phase (Figure 3-3e), the opposite is true. The torso is the most appropriate link to treat as ground (Figure 3-3f) and the actuator model is drawn with the ground on the left side (Figure 3-3g).

Since each link varyingly serves as both the output link and ground link, it is important to consider the forces acting on both sides of the actuator. If, as [52] argues, these forces are not equal for SEAs, there are some major design and control implications. Considering SEAs are most often used as force (or torque) controlled actuators, the actual forces being exerted on the links is a critical parameter for state estimation and control.

Interestingly, the literature places a strong emphasis on the location of the spring element with respect to the assumed "output" side. The early literature on SEAs describes placing the spring element between the motor and the "output" while some more recent designs describe placing the spring element in between the motor and "ground" [52], [65]. Both [52] and [65] refer to their designs as novel implementations specifically because of the spring placement, referring to the new arrangements as Reaction Force Sensing SEA and Force Sensing Compliant Actuator, respectively. Similarly, in [63], the authors use the terms distal compliance and proximal compliance to describe what they consider to be two different SEA implementations based entirely on the location of the spring. However, as we show in Figure 3-3, any SEA in a legged robot will operate in both modes depending on the phase of the walking cycle.

3-3-2 Lumped Mass & Inertia

The second limitation of the conventional lumped model is the lumping of the sprung mass (or inertia) with the reflected inertia of the rotating transmission elements. For a rotary actuator, this would be lumping the motor rotor inertia with the inertia related to the internal rotation of the transmission elements. For a ball screw driven linear actuator, this is lumping the translational inertia of the actuator's sprung mass with the rotational inertia of the motor rotor and ball screw. The model variables in Figure 3-4 are given by

$$m_k = m_l + \frac{J}{r^2}, \tag{3-1}$$

where m_k is the lumped sprung mass, m_l is the translational sprung mass, J is the rotational inertia of the motor and transmission, and r is the gearing ratio given by

$$r = \frac{l}{2\pi N_p}, \tag{3-2}$$

where l is the ball screw lead and N_p is the pulley reduction ratio. For the high gearing ratios typical of robotic actuators, m_k will be much larger than m_l. The linearized motor force F_m is given by

$$F_m = \frac{\tau_m}{r} = \frac{K_\tau}{r} I \,, \tag{3-3}$$

where τ_m is the motor torque, K_τ is the motor torque constant and I is the motor current. Similarly, the linearized motor damping b_m is given by

$$b_m = \frac{b_r}{r^2} \,, \tag{3-4}$$

where b_r is the rotational motor damping due to friction in the motor and transmission.

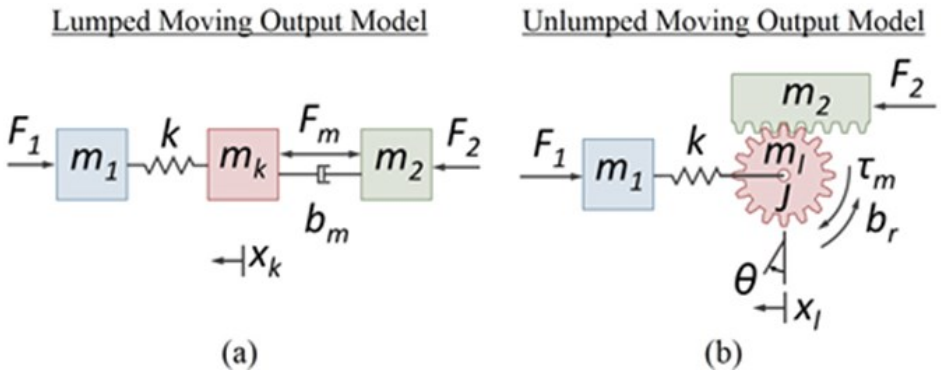

Figure 3-4. Lumped and unlumped models for linear SEAs showing the moving output test case.

3-3-3 Unlumped Model

Instead of lumping the mass and inertia, we propose an unlumped rack & pinion model as shown in Figure 3-4b. In place of the sprung lumped mass m_k, we use a pinion gear with radius r, translational mass m_l, and rotational inertia J. The rack & pinion model provides an intuitive representation of the coupling between the rotational and translational motions as well as the no-slip constraint of the ball screw and ball nut given by

$$x_l - x_2 = \theta r \,. \tag{3-5}$$

An observation worth noting is that even the unlumped model demonstrates a lumped motion behavior in the high impedance test case which is widely used as a simplified test case to derive the dynamics of SEAs. In the high impedance test case, both sides of the actuator are fixed to ground as shown in Figure 3-5 where $\dot{x}_1 = \dot{x}_2 = 0$. The no-slip condition of the rack & pinion model ensures that for a given pinion rotation (motor/transmission rotation) there is a known pinion translation (motor/transmission translation) which is a function of the gearing ratio. In other words, the translational and rotational inertias of the actuator move in a lumped manner for the high impedance test case.

However, this lumped motion does not hold true for the moving output test case, (i.e. if we allow one or both of the output links to move). Consider the unlumped model in Figure 3-4b for an actuator in force control mode with low gains (i.e. slow response). If m_1 is fixed and an impulse load is applied to m_2, the translational mass m_l of the pinion can accelerate immediately with m_2 through the spring deflection, without necessarily accelerating the large rotational inertia J. A

related observation is to note that the lumped model depicts a much larger linear sprung mass ($m_k + m_2$) than the unlumped model ($m_l + m_2$).

3-3-4 High Impedance Model Comparison

The most common approach for evaluating SEA dynamics is to assume the high impedance test case as shown in Figure 3-5.

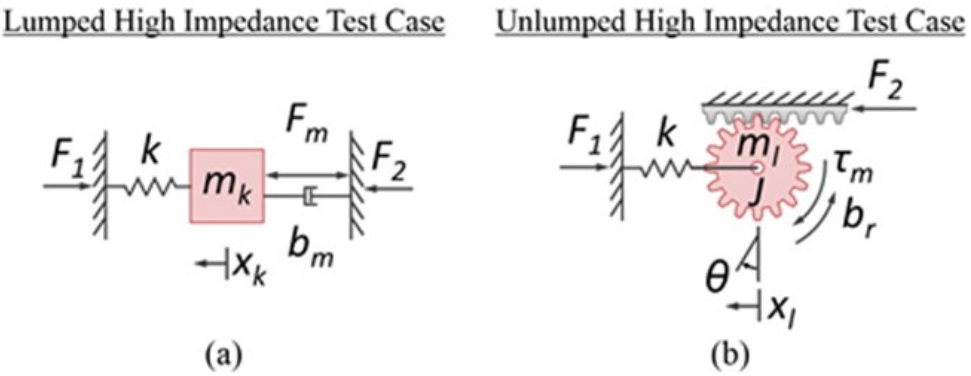

Figure 3-5. Lumped and unlumped models for linear SEAs showing the high impedance test case.

The equivalent experimental setup is shown in Figure 3-6. In this test case we are primarily interested in comparing how the two models relate the input motor force (F_m) to the two forces acting on either end of the actuator (F_1 and F_2) resulting in three responses to consider: F_1/F_m, F_2/F_m, and F_1/F_2. The following transfer functions for the lumped model can be readily derived from standard system dynamics relationships

$$\frac{F_1(s)}{F_m(s)} = \frac{k}{m_k s^2 + b_m s + k}\ , \tag{3-6}$$

$$\frac{F_2(s)}{F_m(s)} = \frac{m_k s^2 + k}{m_k s^2 + b_m s + k}, \tag{3-7}$$

$$\frac{F_1(s)}{F_2(s)} = \frac{k}{m_k s^2 + k}\ . \tag{3-8}$$

The unlumped model dynamics are given by

$$\frac{F_1(s)}{\tau_m(s)} = \frac{kr}{\left(\frac{J}{r^2} + m_l\right)s^2 + \left(\frac{b_r}{r^2}\right)s + k}, \tag{3-9}$$

$$\frac{F_2(s)}{\tau_m(s)} = \frac{(m_l s^2 + k)r}{\left(\frac{J}{r^2} + m_l\right)s^2 + \left(\frac{b_r}{r^2}\right)s + k}\ . \tag{3-10}$$

To allow for a more direct comparison, we can use the linearized motor force F_m instead of τ_m by substituting (3-1), (3-3), and (3-4) into (3-9) and (3-10). The equivalent responses for the unlumped model are then given by

$$\frac{F_1(s)}{F_m(s)} = \frac{k}{m_k s^2 + b_m s + k}, \tag{3-11}$$

$$\frac{F_2(s)}{F_m(s)} = \frac{m_l s^2 + k}{m_k s^2 + b_m s + k}, \tag{3-12}$$

$$\frac{F_1(s)}{F_2(s)} = \frac{k}{m_l s^2 + k}. \tag{3-13}$$

3-3-5 Initial Observations

The lumped and unlumped models have identical F_1/F_m responses in (3-6) and (3-11). This is because even the unlumped rack & pinion model demonstrates a lumped motion behavior in the high impedance test case. It is no surprise, therefore, that the conventional lumped model has been successful in the early development of SEAs. Most of the existing SEA designs use the spring as the force sensor so the actuator dynamics are most commonly derived as the F_k/F_m transfer function, which for our design is equivalent to the F_1/F_m transfer function.

However, as we have shown in Section 3-3-1, it is important to consider the forces acting on both links. Inspecting the remaining equations (3-7), (3-8), (3-12), and (3-13), reveals that while they have a similar form, the two models do predict a different relationship for F_2, the force acting on Link 2. In the lumped model, $F_2(s) = m_k s^2 + k$, while in the unlumped model the same relationship is $F_2(s) = m_l s^2 + k$. For a highly geared actuator, the difference between m_k and m_l can be large, which would lead to very different predictions of the forces acting on Link 2.

3-4 Results

We begin the analysis below by considering the high impedance test case where an actuator is fixed on both ends. Experimental system identification data of a THOR Linear SEA is compared to the responses predicted by the lumped and unlumped models. We then consider the moving output test case where one end of the actuator is fixed and the other end is attached to a pendulum output. Experimental data of a THOR Linear SEA is again compared to the predicted moving output responses of the two models.

To aid the analysis, the most general lumped and unlumped models shown in Figure 3-4a and Figure 3-4b were simulated in Simulink by MathWorks. The Simulink models can be configured to set initial conditions, rigidly fix either or both of the links, and apply external disturbances. Table 5 lists the known and extracted model variables used in the Simulink models. The known variables were either taken from component datasheets or directly measured. The remaining unknown variables were extracted using experimental system identification as described in Section 0.

3-4-1 System Identification – High Impedance Test Case

The system identification was performed with a THOR Linear SEA on a rigid test stand with both ends fixed as shown in Figure 3-6. The actuator's inline load cell at the spring interface directly measures the spring force, which in our configuration is equivalent to measuring F_1. A second load cell was attached at the other fixed end, enabling simultaneous measurement of both F_1 and F_2.

An open loop force excitation (current input with feedforward constant) was applied to the actuator and the force responses of both load cells were recorded at a sampling rate of 400 Hz. The excitation signal consisted of a sinusoidal chirp signal summed with a noise signal. The chirp signal had an amplitude of 480.6 N (2 A) and used a logarithmic sweep from 0.01-150 Hz over 164 seconds. The noise signal was 150 Hz band limited white Gaussian noise with a standard deviation of 240.3 N (1 A). An H1 estimate was used to compute the frequency response function for each of the following three responses: F_1/F_m, F_2/F_m, and F_1/F_2. The results of this experiment are shown in Figure 3-7.

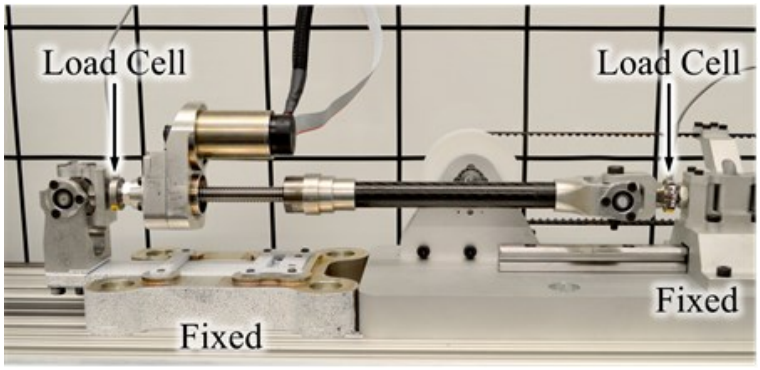

Figure 3-6. THOR Linear SEA in the high impedance test case with a load cell at either end.

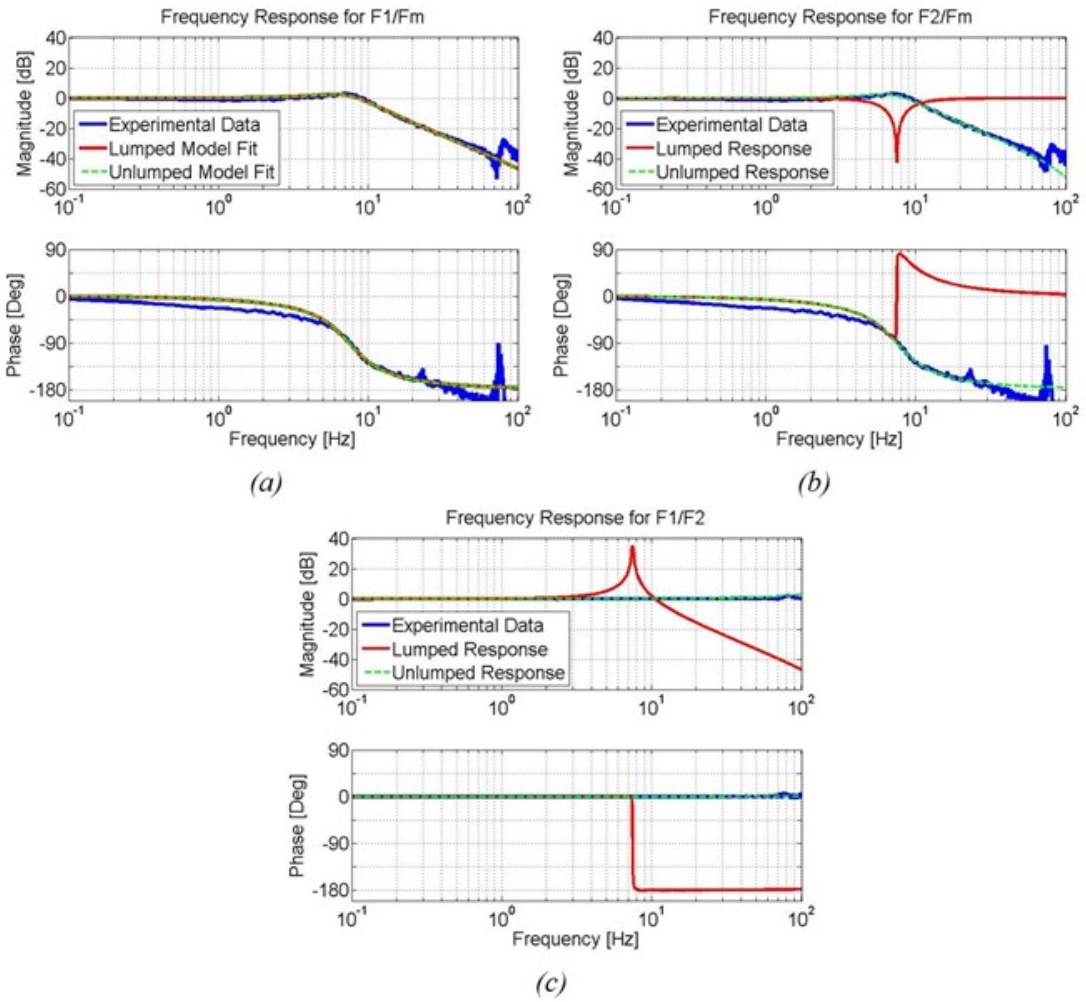

Figure 3-7. Open loop frequency response from Simulink simulations and experimental system identification for F_1/F_m (a), F_2/F_m (b), and F_1/F_2 (c).

3-4-2 Fitting a Model to the F_1/F_m Response

The unknown model variables were extracted using a second order fit of the experimental F_1/F_m response. A second order fit was selected in order to match the form of (3-6) and (3-11). The F_1/F_m response was selected because it is identical in both models, is the response that is most commonly considered in other literature, and is the response that we use in our force controller. The second order fit and its relationship to (3-6) and (3-11) is given by

$$\frac{F_1(s)}{F_m(s)} = \frac{2166}{s^2+39.53s+2279} = \frac{\frac{k}{m_k}}{s^2+\frac{bm}{m_k}s+\frac{k}{m_k}} \quad . \tag{3-14}$$

The unknown variables were extracted by equating the coefficients in (3-14) and using the known variables from Table 5. Figure 3-7 shows the experimental system identification data as well as the two model fits, which as stated before are identical.

Table 5: Known and Extracted model Variables

Known Variables			
m_l	$0.6461\ kg$	K_τ	$0.0255\ Nm/A$
k	$655000\ N/m$	r	$1.061x10^{-4}\ m$
l	$0.002\ m$	N_p	$3:1$
Extracted Variables			
m_k	$294.9\ kg$	J	$3.312x10^{-6}\ kgm^2$
b_m	$11658\ Ns/m$	b_r	$1.3123x10^{-4}\ Nms$

We note here that there is an unexpected resonance in the experimental response at 80 Hz. We believe that this is due to an axial bending mode in the ball screw. When tested, the resonance moved between 70 Hz and 80 Hz depending on the length of the actuator.

3-4-3 Comparing the F_2/F_m and F_1/F_2 Responses

The experimental system identification data and Simulink models were also used to compare the F_2/F_m and F_1/F_2 responses. The F_2/F_m response relates the motor force to the force on Link 2 while the F_1/F_2 response relates the force on Link 2 to the force on Link 1. While the equations for these relationships for the lumped and unlumped models are very similar, Figure 3-7 shows that the simulated responses are very different. It is also apparent in the figure that the experimental data clearly supports the unlumped model as the more appropriate representation of the true actuator dynamics.

F_2/F_m – *Lumped Model:* Figure 3-7 shows that the lumped model predicts a flat response across the frequency spectrum except for a sharp dip at the system's natural frequency. The dip occurs because the lumped mass is moving significantly at the resonance, resulting in a relatively large damping force which counteracts the motor force at Link 2. At every other frequency, the damping force is relatively small and the full motor force is felt at Link 2. Equation 3-7 shows that when $\dot{x_k}$ is small, the transfer function approaches unity. By intuition, $\dot{x_k}$ will be small for low frequencies because the sprung mass is not moving rapidly and will also be small at high frequencies due to the mechanical low pass filtering effects of the spring.

F_2/F_m – *Unlumped Model:* Figure 3-7 shows that the unlumped model predicts a response for F_2/F_m that is very similar to F_1/F_m in Figure 3-7. This effect can be seen by considering (3-12), which for small values of m_l is approximately the same as (3-11).

F_1/F_2 – *Lumped & Unlumped Models:* Figure 3-7 shows that the lumped model predicts very different forces acting on Link 1 and Link 2 while the unlumped model shows an almost completely flat response. Inspecting (3-8) and (3-13) reveals that the difference is that the lumped model has m_k in the denominator while the unlumped model has m_l in the denominator. Note that m_k is about 450 times larger than m_l (294.9 kg vs. 0.6461 kg), resulting in the drastic differences in the responses. In reality, both models predict the same qualitative behavior for F_1/F_2; the only difference is that the unlumped model has a much smaller sprung mass (m_l vs. m_k) so the resonance occurs at a much higher frequency (160 Hz vs. 7.47 Hz). Force controllable actuators

such as SEAs are typically operated at bandwidths between 15-60 Hz, so the unlumped model shows that the F_1/F_2 transfer function can be approximated as being unity across the relevant bandwidth range. An important implication of this finding is that the force sensor can be placed at either end of the actuator and still provide an accurate estimate of the joint torque.

3-4-4 Intuitive Interpretations

Sprung Masses: One way to consider these results intuitively is to consider the sprung mass depicted in each model. In the real system, the sprung mass acts as an intermediate inertia, the dynamics of which influence the forces acting on the two links to which it is attached. The lumped model shows a very large lumped sprung mass which, if accelerating, will result in different forces at the two links. The unlumped model, on the other hand, shows a relatively small sprung mass, keeping the two forces almost equal.

Load Path: The lumped model shows that any forces propagating from Link 1 to Link 2 would have to act through the motor damping or motor force. The rack and pinion model, however, shows an almost direct feedthrough of forces from Link 1 to Link 2. The only difference in forces would be associated with the acceleration of the small translational mass m_l. This can be demonstrated mechanically by considering that there is a direct load path for forces to propagate axially through the ball screw without having to first accelerate the motor and transmission inertias.

Reaction Forces: Another way to consider the results in Figure 3-7 intuitively is by looking at the predicted reaction forces acting on Link 2. In the lumped model, the full motor reaction force F_m acts on Link 2. In the unlumped model, the motor reaction force does not act on Link 2 directly. Instead, it acts out of plane between the pinion gear and a virtual ground. The mechanical equivalent of this effect has to do with how the reaction torque is handled. In a ball screw driven linear actuator, the reaction torque of the motor is counteracted by the anti-rotation mechanism that prevents the ball nut from spinning freely. Various designs implement the anti-rotation mechanism via guiderails, keyways, or in the case of the THOR Linear SEA, with u-joints at either end. Irregardless of the mechanism, the motor reaction torque is either transmitted into torsion in the structure of the actuator or torsion in the structure of the robot.

3-4-5 Moving Output Results

As discussed in Section 3-3-5, both the lumped and unlumped models predict the same F_1/F_m response for the high impedance test case because the pinion's translational and rotational inertias move in a lumped manner due to the no-slip condition. However, this relationship breaks down if either or both of the links are allowed to move. The pinion's translation and rotation are still coupled through the no-slip condition of the ball screw; however, there can now be translation without rotation and vice versa, depending on the loading conditions.

Testing a moving output test case with linear actuators is difficult in practice due to the limited stroke of the actuator. As shown in Figure 3-8, the actuator was set up much like a knee joint in the swing phase. One end of the actuator was rigidly fixed to the test stand (thigh) with the other end acting on a 0.0725 m lever arm (knee) that moved an output link (shin). A 4 kg calibrated weight was added to the end of the 0.300 m output link to simulate the mass of a leg. Because this

added pendulum dynamics to the system, the Simulink models were updated to include a basic pendulum model.

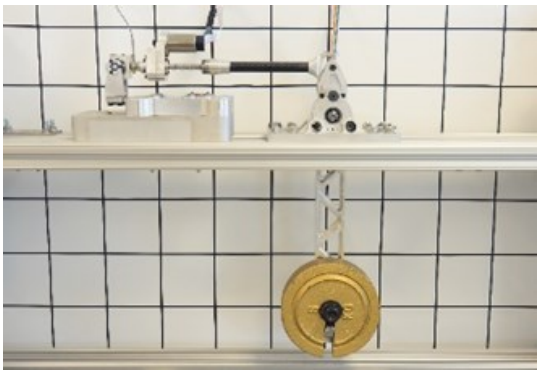

Figure 3-8. THOR Linear SEA in the moving output test case with a 4 kg mass.

The system identification was performed with an open loop force excitation consisting of 150 Hz band limited white Gaussian noise with a standard deviation of 841 N (3 A). The procedure was otherwise identical to the one described in Section 0. Figure 3-9 shows the experimental results as well as the predicted frequency response of the lumped and unlumped models. Note that the dip in the magnitude response at 1 Hz is due to the natural frequency of the pendulum and is reflected in both the simulation and experimental data. As the data clearly shows, the unlumped model is a much closer match to the experimental data, correctly predicting the peak resonance frequency while the lumped model shows a significantly lower frequency for the resonance (15.76 Hz vs. 6.35 Hz) due to the different sprung masses depicted in the models.

This is a significant finding because the lumped model's F_1/F_m response, which was correct in the high impedance test case, is no longer valid for the moving output test case. In [47], [59], [66]–[68], the authors consider the effects of a moving output on the dynamics and performance of compliant actuators but they all use the lumped mass model and only consider the forces acting on one side of the actuator. Given the above results, it would be of significant interest to revisit the findings in those papers using the unlumped rack & pinion model.

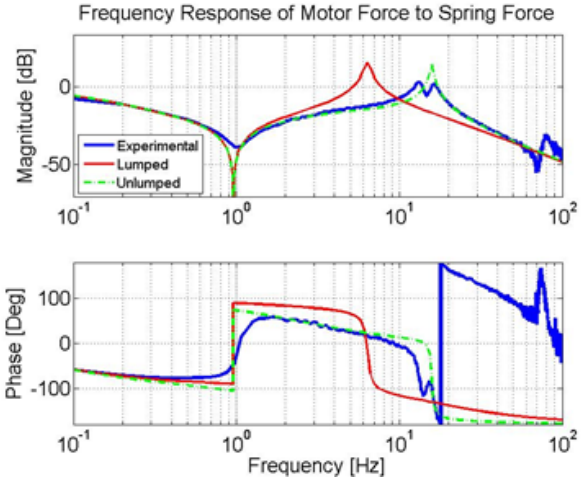

Figure 3-9. Open loop frequency response from Simulink simulations and experimental system identification for F_1/F_m with a moving output.

3-5 Discussion

We can summarize the above findings with the following observations:

1) The widely used lumped model is only a good predictor of the actuator dynamics for the high impedance test case, and even then, only if motor torque is the input and spring force is the output. Since this is how SEAs are commonly considered, it is no surprise that the lumped model has served so well for so long. However, if the desired output is the force acting on the other link (non-spring side) or if either/both of the links are allowed to move, the lumped model is misleading and could lead to poor design decisions and poor force control strategies. Since a moving output is the typical use case for robotic actuators, we argue that the conventional lumped model needs to be amended with more descriptive unlumped models.

2) We have shown in Section 3-3-1 that it is important to consider the forces acting on both links, especially for legged robot applications where the ground contact changes. The literature on SEAs, however, only considers the forces acting on the side with the spring element, which is typically assumed to be the output side. An exception is [52], in which the authors consider both forces for their linear ball screw driven SEA. However, because the conventional lumped model was used, the conclusions incorrectly show drastically different spring forces and output forces, much like the lumped response we show in Figure 3-7.

3) As SEA designs and controllers become more complex and sophisticated, more descriptive models will be necessary to avoid the potential pitfalls of an overly simplified model such as poor design choices (location of spring, transmission selection, etc.) and poor controller strategies (incorrect plant models, poor state estimation, etc.). The unlumped rack & pinion model in this paper is an effort to provide a more descriptive representation of ball screw driven linear Series Elastic Actuators. We believe it provides a good balance between being simple enough to be intuitive yet descriptive enough to capture the most important actuator dynamics. The rack & pinion representation is an intuitive way to capture both the no-slip condition and the coupling between the rotational and translational motions of a ball screw driven actuator. The results above show that the new model matches the experimental data well, not just in the high impedance test case, but also for the moving output test case.

3-6 Future Work

In this paper, we presented a model for ball screw driven linear SEAs and specifically looked at the case where the spring is located in between the motor housing and Link 1. For future work, we are extending this model to investigate other linear SEA designs. For example, if the spring is placed between the ball nut and Link 2, the model would look somewhat different, as shown in Figure 3-10a. We expect to show that the differences in the actuator dynamics are negligible, but are still verifying our results. We are also extending the model to investigate some common rotary SEA designs. Figure 3-10b shows our current concept model for a geared rotary SEA. Our current intuition on rotary SEAs suggests that the conventional lumped model is only appropriate for a direct drive rotary SEA. We also plan to use the newly developed unlumped models to revisit some of the previous findings in the literature that relied heavily on the conventional lumped model.

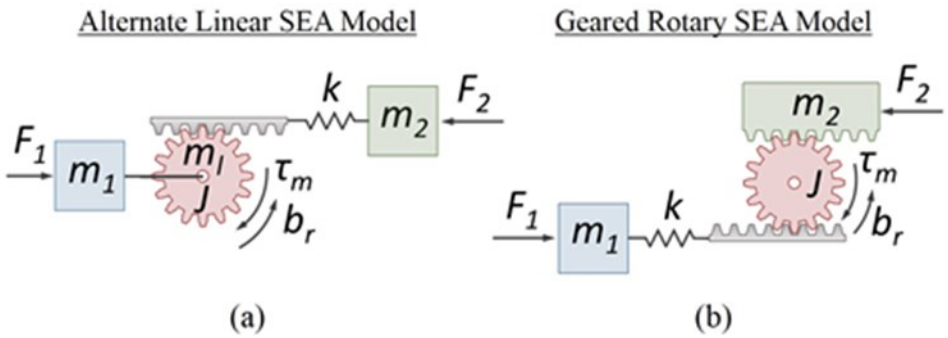

Figure 3-10. Proposed unlumped models for (a) an alternate linear SEA design, and (b) for a geared rotary SEA design.

3-7 Acknowledgment

The authors would like to thank Derek Lahr, Steve Ressler, Bryce Lee, Jordan Neal, and Robert Griffin who contributed valuable insights and feedback for this paper. This material is supported by ONR through grant N00014-11-1-0074 and by DARPA through grant N65236-12-1-1002.

4 Design, Modeling, and Stiffness Selection of Linear Series Elastic Actuators

4-1 Abstract

Series Elastic Actuators (SEAs) are widely used in force-controlled robotic applications where a reliable force source is desired. With careful design, the actuator dynamics can be ignored and the SEA can be treated as a pure force source. However, the designer of high performance Series Elastic Actuators is met with many design choices and tradeoffs. Spring stiffness, spring location, sensor location, motor inertia, motor damping, sprung mass, and supply current are all critical design variables. In this paper we present an approach to evaluating how these critical design variables affect the actuator dynamics and controlled performance of linear SEAs. We first consider the effect of spring location and show that the location of the spring and force sensor does not matter as long as the actuator is used within a certain force bandwidth range. The allowable bandwidth range is shown to depend primarily on the sprung mass within the actuator and a design approach is presented which ensures a desired bandwidth range for the pure force source assumption. We then look at the effect of the design variables on the controlled performance and derive the theoretical maximum force controller bandwidth of an ideal compensator given the motor current saturation limitations. Several important design implications are presented including two stiffness selection criteria which ensure that the actuator can be treated as a pure force source and reliably controlled within a desired force bandwidth range.

4-2 Introduction

Robotic actuators come in many form factors, but the three most common actuator design approaches can be categorized as: rotary, cable, and linear. Rotary actuators are the most widely used and often use an electric motor which transmits torque either directly to a joint or through a transmission which can be a belt reduction, harmonic drive, cycloidal drive, or planetary gearset [25], [50], [51]. Cable Driven Actuators use cables to transmit forces and torques between the actuator and the joint and can be bidirectional [14], or antagonistically actuated [53]. Linear actuators can be hydraulic, pneumatic, or electric, with the common feature being the linear motion of the output [32], [33], [49], [52]. Electric linear actuators often use a rotary motor which acts through a lead screw, ball screw, or roller screw transmission to convert rotary motion into linear motion. Ball screws are the most common choice for high performance robotic applications due to their large gearing ratios, low friction, high efficiency, high backdrivability, low backlash, high precision, and high impact resistance. Roller screws have similar properties and in many cases offer higher performance, but are significantly more expensive than ball screws. In this paper we

specifically address the design of ball screw driven linear actuators, however all of the models and analysis apply equally well to lead screw and roller screw actuators.

The vast majority of robots, and especially electrically powered humanoid robots, use rotary actuators. Electric motors have traditionally been available in a rotary form factor and are relatively straightforward to design into a robotic joint. Linear actuators are less common but offer some significant benefits over rotary actuators. Linear actuators enable a highly biologically-inspired design by enabling a light weight yet rigid internal structure with the actuators packaged externally, much like the skeletal structure of humans. They also enable parallel actuation arrangements in which two or more actuators control two or more degrees of freedom simultaneously. By doing so, multiple actuators can be recruited for high power tasks when necessary. Linear actuators typically act on a lever arm such that the mechanical advantage varies throughout the range of motion. While normally considered an inconvenience, the varying mechanical advantage can be exploited to provide peak force where it is needed within the range of motion. However, the limited stroke of a linear actuator does reduce the range of motion, which can be a significant limitation in robotic applications.

This paper seeks to address the design and modeling of linear Series Elastic Actuators (SEAs) which use screw-type transmissions. SEAs come in a variety of configurations but are generally characterized by a physical spring placed in series with the output. Numerous studies have addressed the various benefits of series elasticity for force control in a variety of robotic applications [25], [46], [47]. Despite the wide variety of actuator designs, SEAs are commonly discussed and depicted using the same basic lumped model shown in Figure 4-1. This same model is used for actuators of varying form factors, transmissions, spring arrangements, and output mechanisms. In [69], we show that the lumped model is inadequate for analyzing linear SEAs with screw-type transmissions and a more descriptive model is needed. In this paper we introduce a more descriptive rack & pinion model and use it to investigate the contributions of important design variables on the actuator dynamics and controlled performance. However, we believe that more descriptive models are also needed for other actuator designs including geared rotary and cable driven actuators.

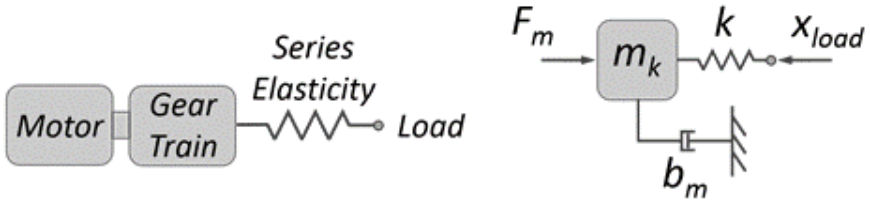

Figure 4-1. Lumped models for SEAs. F_m is the motor force, m_k is the lumped sprung mass, b_m is the lumped damping, and k is the stiffness of the physical spring placed in series.

4-2-1 Depicting Screw-Type Actuators

Figure 4-2a shows a common configuration for ball screw driven actuators. An electric brushless DC motor acts through a belt reduction which in turn rotates a ball screw. As the ball screw is rotated, the ball nut is kept from spinning using an anti-rotation mechanism, resulting in a linear translation of the ball nut, and an overall change in length of the actuator. This same behavior can be visualized and represented with a rack & pinion representation as shown in Figure 4-2b. The

rotation of the pinion gear is converted into the linear translation of the rack through the gearing ratio. The rack & pinion representation is a useful simplification for any screw-based linear actuator because it intuitively depicts this rotary-to-linear transformation.

In the equivalent rack & pinion representation in Figure 4-2b, the pinion gear has a mass m_p and an inertia J corresponding to the combined mass and inertia of the motor, belt reduction, and ball screw. The rack mass m_r corresponds to the combined mass of the ball nut and output mechanism. The radius r of the pinion gear corresponds to the combined gearing ratio of the belt reduction N_p and ball screw lead l

$$r = \frac{l}{2\pi N_p}. \tag{4-1}$$

The motor torque τ_m and motor damping b_m can then be depicted as acting on the equivalent rack & pinion resulting in a rotation θ of the pinion gear and translation x_r of the rack corresponding to the rotation of the motor rotor and change in actuator length, respectively.

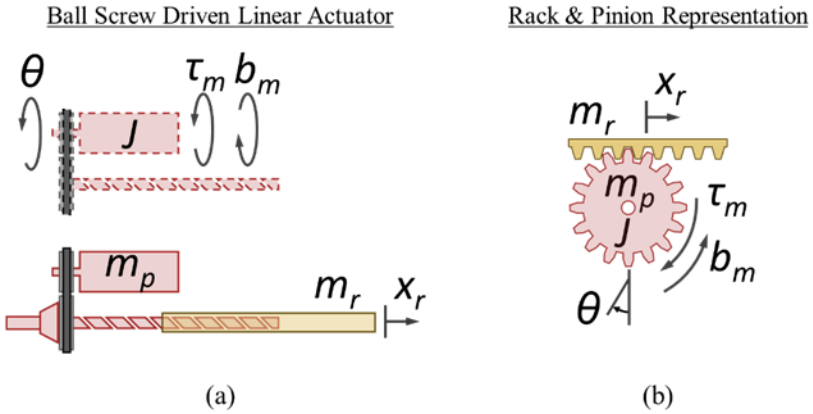

Figure 4-2. Ball screw driven linear actuator showing a belt reduction to ball screw transmission (a) and the equivalent rack & pinion representation (b).

4-2-2 Linear Series Elastic Actuators

The first and most widely used linear SEA design approach is shown in Figure 4-3 [15], [48]. An electric motor directly drives a ball screw transmission with the anti-rotation mechanism implemented via a set of low friction guide rails. A set of preloaded die springs simultaneously serve as the elastic element and force sensor. The force measurement is achieved by measuring the displacement of the springs and calculating the force using a known spring constant.

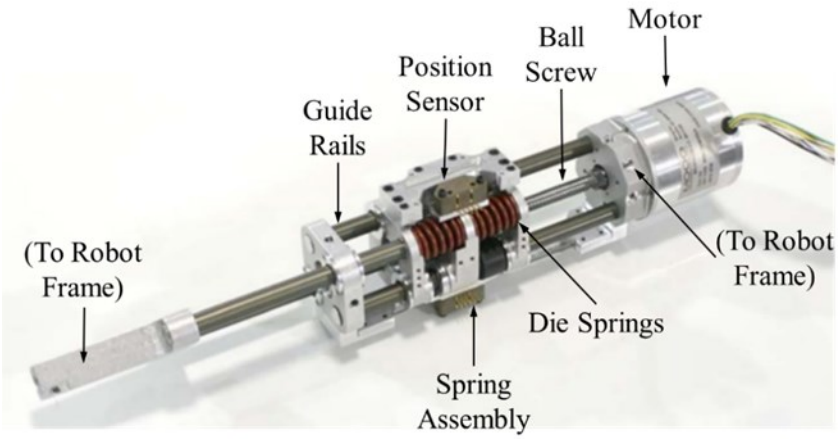

Figure 4-3. Linear Series Elastic Actuator from [15], used with permission of J. Pratt.

An alternative design, shown in Figure 4-4. is used in the lower body of THOR, a 1.78 m 34-DOF torque-controlled humanoid [33]. Similar to other designs, we use a 100 W BLDC motor to drive a ball screw with a 0.002 m lead through a 3:1 belt reduction. This configuration provides a maximum speed of 0.19 m/s, a continuous force of 685 N, and a peak force of 2225 N. Universal joints at either end of the actuator simultaneously serve as the anti-rotation mechanism for the ball nut and as connection points attaching the actuator to the robot structure. One of the novel features of this design is that the spring element is implemented as a cantilevered titanium beam. The actuator applies forces through a lever arm to load the beam with an almost pure moment. The actuator weight 0.938 kg including the compliant beam. While many designs use the spring deflection as a way to estimate force [48], [49], we use a relatively rigid in-line load cell. The spring can be configured in a stiff and a soft compliance setting with stiffness values of 655 kN/m or 372 kN/m respectively. There is also a locked out setting which is accomplished by rigidly bolting the lever arm into a structural member on the robot using two M6 bolts.

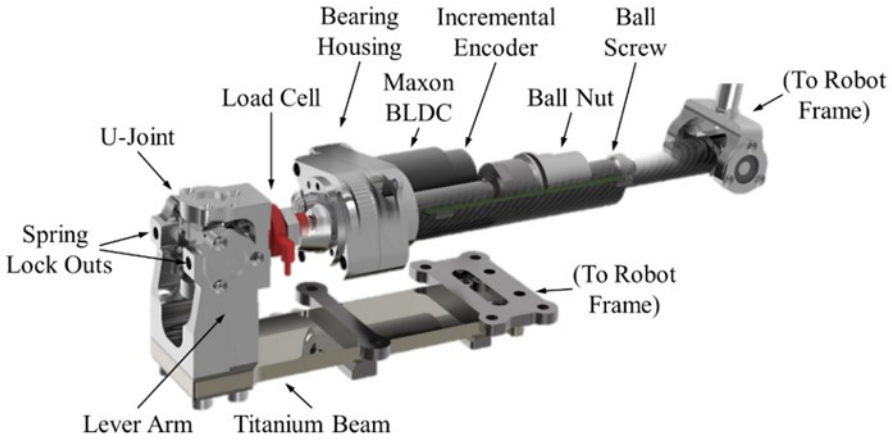

Figure 4-4. Schematic of the THOR Linear SEA used in the lower body of THOR, used with permission of J. Holler.

While these are two design examples, there are numerous other design implementations of linear SEAs which vary in their spring designs, spring locations, force sensing methods, anti-rotation mechanisms, and output mechanisms.

4-2-3 Stiffness Selection

Since their introduction in [25], Series Elastic Actuators have been an exercise in design tradeoffs, especially in selecting the stiffness of the spring element. A softer spring can reduce impact loading, reduce output impedance, and enable high fidelity force control. But these advantages come at the cost of reduced control bandwidth and added complexity and weight. In [48], the authors formulate stiffness selection as a tradeoff between large force bandwidth needing a stiff spring and low impedance performance needing a compliant spring. The design goal in this approach is to develop and tune a very good force controllable actuator which could then be treated as a pure force source. We follow a similar approach in this paper, but introduce a new stiffness criteria which ensures that the actuator can be treated as a pure force source up to a desired bandwidth.

It should be noted, however, that the stiffness selection criteria can depend on the specific design goals that are most relevant for a desired application. Vanderborght et al. have investigated how stiffness affects energy consumption in electrically and pneumatically actuated robots [70], as well as how an active controller can be used in to minimize energy consumption [71]. Grimmer et al. in a three part series of papers have investigated how stiffness affects the energy consumption and peak power for ankle [72], knee [73], and hip Series Elastic Actuators [74]. The results showed that the optimal stiffness depended on which design goal was optimized, energy consumption or peak power. Finally, Tsagarakis et al. find optimized stiffnesses for the COMAN humanoid based on the design goals of maximizing energy storage and avoiding resonances [75].

4-2-4 Modeling Series Elastic Actuators

The conventional approach to modeling SEAs, as shown in Figure 4-1 is to treat one side as being rigidly attached to ground with the other side serving as the output. The actuator dynamics are then derived with the motor torque as the input and the force on the assumed output link as the output. The reaction forces acting on the "ground" link are typically not considered, with [52] being a rare exception. While this approach works well for evaluating an actuator on a test stand, it is not representative of how the actuators are commonly used in robotics applications.

In reality, actuators in legged robots act on two links; each of which are allowed to move subject to joint constraints and have external forces acting on them as shown in Figure 4-5, where m_1 is the linearized inertia of Link 1, m_2 is the linearized inertia of Link 2, F_1 and F_2 are the external forces acting on the links either due to internal dynamics or external disturbances, τ_m is the motor torque, b_m is the damping in the motor and transmission, and k is the spring stiffness.

Depending on the phase of the walking cycle, either side of the actuator can be considered as the most appropriate ground approximation. Since each link varyingly serves as both the output link and ground link, it is important to consider the forces acting on both sides of the actuator, not just the assumed output. Considering SEAs are most often used as force (or torque) controlled actuators, the actual forces being exerted on the links is a critical parameter for state estimation and control.

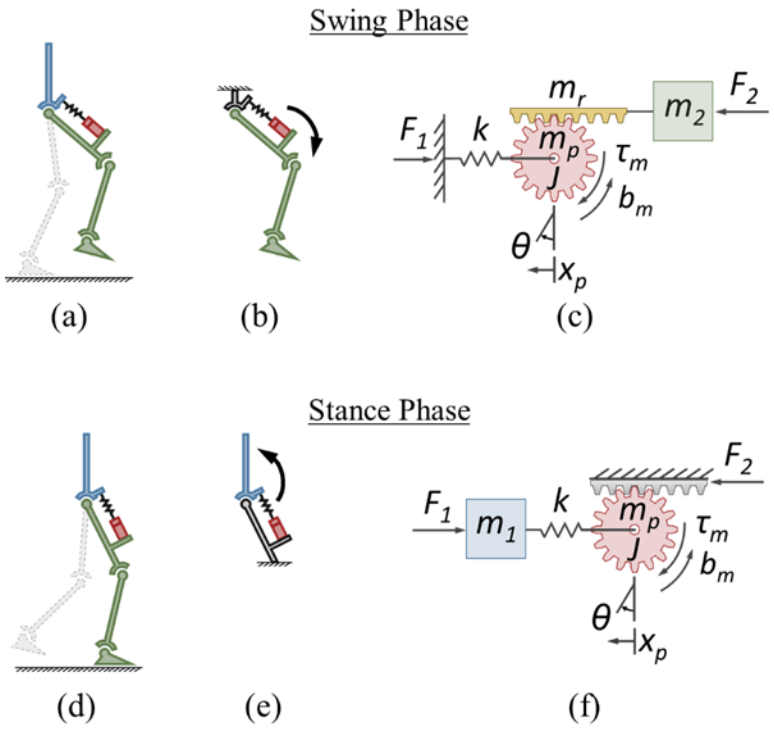

Figure 4-5. The effect of changing ground contacts in the swing phase (a, b, and c), and stance phase (d, e, and f).

4-2-5 Joint Torque & Intermediate Inertias

Series Elastic Actuators are often used in force-control applications where it is common practice to treat them as pure force sources. The assumption is that the SEA serves as a reliable force source and that the internal actuator dynamics can be ignored. For a linear actuator acting across two links, the joint torque can be approximated by measuring the actuator force and knowing the lever arm at the output. This approach assumes that the forces acting at either end of the actuator are identical. However, if the actuator has an internal mass that is accelerating, the intermediate inertia within the actuator can introduce significant additional dynamics, leading to different forces acting on the two links. If the forces acting on the two links are different, then the actuator can no longer be treated as a pure force source, the actuator dynamics can no longer be neglected, and the calculation of joint torque becomes much more complicated.

In [52], the authors argue that the forces acting on either side of an SEA are different due to the internal actuator dynamics. However, they use a lumped model of the actuator, which depicts a very large intermediate inertia, the acceleration of which has dramatic effects on the actuator dynamics. As we show in [69], the actual intermediate inertia is much smaller and is more accurately modeled by the unlumped rack & pinion model which does not lump the translational and rotational inertias. Nevertheless, it is important to consider the effect of internal dynamics to ensure that the pure force source assumption is appropriate for a given application.

4-2-6 Paper Outline

The designer of high performance Series Elastic Actuators is met with many design choices and tradeoffs. Spring stiffness, spring location, sensor location, motor inertia, motor damping, sprung mass, and supply current are all critical design variables. In this paper we present an approach to evaluating how these critical design variables affect the actuator dynamics and controlled performance of linear SEAs. We first consider the effect of spring location by developing models, deriving the equations of motion, and simulating the two most common spring configurations. The models and simulations are used to consider both the high impedance as well as the moving outputs test cases. We use the results to present a design approach which ensures that the actuator can be treated as a pure force source within a desired bandwidth range. We then introduce a simple model for an ideal compensator with motor saturation to investigate the effect of various design variables on the controlled performance. Finally, several important design implications are presented including two stiffness selection criteria.

4-3 Spring Location in Linear SEAs

The spring element in an SEA is of primary importance because it provides the physical compliance, and in many cases, serves as the force sensor. Because of this, the literature places a strong emphasis on the location of the spring element. The early literature on SEAs describes placing the spring element between the motor and the "output" while some more recent designs describe placing the spring element in between the motor and "ground" [52], [65]. Both [52] and [65] refer to their designs as novel implementations specifically because of the spring placement, referring to the new arrangements as Reaction Force Sensing SEA and Force Sensing Compliant Actuator, respectively. Similarly, in [63], the authors use the terms distal compliance and proximal compliance to describe what they consider to be two different SEA implementations based entirely on the location of the spring relative to the output. However, as we show in Section 4-2-4, there is no single "ground" side or "output" side.

There are differences between the various designs, but the differences have to do with how the spring is implemented into the actuator design itself, not which side is being treated as the ground or output. As shown in Figure 4-6, there are two main design approaches for linear actuators based on the location of the spring relative to the motor housing and ball nut. The more popular approach is to place the spring in between one of the links and the ball nut, as shown in Figure 4-6 (a, b) [15]. The other approach is to place the spring in between one of the links and the motor housing, shown in Figure 4-6 (c, d) [26], [52]. We refer to these as sprung motor housing (SMH) and sprung ball nut (SBN) designs. Note that the difference between the two designs is depicted in the rack & pinion models as a sprung pinion gear verses a sprung rack.

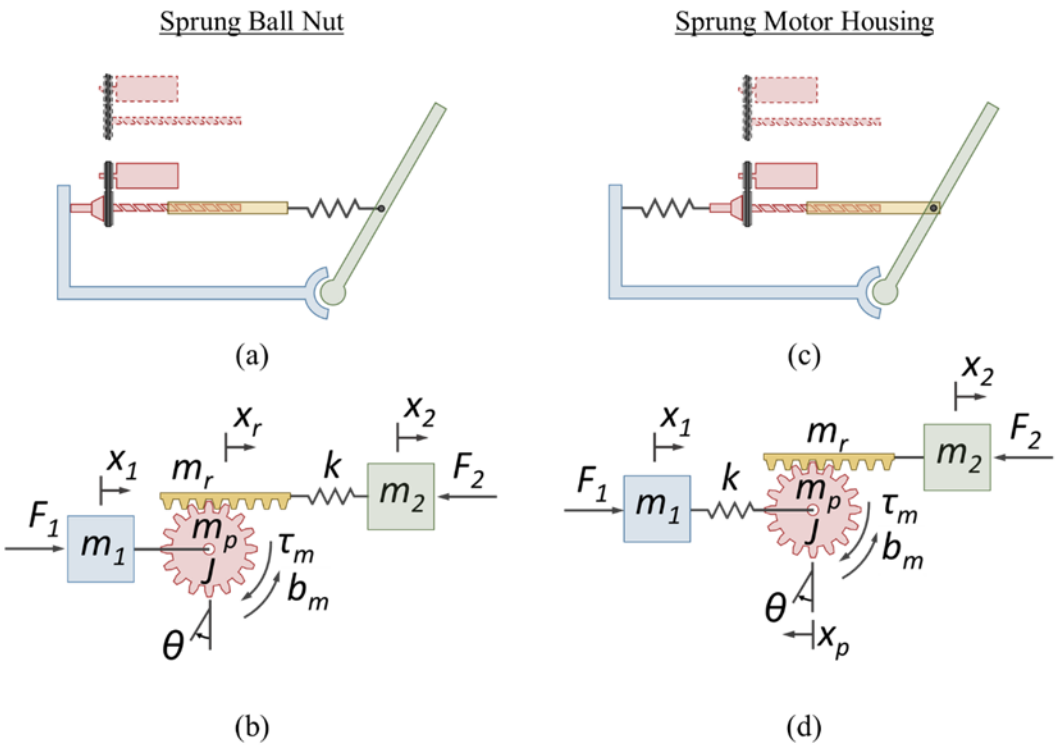

Figure 4-6. Schematic and models of the two most common spring locations for linear Series Elastic Actuators, the Sprung Ball Nut (a, b) and Sprung Motor Housing (c, d).

While both spring locations are viable options, there are some specific advantages to the sprung motor housing design. In the sprung ball nut design, the springs are located within the stroke length of the actuator, reducing how far the ball nut can travel before interfering with the rest of the actuator. In [26], [33], and [52], the authors show how a sprung motor housing leads to a significantly more compact design by moving the spring element away from the main body of the actuator resulting in a shorter overall length for a given stroke. The ratio of actuator length to stroke length is a critical metric for linear actuators as it leads to much more efficient packaging.

There are also instrumentation advantages to the sprung motor housing design. The power and control electronics are usually mounted either on the motor housing or on the link that moves the least relative to the motor housing. A sprung motor housing is therefore easier to instrument because any position or force sensors in the spring element will not move far relative to the power and control electronics. In the sprung ball nut design, the spring moves with the output mechanism along the entire stroke length. Given these differences, we need to address how the spring location affects the actuator dynamics, since it could have significant implications for the mechanical design, instrumentation selection, and control of linear SEAs.

4-4 Model Derivations

In deriving the two models, the forces of interest are motor force F_m, spring force F_k, the internal force in the ball screw F_{bs}, and the forces acting on the two links F_{sprung} and $F_{unsprung}$. We consider two test cases; the widely used high-impedance test case which approximates the stance

phase, balancing, and interacting with the environment; and the moving output test case which approximates the swing phase. We begin with the general equations of motion, derive the high-impedance dynamics analytically, and then simulate the moving output test case.

4-4-1 Equations of Motion | Sprung Ball Nut

Figure 4-7 shows the free body diagrams for the sprung ball nut linear SEA. Note that the translational mass of the pinion gear will always move with the mass of Link 1, so they can be treated as a lumped mass.

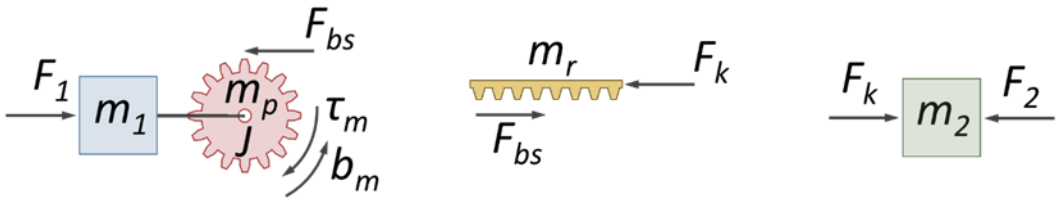

Figure 4-7. Free body diagrams for a sprung ball nut linear SEA actuator model.

The resulting equations of motion for the SBN model are given by

$$\sum F_{x\_m_1}: \quad (m_1 + m_p)\ddot{x}_1 = F_1 - F_{bs}\ , \tag{4-2}$$

$$\sum F_{x\_m_r}: \quad m_r\ddot{x}_r = F_{bs} - F_k\ , \tag{4-3}$$

$$\sum F_{x\_m_2}: \quad m_2\ddot{x}_2 = F_k - F_2\ , \tag{4-4}$$

$$\sum M_{\theta\_m_p}: \quad J\ddot{\theta} = \tau_m - b_m\dot{\theta} - rF_{bs}\ , \tag{4-5}$$

$$motor: \quad \tau_m = rF_m\ , \tag{4-6}$$

$$no\ slip: \quad x_r - x_1 = r\theta, \quad \dot{x}_r - \dot{x}_1 = r\dot{\theta}, \quad \ddot{x}_r - \ddot{x}_1 = r\ddot{\theta}\ . \tag{4-7}$$

4-4-2 Equations of Motion | Sprung Motor Housing

Similarly, Figure 4-8 shows the free body diagrams for the sprung motor housing linear SEA. Note that the translational mass of the rack will always move with the mass of Link 2, so they can be treated as a lumped mass.

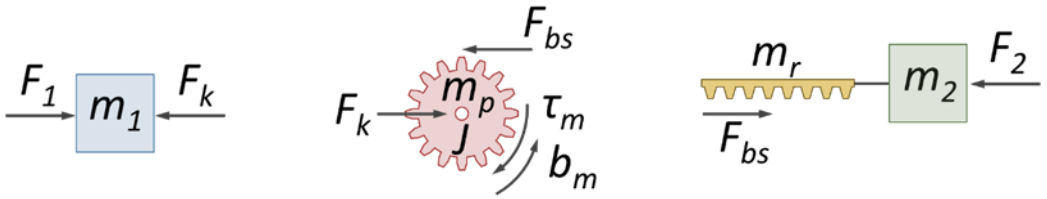

Figure 4-8. Free body diagrams for a sprung motor housing linear SEA actuator model.

The resulting equations of motion for the SMH model are given by

$$\sum F_{x\_m_1}: \quad m_1\ddot{x}_1 = F_1 - F_k \,, \tag{4-8}$$

$$\sum F_{x\_m_p}: \quad m_p\ddot{x}_p = F_{bs} - F_k \,, \tag{4-9}$$

$$\sum F_{x\_m_2}: \quad (m_r + m_2)\ddot{x}_2 = F_{bs} - F_2 \,, \tag{4-10}$$

$$\sum M_{\theta\_m_p}: \quad J\ddot{\theta} = \tau_m - b_m\dot{\theta} - rF_{bs} \,, \tag{4-11}$$

$$motor: \quad \tau_m = rF_m \,, \tag{4-12}$$

$$no\ slip: \quad x_p + x_2 = r\theta, \quad \dot{x_p} + \dot{x_2} = r\dot{\theta}, \quad \ddot{x_p} + \ddot{x_2} = r\ddot{\theta} \,. \tag{4-13}$$

4-4-3 Simulink Models

The above equations of motion were used to develop simulation models for both of the spring configurations shown in Figure 4-6b and Figure 4-6d. The models were developed in Simulink by MathWorks and enable the user to fix either of the links or to adjust any of the design variables to investigate their relative effects. But more importantly, the Simulink models can simulate experiments that would otherwise be difficult to perform in real hardware due to the limited stroke of the actuator, sensor noise issues, motor thermal limits, and communication bandwidth. As an example, the CAN communication used on the THOR-Linear actuator test stand is limited to 400 Hz, which limits the system identification bandwidth to 200 Hz. Table 6 lists the known and extracted model variables used in the Simulink models. The known variables were either taken from component datasheets or directly measured. The remaining unknown variables were extracted using the experimental system identification procedure described in [69].

Table 6: Known and Extracted model Variables

Known Variables			
m_l	$0.6461\ kg$	K_τ	$0.0255\ Nm/A$
k	$655000\ N/m$	r	$1.061x10^{-4}\ m$
l	$0.002\ m$	N_p	$3:1$
Extracted Variables			
m_k	$294.9\ kg$	J	$3.312x10^{-6}\ kgm^2$
b_m	$11658\ Ns/m$	b_r	$1.3123x10^{-4}\ Nms$

4-4-4 High Impedance Test Case

A common approach for evaluating SEA dynamics is to consider the high impedance test case, shown in Figure 4-9, where both output links are rigidly fixed. This test case is intended to approximate both sides of the actuator acting on ground or attached to a link with a large inertia. An example of this scenario is the ankle actuator in the single stance phase where one side (foot) is acting against ground while the other side (shin) is accelerating the entire weight of the robot.

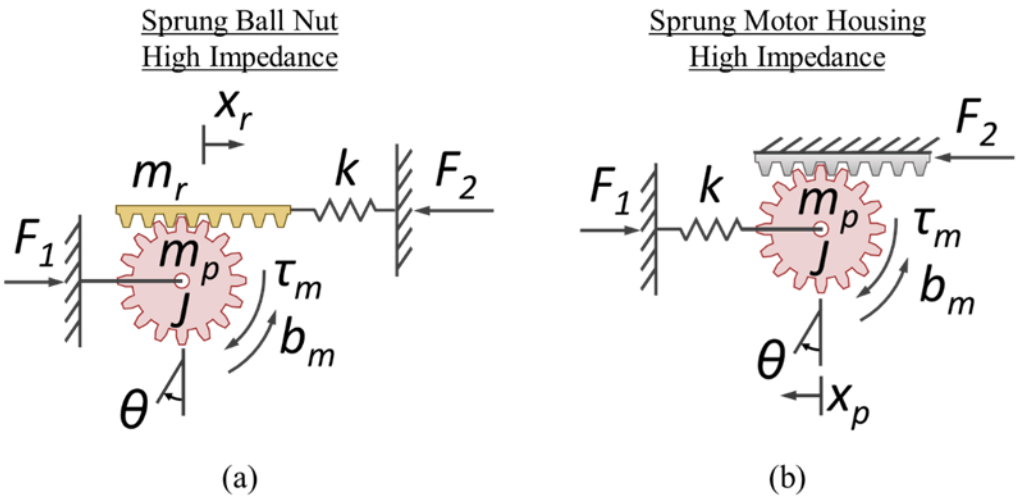

Figure 4-9. High impedance test case for the two spring configurations.

Note that the two models show a fundamental difference in the designs beyond just the location of the spring. In the SMH model, the pinion mass m_p (combined mass of the motor, motor housing, and ball screw) moves with the spring deflection while the rack mass m_r (mass of the ball nut and output tube) is lumped in with Link 2 and treated as fixed to ground. However, in the SBN model, the rack mass moves with the spring deflection and the pinion mass is lumped with Link 1. Another way of articulating this observation is that the two models depict a different sprung mass; the SMH model's sprung mass is m_p while the SBN model's sprung mass is m_r. Depending on the design of the actuator, these masses can potentially be very different. This is an interesting and important observation because it relates how the two designs physically behave differently due to the spring placement.

In this test case we are primarily interested in comparing how the two models relate the input motor force (F_m) to the two forces acting on either end of the actuator (F_1 and F_2). We are also interested in relating how the forces acting on the two links compare to one another, resulting in three responses to consider: F_1/F_m, F_2/F_m, and F_1/F_2. However, instead of looking at F_1/F_2, a similar but more general formulation would be $F_{sprung}/F_{unsprung}$ where F_{sprung} is the force acting on the link closest to the spring and $F_{unsprung}$ is the force acting on the link furthest from the spring.

For the high impedance test case, where $\ddot{x}_1 = 0$ and $\ddot{x}_2 = 0$, Equations 4-2 - 4-7 for the SBN model can be further simplified to

$$F_1 = F_{bs} \,, \tag{4-14}$$

$$m_r \ddot{x}_r = F_{bs} - F_k = F_{bs} - kx_r \,, \tag{4-15}$$

$$F_2 = F_k = kx_r \,, \tag{4-16}$$

$$J\ddot{\theta} = \tau_m - b_m\dot{\theta} - rF_{bs} \,, \tag{4-17}$$

$$motor: \quad \tau_m = rF_m \,, \tag{4-18}$$

$$no\ slip: \quad x_r = r\theta, \quad \dot{x}_r = r\dot{\theta}, \quad \ddot{x}_r = r\ddot{\theta}\ . \tag{4-19}$$

From which we can derive the following transfer functions

$$\frac{F_1(s)}{F_m(s)} = \frac{F_{bs}(s)}{F_m(s)} = \frac{m_r s^2 + k}{\left(\frac{J}{r^2} + m_r\right)s^2 + \left(\frac{b_m}{r^2}\right)s + k}\ , \tag{4-20}$$

$$\frac{F_2(s)}{F_m(s)} = \frac{F_k(s)}{F_m(s)} = \frac{k}{\left(\frac{J}{r^2} + m_r\right)s^2 + \left(\frac{b_m}{r^2}\right)s + k}\ , \tag{4-21}$$

$$\frac{F_{sprung}(s)}{F_{unsprung}(s)} = \frac{F_k(s)}{F_{bs}(s)} = \frac{F_2(s)}{F_1(s)} = \frac{k}{m_r s^2 + k}\ . \tag{4-22}$$

Similarly, Equations 4-8 - 4-13 for the SMH model can be simplified to

$$F_1 = F_k = kx_p\ , \tag{4-23}$$

$$m_p\ddot{x}_p = F_{bs} - F_k = F_{bs} - kx_p\ , \tag{4-24}$$

$$F_2 = F_{bs}\ , \tag{4-25}$$

$$J\ddot{\theta} = \tau_m - b_m\dot{\theta} - rF_{bs}\ , \tag{4-26}$$

$$motor: \quad \tau_m = rF_m\ , \tag{4-27}$$

$$no\ slip: \quad x_p = r\theta, \quad \dot{x}_p = r\dot{\theta}, \quad \ddot{x}_p = r\ddot{\theta}\ . \tag{4-28}$$

From which we can derive the following transfer functions:

$$\frac{F_1(s)}{F_m(s)} = \frac{F_k(s)}{F_m(s)} = \frac{k}{\left(\frac{J}{r^2} + m_p\right)s^2 + \left(\frac{b_m}{r^2}\right)s + k}\ , \tag{4-29}$$

$$\frac{F_2(s)}{F_m(s)} = \frac{F_{bs}(s)}{F_m(s)} = \frac{m_p s^2 + k}{\left(\frac{J}{r^2} + m_p\right)s^2 + \left(\frac{b_m}{r^2}\right)s + k}\ , \tag{4-30}$$

$$\frac{F_{sprung}(s)}{F_{unsprung}(s)} = \frac{F_k(s)}{F_{bs}(s)} = \frac{F_1(s)}{F_2(s)} = \frac{k}{m_p s^2 + k}\ . \tag{4-31}$$

4-4-5 Discussion | Actuator Models

Table 7 summarizes the key transfer functions for both models. In comparing the two spring locations, it is clear that there are some differences but that the forms of the transfer functions are very similar, with 4-20 matching the form of 4-30, 4-21 matching 4-29, and 4-22 matching 4-31. These pairs have to do with the fact that 4-20 and 4-30 relate the motor force to the force on the unsprung side of the actuator, while 4-21 and 4-29 relate the motor force to the force on the sprung side of the actuator. The only difference in the paired transfer functions is the occurrence of m_r verses m_p. These results demonstrate that the essential actuator dynamics of the two designs are intrinsically different, but have similar forms. The differences can be entirely attributed to the

difference in the sprung mass of the two designs; the SMH model's sprung mass is the pinion while the SBN model's sprung mass is the rack.

Table 7: Actuator Model Transfer Functions

Sprung Ball Nut		Sprung Motor Housing	
$\frac{F_1(s)}{F_m(s)} = \frac{m_r s^2+k}{\left(\frac{J}{r^2}+m_r\right)s^2+\left(\frac{bm}{r^2}\right)s+k}$	(4-20)	$\frac{F_1(s)}{F_m(s)} = \frac{k}{\left(\frac{J}{r^2}+m_p\right)s^2+\left(\frac{bm}{r^2}\right)s+k}$	(4-29)
$\frac{F_2(s)}{F_m(s)} = \frac{k}{\left(\frac{J}{r^2}+m_r\right)s^2+\left(\frac{bm}{r^2}\right)s+k}$	(4-21)	$\frac{F_2(s)}{F_m(s)} = \frac{m_p s^2+k}{\left(\frac{J}{r^2}+m_p\right)s^2+\left(\frac{bm}{r^2}\right)s+k}$	(4-30)
$\frac{F_{sprung}(s)}{F_{unsprung}(s)} = \frac{k}{m_r s^2+k}$	(4-22)	$\frac{F_{sprung}(s)}{F_{unsprung}(s)} = \frac{k}{m_p s^2+k}$	(4-31)

4-5 Actuator Dynamics

To further investigate the effect of spring location, the two models were simulated in Simulink, first in the high impedance test case (Section 4-5-1) and then with moving outputs (Section 4-5-2). The three key transfer functions F_1/F_m, F_2/F_m, and $F_{sprung}/F_{unsprung}$ were considered as they relate the motor current input to the forces being exerted on either link.

4-5-1 High Impedance Test Case

Figure 4-10 shows the simulated frequency response functions for the high impedance test case for both the SBN and the SMH spring locations. Both Link 1 and Link 2 were fixed and the values in Table 6 were used for the model variables. In order to isolate only the effect of spring location, both m_r and m_p were set equal to m_k, the sprung mass.

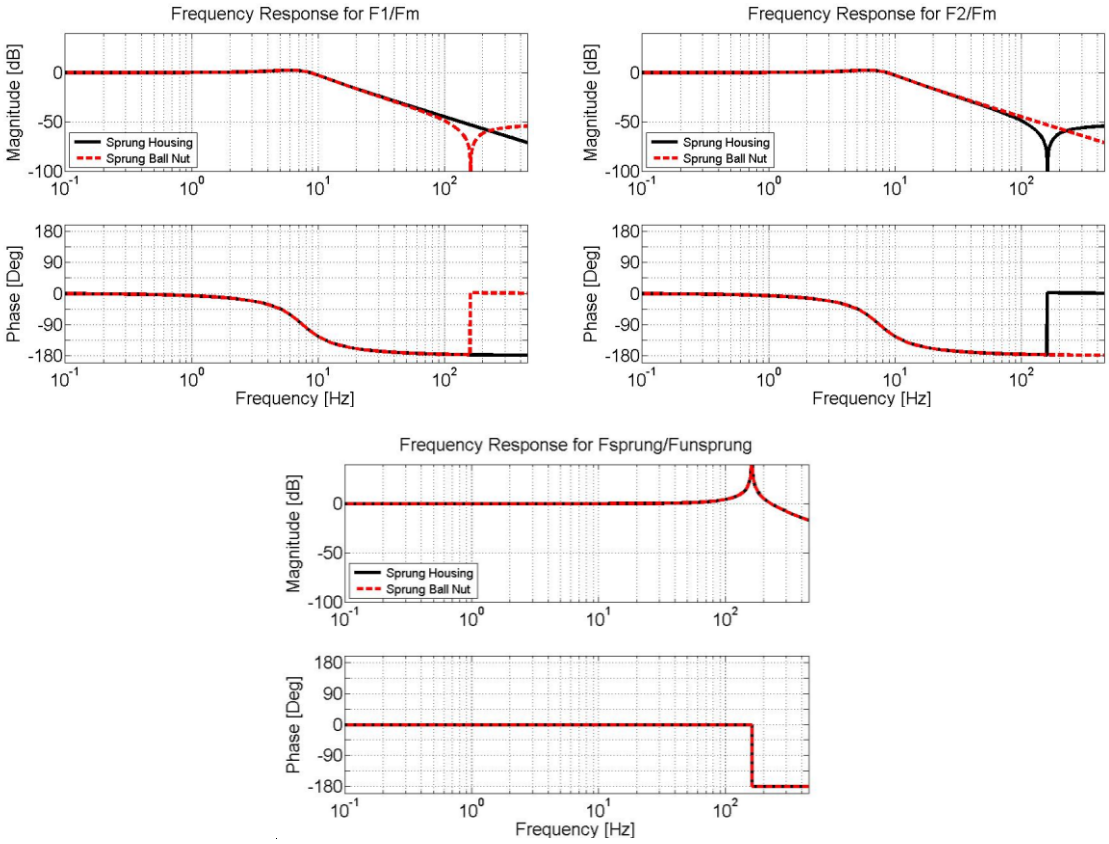

Figure 4-10. Open loop frequency response for the high impedance test case from Simulink simulations for F_1/F_m, F_2/F_m, and $F_{sprung}/F_{unsprung}$.

The simulation plots confirm the observations made from the analytical models in Section 4-4-5. For the same m_k, the only effect of spring location is that of reversing the F_1/F_m and F_2/F_m responses and the $F_{sprung}/F_{unsprung}$ responses are identical. The F_1/F_m and F_2/F_m responses show that the unsprung side of the actuator experiences an anti-resonance at the natural frequency of the mass-spring system given by

$$\omega_n = \frac{1}{2\pi}\sqrt{\frac{k}{m_k}}, \tag{4-32}$$

which is also the frequency of the resonance shown in the $F_{sprung}/F_{unsprung}$ responses.

The $F_{sprung}/F_{unsprung}$ transfer functions are important because they relate the forces being applied on either end of the actuator. In order to treat the actuator as a pure force source, the forces on either end should be close enough to be treated as equal over the desired operational frequency range of the actuator. The $F_{sprung}/F_{unsprung}$ responses in Figure 4-10 show that this condition can only be met over a limited frequency range due to the internal actuator dynamics. The pure force source assumption is valid only over the range of frequencies where the $F_{sprung}/F_{unsprung}$ response is flat.

We here define ω_{max} as the frequency at which the $F_{sprung}/F_{unsprung}$ response first crosses 3 dB, which we define as the upper limit of the acceptable frequency range for the pure force source assumption. A magnitude of 3 dB in the frequency response is a common standard reference because it corresponds to an output amplitude which is 1.413 times larger than the input amplitude, or in terms of power, an output power which is 2 times larger than the input power.

The magnitude of the $F_{sprung}/F_{unsprung}$ response is given by:

$$\left|\frac{F_{sprung}(j\omega)}{F_{unsprung}(j\omega)}\right|_{db} = 20\log_{10}\left(\left|\frac{k}{-m_k\omega^2+k}\right|\right), \tag{4-33}$$

which can then be equated to 3 dB, simplified, and solved for ω_{max} (in rad/s)

$$3\ dB = 20\log_{10}\left(\left|\frac{k}{-m_k\omega_{max}^2+k}\right|\right), \tag{4-34}$$

$$1.413 = \left|\frac{k}{-m_k\omega_{max}^2+k}\right|, \tag{4-35}$$

$$\omega_{max} = \sqrt{\frac{k}{m_k}(\pm 0.708 + 1)}, \tag{4-36}$$

where the first 3 dB magnitude crossing (in Hz) is given by

$$\omega_{max} = \frac{1}{2\pi}\sqrt{\frac{0.292k}{m_k}}. \tag{4-37}$$

Equation 4-37 relates how the mechanical design of the actuator affects the validity of the pure force assumption. At frequencies below ω_{max} the pure force assumption is valid but at frequencies above ω_{max} the internal dynamics of the actuator result in drastically different forces acting on either end of the actuator and the pure force assumption is not valid. Figure 4-11 shows a contour plot of the relationship between m_k, k, and ω_{max}.

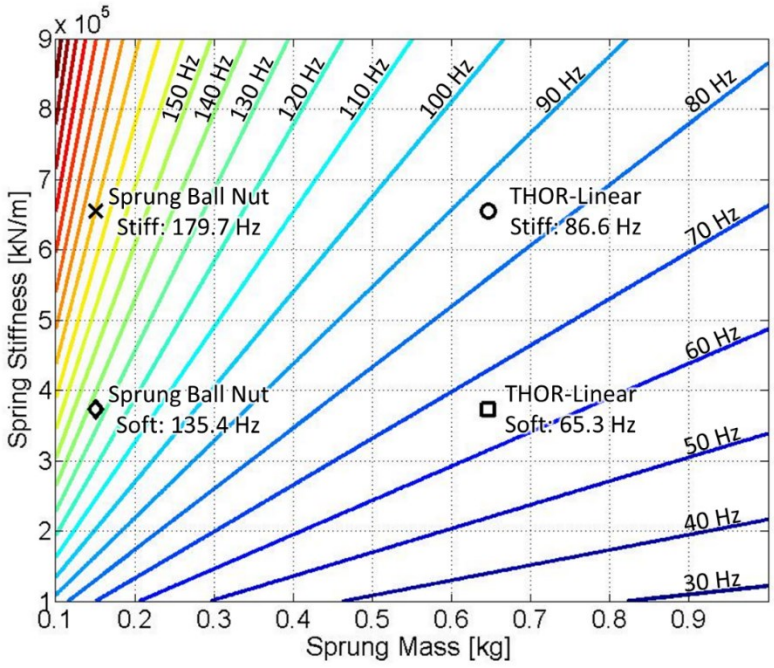

Figure 4-11. Contour plot showing the relationship between spring stiffness, actuator sprung mass, and the maximum bandwidth for which the pure force source assumption is valid.

It is important to note here that Equation 4-37 and Figure 4-11 are related to the actuator dynamics and do not depend on the controller. The ω_{max} values are based solely on the mechanical design of the actuator, highlighting the importance of good design for high performance linear SEAs. Keeping the sprung mass as small as possible and the stiffness as large as possible results in the broadest allowable frequency range for the pure force source assumption. This is where the sprung motor housing design is at a disadvantage. In the SMH design, the sprung mass is m_p, which is the combined mass of the motor, motor housing, and ball screw. In the SBN design, the sprung mass is m_r, which is the much smaller combined mass of the ball nut and output tube. Because of this, keeping the overall actuator mass small is more important for a SMH design since a larger percentage of the actuator mass becomes the sprung mass.

The circle and square markers in Figure 4-11 show the THOR-Linear actuator's values at the stiff and soft setting respectively. The cross and diamond markers show what the THOR-Linear actuator's values would be if we used a sprung ball nut design, which we estimate could be achieved with a sprung mass of 0.15 kg (combined mass of ball nut, spring carrier, and output tube). As the markers show, there is a clear advantage to the SBN design, namely, a much higher ω_{max} for a given spring stiffness due to the much lower sprung mass. Nevertheless, Figure 4-11 shows that the ω_{max} for the THOR-Linear design in both its stiff and soft settings is within our typical force bandwidth of around 60 Hz. We can therefore safely use the SMH design to take advantage of the benefits described in Section 4-3. More importantly, Equation 4-37 and Figure 4-11 can be used by any future designers to ensure that their design meets the pure force source assumption up to a desired operational force bandwidth.

4-5-2 Moving Output Test Case

The Simulink models were also used to investigate the effect of moving outputs on the pure force source assumption. The moving output models, shown in Figure 4-6b and Figure 4-6d, were simulated with either one or both of the outputs moving. Figure 4-12 shows the simulated open loop frequency response functions for the moving output test case. The values in Table 6 were again used for the model variables with both m_r and m_p being set equal to m_k.

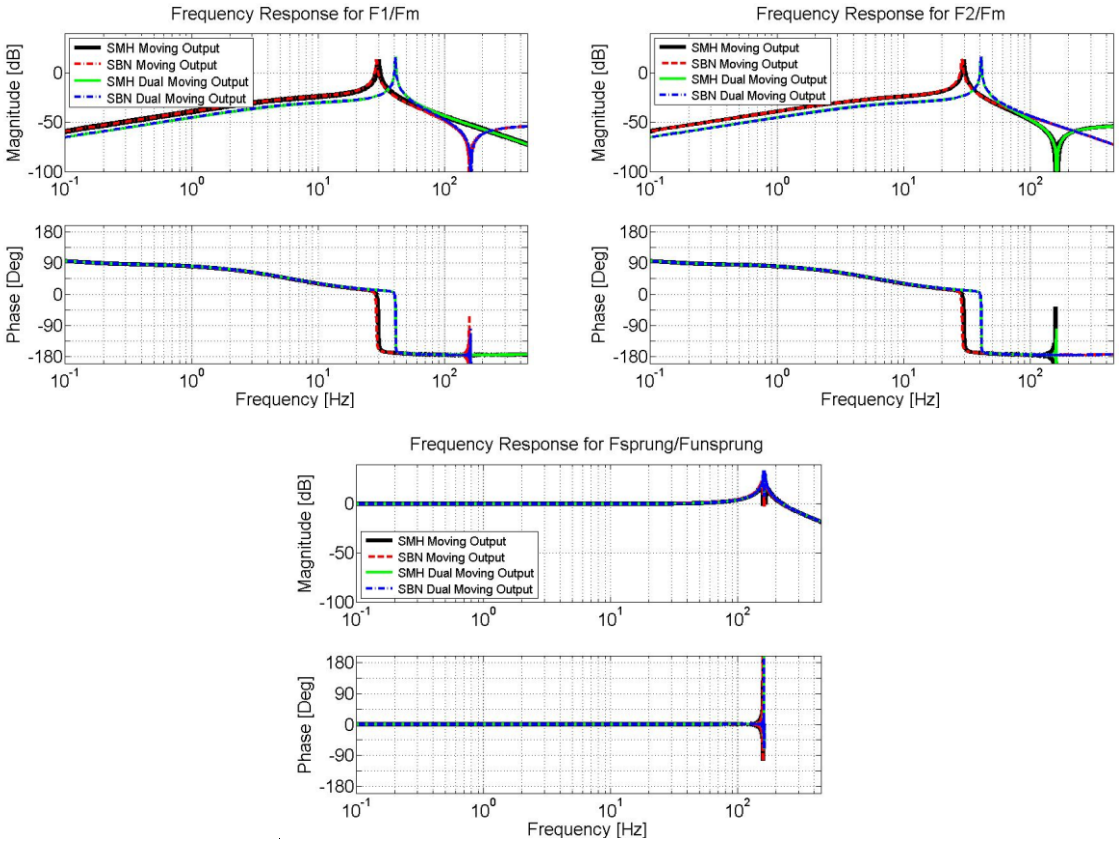

Figure 4-12. Open loop frequency response of the moving output test case from Simulink simulations for F_1/F_m, F_2/F_m, and $F_{sprung}/F_{unsprung}$.

A number of observations could be made with regard to the plots in Figure 4-12, however, the most important observation is that a change in output impedances does affect the F_1/F_m and F_2/F_m responses, but does not affect the $F_{sprung}/F_{unsprung}$ responses. In fact, the $F_{sprung}/F_{unsprung}$ responses perfectly match the high impedance test case responses of Figure 4-10.

The importance of this finding is that even for changing output impedances, the relationship of Equation 4-37 still holds, and the pure force source assumption is still valid up to a bandwidth of ω_{max}. As long as the actuator is controlled within this operational force bandwidth, the internal actuator dynamics are negligible and the actuator can be treated as a pure force source, simplifying the computational burden on higher level joint controllers and whole body controllers. There are also implications for the location of the force sensor. As long as the pure force source criteria is met, the force sensor (such as the load cell in the THOR-Linear design) can be placed anywhere within the line of action of the actuator, not just at the spring interface.

4-5-3 Stiffness Criteria | Pure Force Source Approximation

The above results lead to a new minimum stiffness criteria based on the actuator dynamics of the two linear SEA designs. Solving Equation 4-37 for k yields

$$k_{min} = 134.84\, m_k\, {\omega_{max}}^2\,, \tag{4-38}$$

where ω_{max} is the upper bound of the desired force bandwidth of the pure force source assumption. Since m_k is different for the two designs, we can define

$$k_{min,SBN} = 134.84\, m_r\, {\omega_{max}}^2\,, \tag{4-39}$$

$$k_{min,SMH} = 134.84\, m_p\, {\omega_{max}}^2\,. \tag{4-40}$$

While Equations 4-39 and 4-40 are very similar, the values of m_r and m_p can be significantly different depending on the design. For the THOR-Linear SEA, m_p is 0.6461 kg while m_r would be approximately 0.15 kg for a notional SBN design. Summarizing the results of this section, we can conclude that the THOR-Linear actuator can be treated as a pure force source up to a force bandwidth of 86.6 Hz in the stiff spring configuration and 65.3 Hz in the soft spring configuration which is within our typical 30-60 Hz force bandwidth range.

4-6 Controlled Performance

Sections 4-4 and 4-5 investigated the effect of spring location, spring stiffness, and sprung mass on the actuator dynamics. In this section, we consider how SEA design affects the controlled performance. A theoretical ideal compensator is introduced and used in the analysis so that the results are generalized for any controller and provide an upper bound for the maximum performance that can reasonably be expected. The performance metric used in this analysis is force bandwidth which is widely considered to be a key performance metric in force controlled applications.

4-6-1 Ideal Compensator

For a given plant model, an ideal compensator can be defined as the inverse of the plant model. Such a compensator would perfectly compensate for the plant dynamics across all frequencies but is impractical due to the detrimental effects of modeling errors, noise, and nonlinearities which would all lead to control and stability issues. There is also the practical consideration of motor saturation due to power supply limits. The inverse plant has to compensate for the low pass filtering effect of the spring element by counteracting the -40 dB/decade roll off above the cutoff frequency. There is a practical limit, however, to the achievable compensator gain due to power limitations.

Figure 4-13 shows the diagram of an ideal inverse plant compensator with current saturation where F_d is the desired force and F_m is the force of the actuator output. The limitation of the power supply is implemented with a simple current saturation model where the command signal is clipped above the max allowable current.

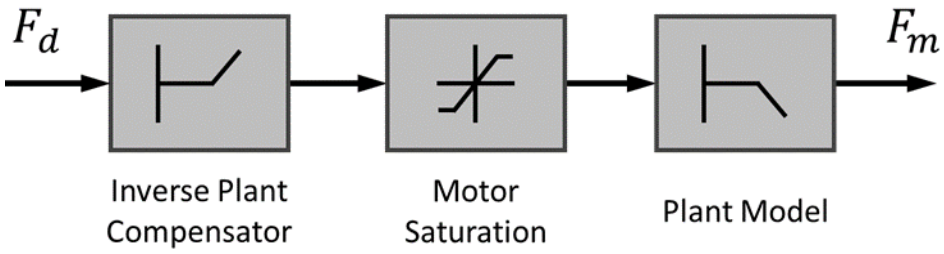

Figure 4-13. Control diagram of an ideal inverse plant compensator with motor saturation.

Figure 4-14 shows the effect of motor saturation on the controlled performance of an SEA.

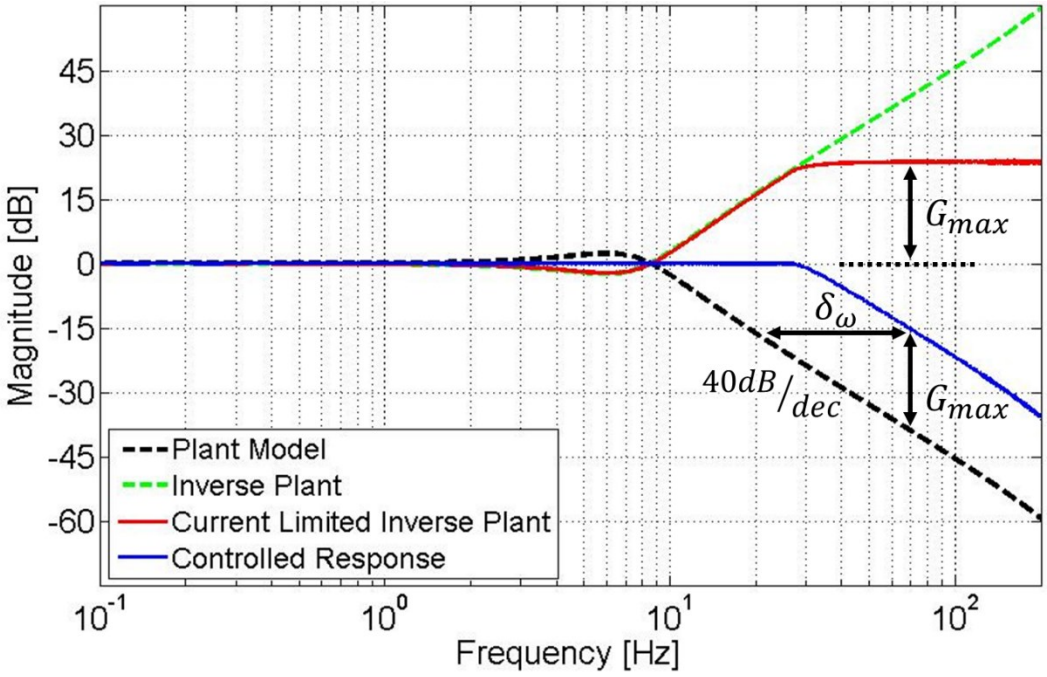

Figure 4-14. Controlled performance of an ideal compensator with motor saturation. The input force amplitude is 200 N and the current limit is 10 A.

The 2nd order plant model, shown in black, is given by

$$P = \frac{F_m(s)}{F_d(s)} = \frac{k}{\left(\frac{J}{r^2}+m_k\right)s^2+\left(\frac{b_m}{r^2}\right)s+k}, \tag{4-41}$$

the ideal inverse plant compensator, shown in green, is given by

$$P^{-1} = \frac{\left(\frac{J}{r^2}+m_k\right)s^2+\left(\frac{b_m}{r^2}\right)s+k}{k}, \tag{4-42}$$

and the motor saturation model is defined by

$$F_{sat} = \begin{cases} F_{max} & , F > F_{max} \\ F & , F_{max} > F > -F_{max} \\ -F_{max} & , F < -F_{max} \end{cases}. \tag{4-43}$$

Without saturation, the controlled performance of the ideal compensator would result in a flat response at all frequencies. With saturation, the inverse plant compensator, shown in red, levels off at higher frequencies once it reaches the maximum gain G_{max}. The value of G_{max} can be found by considering the effect of saturation on the control signal, shown in Figure 4-15. Without saturation, the root mean square (RMS) value of a sinusoidal control signal with amplitude A is

$$A_{RMS} = \frac{A}{\sqrt{2}}, \tag{4-44}$$

and the controller gain in decibels is given by

$$G_{dB} = 20\log_{10}\left(\frac{A_{output}/\sqrt{2}}{A_{input}/\sqrt{2}}\right), \tag{4-45}$$

which for an SEA can be expressed either in terms of force or current amplitude

$$G_{dB} = 20\log_{10}\left(\frac{A_{F_{output}}/\sqrt{2}}{A_{F_d}/\sqrt{2}}\right), \tag{4-46}$$

$$G_{dB} = 20\log_{10}\left(\frac{A_{I_{output}}/\sqrt{2}}{A_{I_d}/\sqrt{2}}\right), \tag{4-47}$$

where the relationship between I and F is given by

$$F = \frac{K_\tau}{r} I\,. \tag{4-48}$$

If the control signal exceeds the current limit of the power supply, the signal will be clipped as shown in Figure 4-15. As the control signal amplitude continues to increase, its RMS value begins to approach that of a square wave with an amplitude of A_{max}. Thus, the maximum gain that an ideal compensator can supply due to motor saturation is given by

$$G_{max} = 20\log_{10}\left(\frac{A_{F_{max}}}{A_{F_d}/\sqrt{2}}\right), \tag{4-49}$$

$$G_{max} = 20\log_{10}\left(\frac{A_{I_{max}}}{A_{I_d}/\sqrt{2}}\right). \tag{4-50}$$

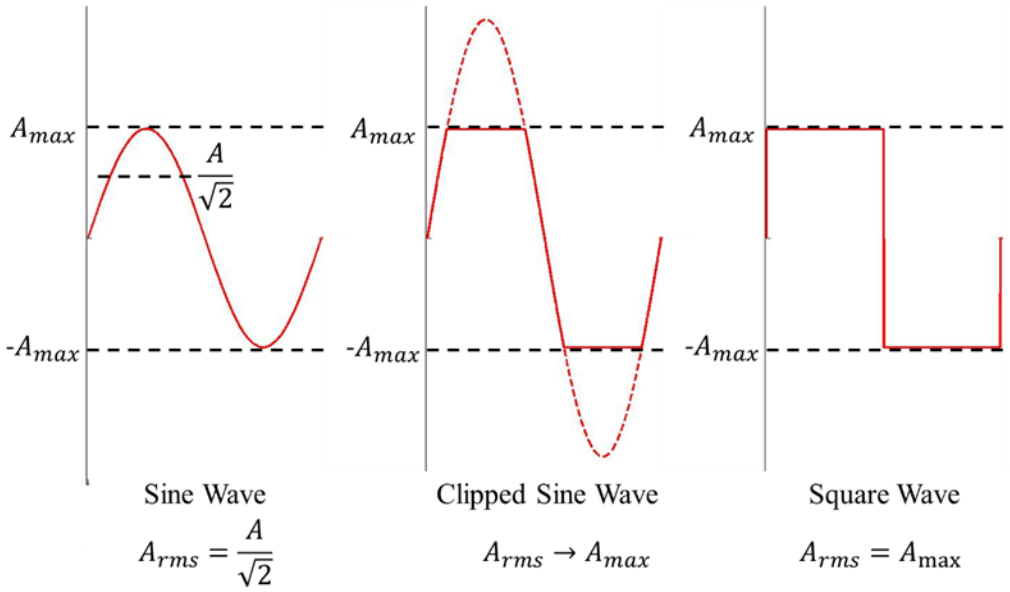

Figure 4-15. Signal clipping due to saturation of a control signal.

For a 2nd order system with a -40 dB/dec roll-off, a controller gain of G_{max} will correspond to an increase in force bandwidth given by

$$\delta_\omega = \frac{G_{max}}{40\ dB/dec}, \tag{4-51}$$

$$\delta_\omega = \frac{1}{2}\log_{10}\left(\frac{A_{I_{max}}}{A_{I_d}/\sqrt{2}}\right) = \frac{1}{2}\log_{10}\left(\frac{A_{F_{max}}}{A_{F_d}/\sqrt{2}}\right), \tag{4-52}$$

where δ_ω is the increase in force bandwidth in decades. Equation 4-52 allows a designer to determine the max possible force bandwidth of a known plant model at a desired force amplitude with known current limits. It should be noted that Equation 4-52 relates the max theoretical performance, assuming a fully saturated control signal and a 2nd order system with -40 dB/dec roll-off. In practice there will be some variation due to the sharpness of the roll-off, model accuracy, and how close the motor is to being fully saturated. Nevertheless, this serves as a useful upper bound on the expected force bandwidth increase of an ideal controller.

4-6-2 Effect of Current Limit

Equations 4-49 and 4-50 shows that G_{max} is a function of both the current limit as well as the input amplitude. Figure 4-16 shows how the controlled response changes as the current limit is varied from 2.5 A to 40 A with a constant input amplitude of 200 N. As the current limit is lowered, less current is available to the controller to compensate for the plant's roll-off, resulting in a reduced force bandwidth. These findings provide a clear argument for SEAs with power supplies capable of providing as much power to the actuator as possible in order to maximize the force bandwidth. However, even with large supply currents, the practical limitations of motor windings and heat dissipation should be considered.

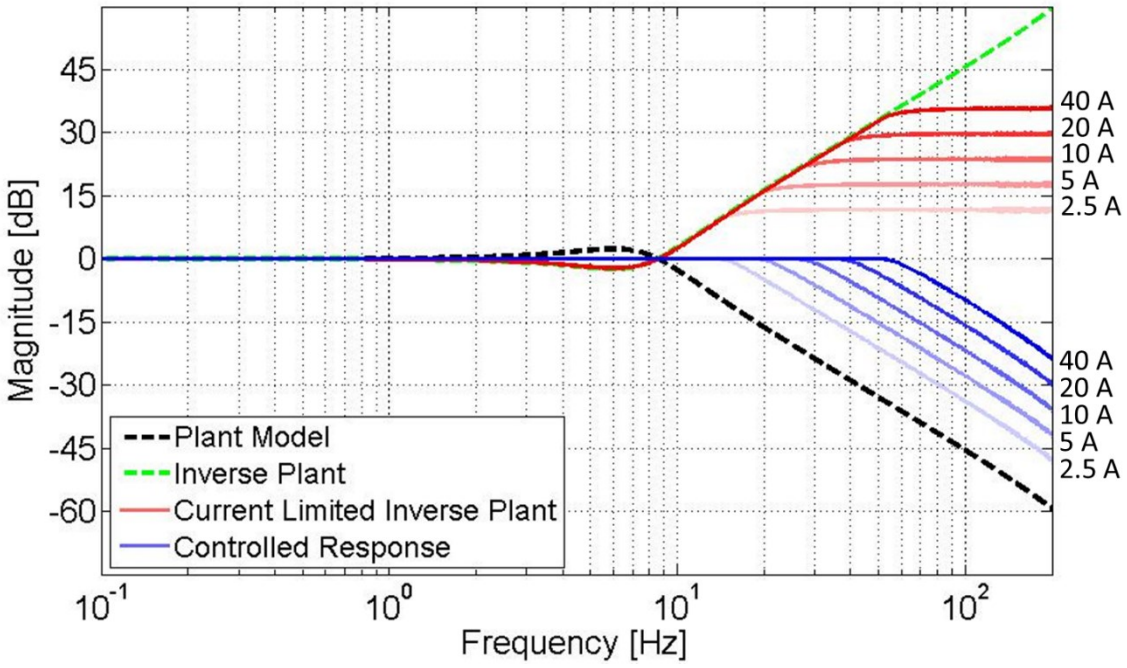

Figure 4-16. Effect of current limit on the controlled performance of an ideal compensator with motor saturation. The input force amplitude is 200 N and the current limit is varied from 2.5 A to 40 A.

4-6-3 Effect of Input Amplitude

Figure 4-17 shows how the controlled response changes as the input amplitude is varied from 50 N to 800N with a constant current limit of 10 A. For smaller input amplitudes, the force bandwidth is higher since there is more available current margin for the controller to compensate for the actuator dynamics. At larger input amplitudes, the controller has less current margin to work with, resulting in a reduced bandwidth. One noteworthy observation is that the set of responses in Figure 4-17 and Figure 4-16 are nearly identical, as can be expected by inspecting Equations 4-49 and 4-50. A twofold increase the available current has the same effect as a twofold decrease in the input amplitude.

Figure 4-17 also reveals an important observation for reporting actuator performance metrics. Even for an ideal compensator, the controlled performance metric of force bandwidth drastically depends on the input amplitude. Interestingly, the vast majority of papers in the SEA literature do not specify the amplitudes for their reported force bandwidth results, which can lead to drastically misleading results. The SEA literature would benefit greatly if authors would include a set of experimental results over a range of input amplitudes to avoid misleading performance metrics.

These findings also raise the issue of what input amplitude a SEA should be designed for when considering the desired force bandwidth. For the THOR-Linear SEA, we have used an amplitude of 200 N for most of our experiments and tuning because our control signals typically oscillate within that range, which is roughly 30 percent of our rated continuous force. However, the SEA literature provides no guidelines or design principles on this topic since most of the literature overlooks the effect of input amplitude altogether.

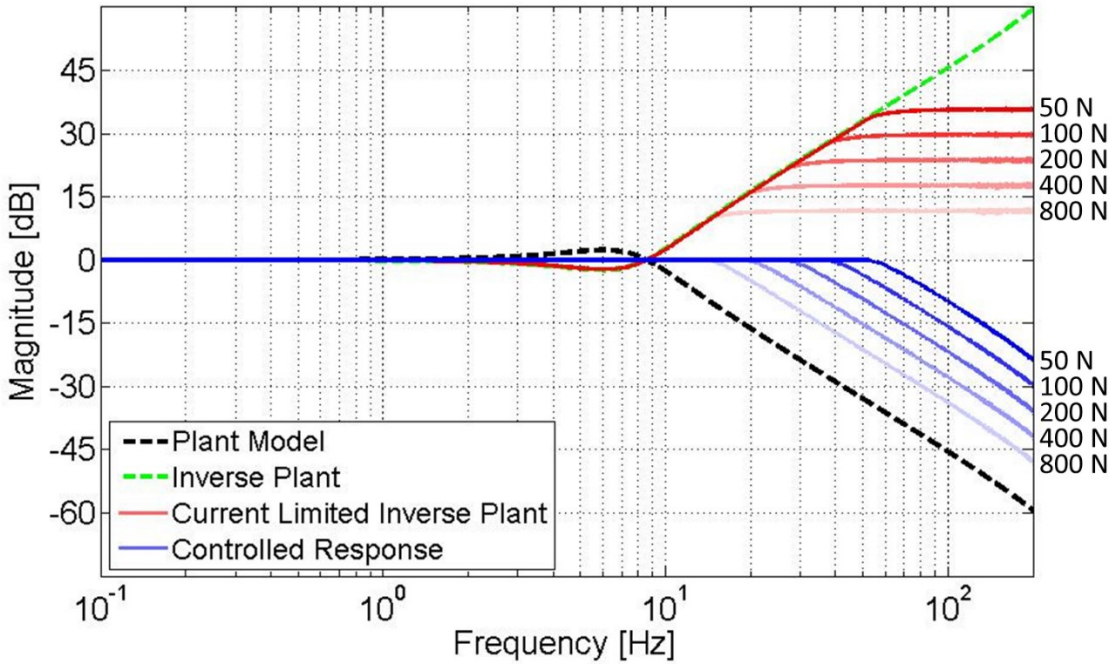

Figure 4-17. Effect of input force amplitude on the controlled performance of an ideal compensator with motor saturation. The current limit is 10 A while the input force amplitude is varied from 50 N to 800 N.

4-6-4 Design Implications | Force Bandwidth

The above results can be summarized by considering the plant model P which relates the baseline actuator dynamics and δ_ω which provides an upper bound on the increase in force bandwidth of an ideal compensator with motor saturation

$$P = \frac{F_m(s)}{F_d(s)} = \frac{k}{\left(\frac{J}{r^2}+m_k\right)s^2+\left(\frac{bm}{r^2}\right)s+k}, \tag{4-41}$$

$$\delta_\omega = \frac{1}{2}\log_{10}\left(\frac{A_{I_{max}}}{A_{I_d}/\sqrt{2}}\right) = \frac{1}{2}\log_{10}\left(\frac{A_{F_{max}}}{A_{F_d}/\sqrt{2}}\right). \tag{4-52}$$

Together, these equations relate the effect of spring stiffness, reflected inertia, damping, current saturation, and input amplitude on the force bandwidth of a Series Elastic Actuator. Actuator improvements can be achieved either by modifying the plant model or the controller. The mechanical design variables define the baseline plant model while the power supply limitations and controller design determine the bandwidth increase of the controlled response. The spring stiffness k only directly impacts the plant model P but this too can serve as a stiffness selection criteria based on a desired force bandwidth. Given the known power limitations and known actuator design variables, a spring stiffness can readily be selected to ensure a desired force bandwidth.

Note that the plant model P would be slightly different for the two spring configurations of the SBN and SMH designs. Namely, the value of m_k would either be m_r or m_p depending on the location of the spring. However, the contribution of the m_k term is very small compared to the reflected motor inertia, due to the high gearing ratios of most SEA designs. Therefore, the spring

location has a negligible effect on the controlled force bandwidth, as long as the pure force source bandwidth ω_{max} is higher than the desired controlled force bandwidth.

These results promote the development of power supplies capable of higher power output as well as accompanying methods for active cooling of motors. It also promotes lower impedance actuators since a smaller reflected inertia leads to a plant model with higher bandwidth while keeping the spring stiffness low for managing impacts and enabling high fidelity force control. These equations also reveal the benefit of a nonlinear spring, specifically, a hardening spring which gets stiffer with increasing force. Since the controlled bandwidth decreases with increasing amplitude, a hardening spring would mechanically increase the plant model's bandwidth for higher forces where the controller has less available current. Finally, since Equations 4-41 and 4-52 give the theoretical maximum performance, they can be used to evaluate the relative performance of an existing controller to determine if there are any more improvements to be made.

4-7 Spring Location & Stiffness Selection

Two of the key design decisions for an SEA is choosing the spring location and spring stiffness. Taking the THOR-Linear SEA as an example, Table 8 provides a comparison of a few possible design configurations. Two stiffness settings are considered, a soft spring constant of 372 kN/m, and a stiff spring constant of 655 kN/m. Both of these settings are achievable with the existing THOR-Linear SEA as described in Section 4-2-2. We also consider the sprung motor housing spring configuration (as designed) as well as a notional sprung ball nut design. The sprung mass for the SMH design is 0.6461 kg while the sprung mass for the SBM design was approximated as 0.150 kg. All of the calculations assume a power supply operating at 48 V with a current limit of 10 A and an input force amplitude of 200 N.

Table 8: Comparison of Different Configurations of the THOR-Linear SAE (I_{max} = 10 A, A_{F_d} = 200 N)

	Soft Spring (372 kN/m)		**Stiff Spring (655 kN/m)**	
	SBN	**SMH**	**SBN**	**SMH**
Plant Natural Frequency: ω_n	6.84 Hz		10.19 Hz	
Pure Force Source Assumption: ω_{max}	65.3 Hz	86.6 Hz	135.4 Hz	179.7 Hz
Ideal Compensator Bandwidth Increase: δ_ω	0.6151 decade			
Maximum Force Bandwidth w/ Ideal Compensator: $\omega_{c_{max}}$	28.2 Hz		42.0 Hz	

Several observations can be made from the values in Table 8.

The bandwidth increase (δ_ω) of the ideal compensator is the same for all four configurations because it depends only on the maximum current and the amplitude of the input.

The plant dynamics for the SBN and SMH designs, given in Table 7 in Section 4-4-5, show that the actuator dynamics differ in that the SBN dynamics use m_r as the sprung mass while the SMH

dynamics use m_p. However, because the plant dynamics are dominated by the much larger reflected inertia J, the natural frequencies ω_n end up being the same for both spring locations, varying only with spring stiffness.

The pure force source assumption, however, depends greatly on the sprung mass, resulting in four different values of ω_{max}. Interestingly, we see that for the THOR-Linear design this is not the limiting factor and all four configurations can reliably be treated as pure force sources.

The limiting factor for the THOR-Linear SEA is the controlled force bandwidth $\omega_{c_{max}}$. The controlled force bandwidth depends only on the plant model and the compensator gain, so the spring location has no effect on the controlled performance. Because of this, we can safely use the sprung motor housing design to take advantage of the benefits outlined in Section 4-3.

Our design goal for the THOR-Linear SEA was to achieve a 15 Hz bandwidth for the soft spring setting and a 30 Hz bandwidth for the stiff spring setting, which is within a manageable margin of the predicted theoretical performance (28.2 Hz and 42.0 Hz, respectively). While the values in Table 8 validate the THOR-Linear SEA design, we should note that an actuator with a larger sprung mass, softer spring, lower current limit, or higher force amplitude demands could have significantly different results.

Of the many design variables for SEAs, spring stiffness is one of the most important because it strongly influences the plant dynamics that the controller has to compensate for. Both the pure force assumption (ω_{max}) and the controlled force bandwidth ($\omega_{c_{max}}$) place a lower bound on the spring stiffness, therefore both should be considered to determine which is the limiting factor. In the THOR-Linear design, the limiting factor is the controlled force bandwidth criteria which is the same for both spring configurations, so the spring location does not matter. In some cases, however, the pure force source assumption could be the limiting factor, in which case the spring location would be much more important.

Therefore, the design approach should be to first determine the desired force amplitude and bandwidth of the application. Then with a known actuator mass and power source, consider both ω_{max} and $\omega_{c_{max}}$ to determine which condition is the limiting factor. Both ω_{max} and $\omega_{c_{max}}$ place a lower bound on the spring stiffness, so as long as both conditions are met, the designer is free to select any stiffness above that minimum. Since a softer spring provides higher force fidelity and reduces impact forces, most force-control applications would benefit from the softest possible spring.

4-8 Conclusions

Series Elastic Actuators are widely used in force-controlled robotic applications where a reliable force source is desired. Unfortunately, the spring element acts as a mechanical filter which limits the force bandwidth and introduces significant actuator dynamics into the control system. Nevertheless, with careful design, the actuator can be treated as a pure force source while also achieving a desired force bandwidth. In this paper we have investigated how the spring stiffness, spring location, sprung mass, and motor saturation of linear Series Elastic Actuators impacts the actuator dynamics and controlled performance. In this section we summarize our findings and show how they relate to these critical design variables.

4-8-1 Model Derivations

Actuator Forces: In legged robot applications, the effective ground link of an actuator changes depending on the phase of the walking cycle. Therefore, it is important to consider the forces acting on both sides of the actuator.

Modeling: Actuators with screw-type transmissions that convert rotary motion into linear motion can intuitively be modeled using a rack & pinion depiction. Unlike the widely used lumped model, the rack & pinion model accurately predicts the forces acting on both sides of the actuator.

Spring Location: The literature on SEAs often discusses the location of the spring element relative to the output. However, we have shown that the output side changes with the phase of the walking cycle. Instead, the distinction should be related to where the spring is located within the internal design of the actuator itself. We show how the two most common linear SEA designs vary by introducing and analyzing the Sprung Ball Nut and Sprung Motor Housing designs.

4-8-2 Actuator Dynamics

Pure Force Source Assumption: While it is common to treat SEAs as a pure force source, we show that this assumption is not always valid by comparing the forces acting on either side of the actuator. We derive the relationship between the two output forces and show that the allowable force bandwidth for the pure force source assumption is independent of the controller and depends only on the spring stiffness and the effective sprung mass of the actuator.

Sprung Mass: The key difference between the two common linear SEA designs is the effective sprung mass. The Sprung Ball Nut design can be designed to have a much smaller sprung mass, which offers some significant benefits, but is more difficult to instrument and reduces the linear stroke of the actuator. The Sprung Motor Housing design has a larger sprung mass, however, we show that for common linear SEA designs the SMH design approach can achieve an acceptable force bandwidth range.

Sensor Location: As long as the actuator is used within the allowable bandwidth given by the pure force source assumption, the force sensor location does not matter. It is common to measure the spring deflection to estimate force, but an inline load cell can also be used either at the spring interface or anywhere else within the actuator's line of action.

Stiffness Selection Criteria: We use the pure force source assumption to propose a new stiffness selection criteria which ensures that the actuator can be treated as an idea force source up to a desired bandwidth. This criteria is independent of the controller and depends only on the desired force bandwidth and internal sprung mass of the actuator.

4-8-3 Controlled Performance

Motor Saturation: The effect of motor saturation is that of leveling off the compensator response at high frequencies. Due to saturation, there is a maximum gain that an ideal controller can achieve for a given input signal amplitude. This finding promotes the development of more capable power supplies.

Improving Bandwidth: Increasing spring stiffness, reducing reflected inertia, and increasing available power are all ways to increase the force bandwidth of an SEA but should be balanced against other practical considerations including impact loading of the actuator and heat dissipation limitations.

Reporting Metrics: Since the controlled bandwidth depends on the input amplitude, the SEA literature should specify force amplitude when reporting performance metrics such as bandwidth. Actuators should also be evaluated over a range of force amplitudes to more accurately represent their performance.

4-8-4 Future Work

While the analysis in this paper is specific to linear Series Elastic Actuators, the same approach should be extended to consider other designs, especially geared rotary and cable driven SEAs. The results in this paper also support the development of actuators with nonlinear stiffness springs as well as more capable power supplies in an effort to improve the performance of linear SEAs.

5 Conclusions

As the field of compliant humanoids advances, there will be a continuing need for higher performance actuators, more descriptive dynamic models, and more sophisticated controllers. This dissertation has sought to contribute to the field by addressing the design and modeling of linear Series Elastic Actuators as they are used in humanoid robots.

Chapter 2 introduced Configurable Compliance, a novel approach to implementing compliance in linear Series Elastic Actuators. Two versions of Configurable Compliance were presented; a perpendicular orientation which is used for the SAFFiR-Linear SEA and a parallel orientation which is used in the THOR-Linear SEA. Both designs use a cantilevered beam design which incorporates a movable or removable pivot in order to adjust the passive stiffness of the spring element.

Chapter 3 developed a new unlumped model for linear Series Elastic Actuators. The model uses a rack & pinion representation which is intuitive to understand and faithfully predicts the true actuator dynamics, unlike the widely used lumped model.

Chapter 4 used the new rack & pinion model to investigate the design implications of several critical design variables including spring location, spring stiffness, sprung mass, motor saturation, and input amplitude.

The work presented in this dissertation is based on several years of developing novel actuators for two full-scale compliant humanoid robots, SAFFiR and THOR. Both are state-of-the-art robots intended to operate in the extremely challenging real world scenarios of shipboard firefighting and disaster response.

5-1 SAFFiR & THOR Results

Figure 5-1 and Figure 5-2 show walking results for SAFFiR and THOR. The SAFFiR Prototype robot has demonstrated stable compliant walking on loose gravel, 50 mm thick synthetic grass, and loose sand. To date, SAFFiR is the first and only full-scale biped to demonstrate walking in sand. The THOR robot has demonstrated stable compliant walking on level concrete, loose gravel, and 50 mm thick synthetic grass.

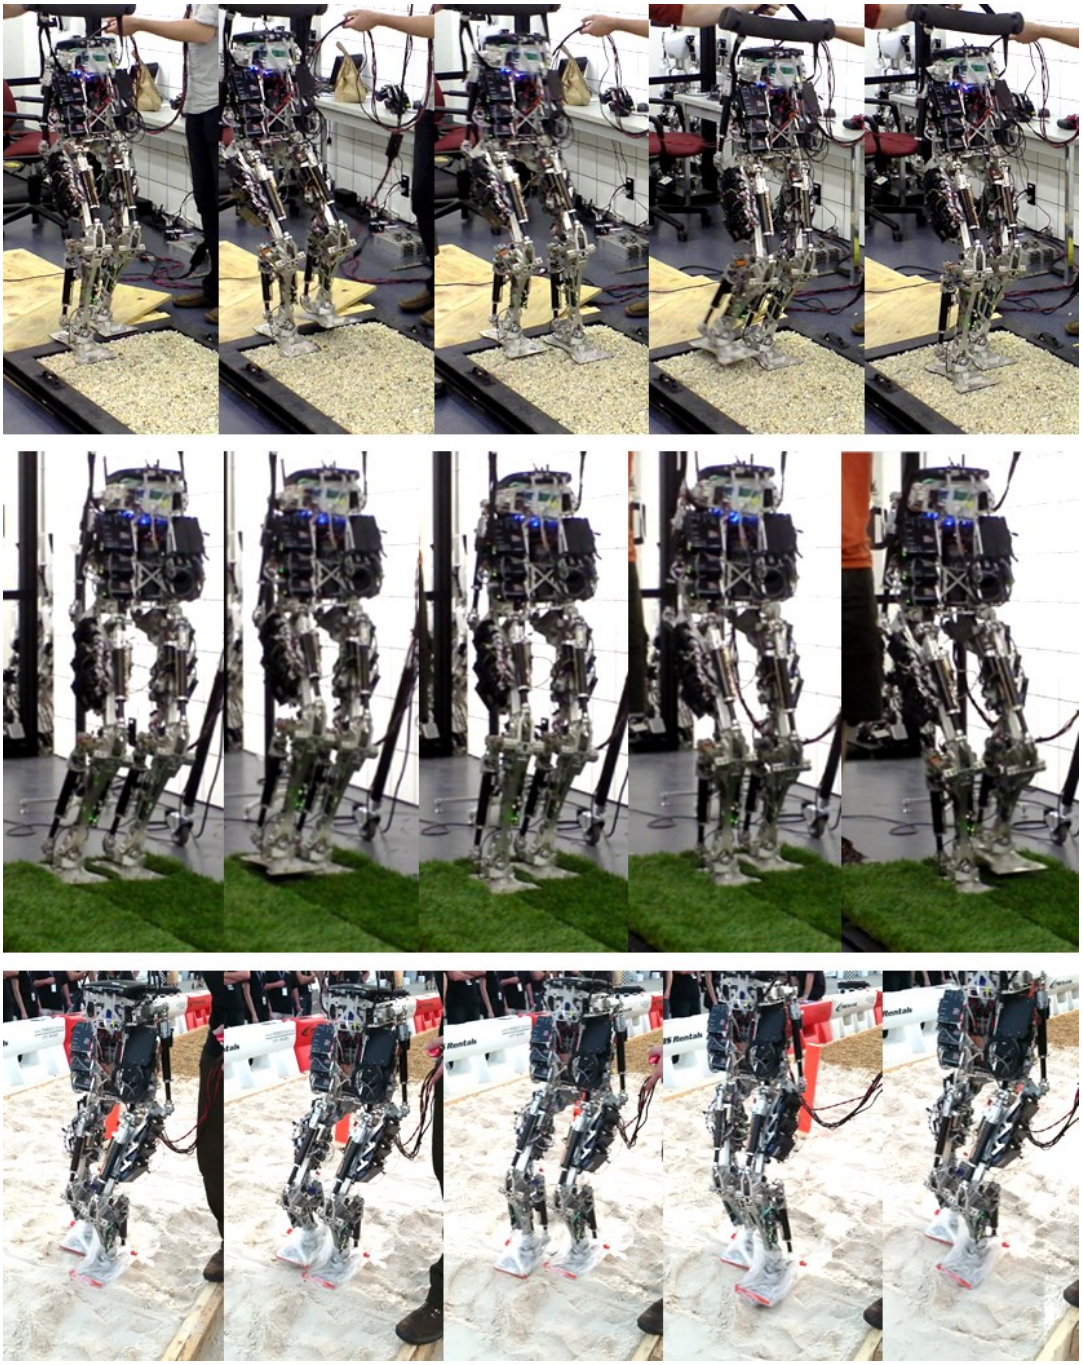

Figure 5-1. SAFFiR Prototype robot demonstrating stable walking on gravel, grass, and sand. Images used with permission of D. Lahr.

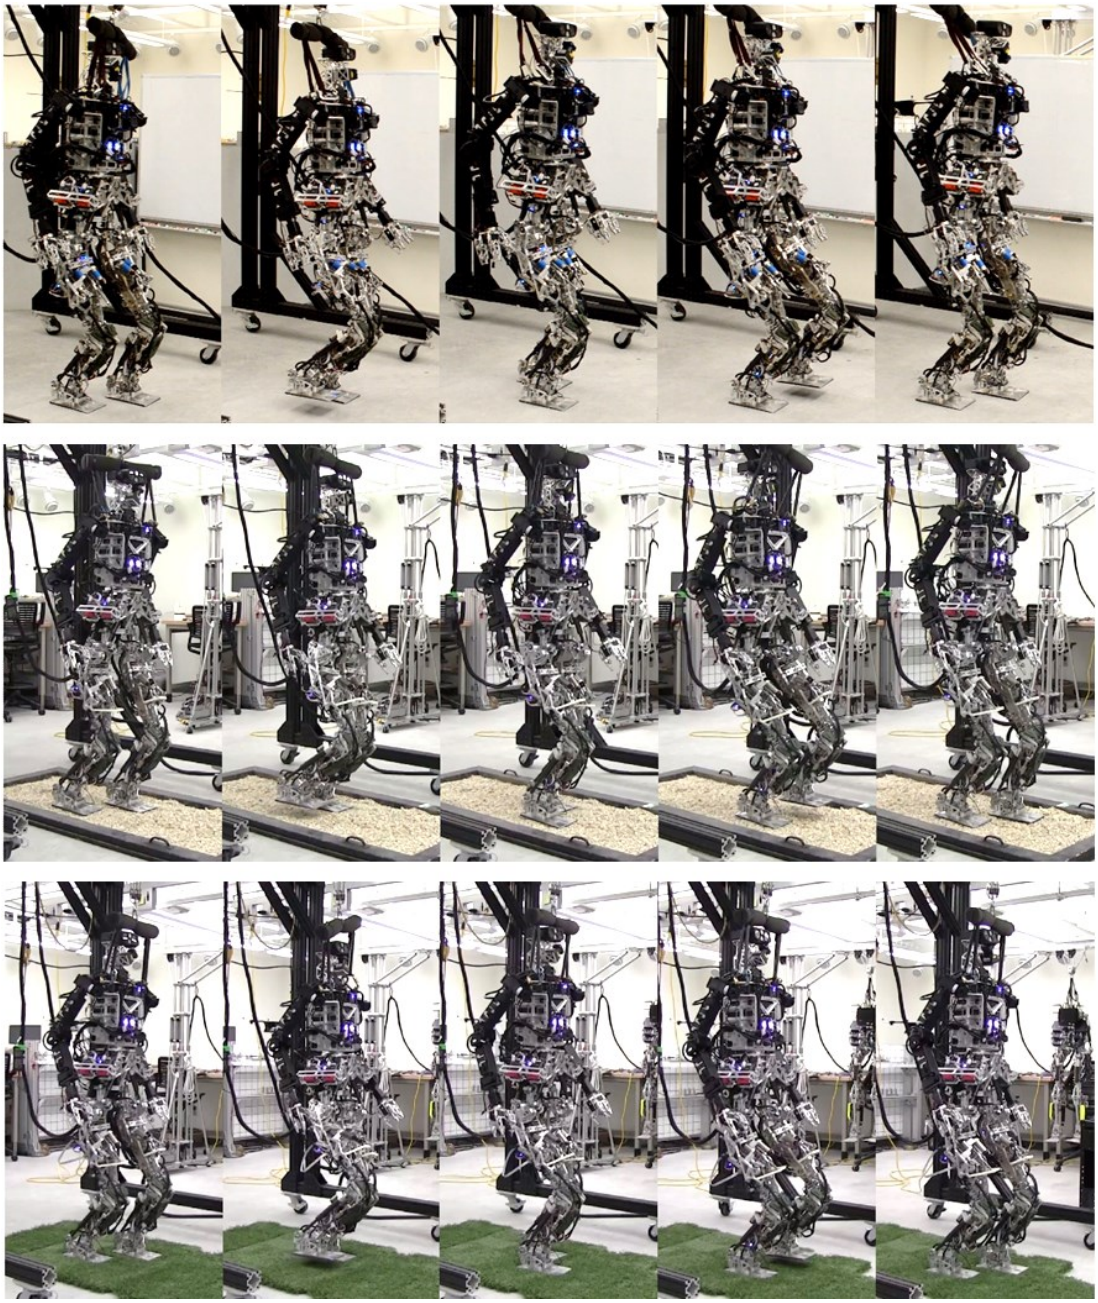

Figure 5-2. THOR robot demonstrating stable walking on concrete, gravel, and grass. Images used with permission of M. Hopkins.

5-2 Future Work

The results of this work provide some useful insights and motivation for future work in compliant actuators.

Rotary Series Elastic Actuators: While linear actuators provide several benefits, rotary actuators are more widely used in robotics. The same modeling and analysis approach in this dissertation should be extended to consider rotary SEA designs. Based on the work in this dissertation it is

believed that the conventional lumped model may also be a poor representation for rotary actuators with transmissions and may only be valid for direct drive rotary actuators.

Design Features: Since so much of the literature on SEAs uses the lumped model, it would be of significant interest to revisit the existing literature that has investigated the role of non-linear springs, parallel springs, variable stiffness, series/parallel damping, variable damping, clutches, and multi-speed transmissions.

Power Supplies & Active Cooling: A major limitation of the existing SEA designs is in the power limitations due to motor saturation. More capable power supplies could possibly be the lowest hanging fruit for improving performance rather than more complex designs or more sophisticated controllers. Along with more power, however, there will likely be a need to actively cool the motor due to high thermal losses in the windings at high duty cycles.

Nonlinear Springs & Damping: Of the many possible design features, nonlinear springs and series/parallel damping seem to be the most fruitful avenues to pursue next. A nonlinear spring which is stiffer at high loads would be able to counteract the drop in bandwidth at higher force amplitudes while retaining force fidelity at low amplitudes. Parallel damping would have a similar effect except the hardening behavior would be a function of velocity rather than force. Damping has the benefit of reducing resonances but would also result in higher impact loads on the actuator.

Parallel Actuators: Most of the literature on SEAs considers a single actuator acting across a single degree of freedom joint. It would be useful to extend the work in this dissertation to consider parallelly actuated milti-DOF joints like the ankles and hips in SAFFiR and THOR. Parallelly actuated joints may introduce additional resonance modes due to multiple spring elements as well as different effective stiffnesses at each DOF depending on the actuator arrangements. Similarly, it would be useful to consider how the compliance, actuator dynamics, and intermediate inertias of SEAs affects the whole body dynamics of a compliant humanoid.

Dual Elastic Elements: When the spring element on the THOR-Linear SEA was locked out on a test stand, the measured stiffness of the actuator only increased by a factor of four. This was due to deflections within the actuator's internal components, but especially due to deflections in the u-joints at either end. A novel design approach could be to intentionally design the u-joints (or other attachment methods) at either end to serve as two series elastic members, one on either side of the actuator. The rack & pinion model could be extended to consider the effect of a spring on both sides of the actuator.

Wobbling Masses: In contrast to the rigidity of robots, many biological systems are characterized by a skeletal structure surrounded by soft-tissue masses. Recent work [76] has looked at how the bulk motion of soft-tissue, termed "wobbling masses", contributes to the dissipation of energy from impacts. The same behavior could be leveraged for reducing impact loads in humanoids, either by introducing soft-tissue, or by intelligently using the deflection of compliant actuators to achieve similar effects.

References

[1] D. Lahr, V. Orekhov, B. Lee, and D. Hong, "Early developments of a parallelly actuated humanoid, SAFFiR," in *ASME International Design Engineering Technical Conferences*, 2013.

[2] B. Lee, C. Knabe, V. Orekhov, and D. Hong, "Design of a human-like range of motion hip joint for humanoid robots," in *ASME International Design Engineering Technical Conferences*, 2013.

[3] S. Behnke, "Humanoid robots – from fiction to reality?," *Künstliche Intelligenz H.*, vol. 4, pp. 5–9, 2008.

[4] S. H. Collins and A. D. Kuo, "Recycling energy to restore impaired ankle function during human walking," *PLoS One*, vol. 5, no. 2, 2010.

[5] S. K. Au and H. M. Herr, "Powered ankle-foot prosthesis," *IEEE Robotics & Automation Magazine*, vol. 15, no. 3, pp. 52–59, Sep-2008.

[6] R. Alexander, "Three Uses for Springs in Legged Locomotion," *Int. J. Rob. Res.*, vol. 9, no. 2, pp. 53–61, 1990.

[7] S. N. Patek, D. M. Dudek, and M. V. Rosario, "From bouncy legs to poisoned arrows: elastic movements in invertebrates," *J. Exp. Biol.*, vol. 214, pp. 1973–1980, 2011.

[8] T. J. Roberts and E. Azizi, "Flexible mechanisms: the diverse roles of biological springs in vertebrate movement," *J. Exp. Biol.*, vol. 214, pp. 353–361, 2011.

[9] K. W. Hollander, R. Ilg, T. G. Sugar, and D. E. Herring, "An efficient robotic tendon for gait assistance," *J. Biomech. Eng.*, vol. 128, pp. 788–791, 2006.

[10] G. Nelson, A. Saunders, N. Neville, B. Swilling, J. Bondaryk, D. Billings, C. Lee, R. Playter, and M. Raibert, "PETMAN: A humanoid robot for testing chemical protective clothing," *J. Robot. Soc. Japan*, vol. 30, no. 4, pp. 372–377, 2012.

[11] "Atlas - The Agile Anthropomorphic Robot." [Online]. Available: http://www.bostondynamics.com/robot_Atlas.html.

[12] 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 *Conference on Intelligent Robots and Systems*, 2011, pp. 4400–4407.

[13] Y. Sakagami, R. Watanabe, C. Aoyama, S. Matsunaga, N. Higaki, and K. Fujimura, “The intelligent ASIMO: system overview and integration,” in *International Conference of Intelligent Robots and Systems*, 2002, pp. 2478–2483.

[14] D. Hobbelen, T. de Boer, and M. Wisse, “System overview of bipedal fobots Flame and TUlip: tailor-made for limit cycle walking,” in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2008, pp. 2486–2491.

[15] J. E. Pratt and B. T. Krupp, “Design of a bipedal walking robot,” *Proc. SPIE*, vol. 6962, 2008.

[16] L. Colasanto, N. G. Tsagarakis, and D. G. Caldwell, “A compact model for the compliant humanoid robot COMAN,” *2012 4th IEEE RAS EMBS Int. Conf. Biomed. Robot. Biomechatronics*, pp. 688–694, Jun. 2012.

[17] C. Ott, C. Baumgartner, J. Mayr, M. Fuchs, R. Burger, D. Lee, O. Eiberger, A. Albu-Schaffer, M. Grebenstein, and G. Hirzinger, “Development of a biped robot with torque controlled joints,” *IEEE-RAS Int. Conf. Humanoid Robot.*, pp. 167–173, 2010.

[18] N. Paine, J. S. Mehling, J. Holley, N. A. Radford, G. Johnson, C.-L. Fok, and L. Sentis, “Actuator control for the NASA-JSC Valkyrie humanoid robot : a decoupled dynamics approach for torque control of series elastic robots,” *J. F. Robot.*, 2014.

[19] N. A. Paine, “High-Performance Series Elastic Actuation,” University of Texas at Austin, 2014.

[20] S. Lohmeier, T. Buschmann, and H. Ulbrich, “Humanoid robot LOLA,” in *IEEE International Conference on Robotics and Automation*, 2009, pp. 775–780.

[21] F. Pfeiffer, “The TUM walking machines,” *Philos. Trans. R. Scociety A*, vol. 365, pp. 109–131, Jan. 2007.

[22] B. Vanderborght, “Dynamic stabilisation of the biped Lucy powered by actuators with controllable stiffness,” Vrije Universiteit Brussel, 2007.

[23] J. A. Grimes and J. W. Hurst, “The design of ATRIAS 1.0 a unique monopod, hopping robot,” in *International Conference on Climbing and Walking Robots*, 2012, pp. 548–554.

[24] J. W. Hurst, “The role and implementation of compliance in legged locomotion,” Carnegie Mellon University, 2008.

[25] G. A. Pratt and M. M. Williamson, “Series elastic actuators,” *Int. Conf. Intell. Robot. Syst.*, vol. 1, pp. 399–406, 1995.

[26] V. Orekhov, D. Lahr, B. Lee, and D. Hong, “Configurable compliance for series elastic actuators,” in *ASME International Design Engineering Technical Conferences*, 2013.

[27] D. P. Ferris and C. T. Farley, “Interaction of leg stiffness and surface stiffness during human hopping,” *J. Appl. Physiol.*, vol. 82, pp. 15–22, 1997.

[28] D. P. Ferris, M. Louie, and C. T. Farley, “Running in the real world: adjusting leg stiffness for different surfaces,” *R. Soc. London*, vol. 265, pp. 989–994, 1998.

[29] J. W. Hurst, J. E. Chestnutt, and A. A. Rizzi, “The actuator with mechanically adjustable series compliance,” *IEEE Trans. Robot.*, vol. 26, no. 4, pp. 597–606, 2010.

[30] R. Van Ham, T. G. Sugar, B. Vanderborght, K. W. Hollander, and D. Lefeber, “Compliant actuator designs,” *IEEE Robot. Autom.*, pp. 81–94, 2009.

[31] B. Vanderborght, A. Albu-Schaeffer, A. Bicchi, E. Burdet, D. G. Caldwell, R. Carloni, M. Catalano, O. Eiberger, W. Friedl, G. Ganesh, M. Garabini, M. Grebenstein, G. Grioli, S. Haddadin, H. Hoppner, A. Jafari, M. Laffranchi, D. Lefeber, F. Petit, S. Stramigioli, N. Tsagarakis, M. Van Damme, R. Van Ham, L. C. Visser, and S. Wolf, “Variable impedance actuators: a review,” *Rob. Auton. Syst.*, vol. 61, pp. 1601–1614, 2013.

[32] B. Lee, V. Orekhov, L. Derek, and D. Hong, “Design and measurment error analysis of a low-friction, lightweight linear series elastic actuator,” in *ASME International Design Engineering Technical Conferences*, 2013.

[33] C. Knabe, B. Lee, V. Orekhov, and D. Hong, “Design of a compact, lightweight, electromechanical linear series elastic actuator,” in *ASME International Design Engineering Technical Conferences*, 2014.

[34] C. Knabe, B. Lee, and D. Hong, “An inverted straight line mechanism for augmenting joint range of motion in a humanoid robot,” in *ASME International Design Engineering Technical Conferences*, 2014.

[35] O. Tabata, S. Konishi, P. Cusin, Y. Ito, F. Kawai, S. Hirai, and S. Kawamura, “Microfabricated tunable bending stiffness device,” *Proc. IEEE Int. Conf. Micro Electro Mech. Syst.*, pp. 23–27, 2000.

[36] S. Kawamur, T. Yamamoto, D. Ishida, T. Ogata, Y. Nakayama, O. Tabata, and S. Sugiy, “Development of passive elements with variable mechanical impedance for wearable robots,” pp. 248–253, 2002.

[37] T. Morita and S. Sugano, “Design and development of a new robot joint using a mechanical impedance adjuster,” in *IEEE International Conference on Robotics and Automation*, 1995, pp. 2469–2475.

[38] T. Morita and S. Sugano, “Development of an anthropomorphic force-controlled manipulator WAM-10,” in *International Conference on Advanced Robotics*, 1997, pp. 701–706.

[39] R. Wang and H. Huang, “ADEA—Active variable stiffness differential elastic actuator: design and application for safe robotics,” in *IEEE International Conference on Robotics and Biomimetics*, 2011, pp. 2768–2773.

[40] J. Choi, S. Hong, W. Lee, S. Kang, and M. Kim, “A robot joint with variable stiffness using leaf springs,” *IEEE Trans. Robot.*, vol. 27, no. 2, pp. 229–238, 2011.

[41] R.-J. Wang and H.-P. Huang, “AVSER — Active variable stiffness exoskeleton robot system: design and application for safe active-passive elbow rehabilitation,” *IEEE/ASME Int. Conf. Adv. Intell. Mechatronics*, pp. 220–225, 2012.

[42] A. De and U. Tasch, “A two-DOF manipulator with adjustable compliance capabilities and comparison with the human finger,” *J. Robot. Syst.*, vol. 13, no. 1, pp. 25–34, 1996.

[43] E. Campbell, W. Hered, A. J. Lynch, M. K. O’Malley, and J. McLurkin, “Design of a low-cost series elastic actuator for multi-robot manipulation,” in *IEEE International Conference on Robotics and Automation*, 2011, pp. 5395–5400.

[44] L. L. Howell, *Compliant Mechanisms*. New York, NY: John Wiley & Sons, 2001, p. 449.

[45] R. C. Juvinall and K. M. Marshek, *Fundamentals of Machine Component Design*, 4th Editio. 2006, pp. 814–817.

[46] J. Pratt, B. Krupp, and C. Morse, “Series elastic actuators for high fidelity force control,” *Ind. Rob.*, vol. 29, no. 3, pp. 234–241, 2002.

[47] J. W. Hurst, A. A. Rizzi, and D. Hobbelen, “Series elastic actuation: potential and pitfalls,” in *International Conference on Climbing and Walking Robots*, 2004.

[48] D. W. Robinson, J. E. Pratt, D. J. Paluska, and G. A. Pratt, “Series elastic actuator development for a biomimetic walking robot,” in *IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, 1999, pp. 561–568.

[49] G. A. Pratt, “Low impedance walking robots,” *Integr. Comp. Biol.*, vol. 42, pp. 174–181, 2002.

[50] K. Kong, J. Bae, and M. Tomizuka, “A compact rotary series elastic actuator for human assistive systems,” *IEEE/ASME Trans. Mechatronics*, vol. 17, no. 2, pp. 288–297, 2012.

[51] D. Accoto, G. Carpino, F. Sergi, N. L. Tagliamonte, L. Zollo, and E. Guglielmelli, "Design and characterization of a novel high-power series elastic actuator for a lower limb robotic orthosis," *Int. J. Adv. Robot. Syst.*, vol. 10, no. 359, pp. 1–12, 2013.

[52] N. Paine, S. Oh, and L. Sentis, "Design and control considerations for high performance series elastic actuators," *IEEE/ASME Trans. Mechatronics*, vol. 19, no. 3, pp. 1080–1091, 2014.

[53] S. A. Migliore, E. A. Brown, and S. P. DeWeerth, "Biologically inspired joint stiffness control," in *International Conference on Robotics and Automation*, 2005, pp. 4508–4513.

[54] M. Laffranchi, N. Tsagarakis, and D. G. Caldwell, "A compact compliant actuator (CompAct) with variable physical damping," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 4644–4650.

[55] S. Wolf and G. Hirzinger, "A new variable stiffness design: matching requirements of the next robot generation," in *IEEE International Conference on Robotics and Automation*, 2008, pp. 1741–1746.

[56] L. Mooney and H. Herr, "Continuously-variable series-elastic actuator," in *International Conference on Rehabilitation Robotics*, 2013, pp. 1–6.

[57] E. J. Rouse, L. M. Mooney, E. C. Martinez-Villalpando, and H. M. Herr, "Clutchable series-elastic actuator: design of a robotic knee prosthesis for minimum energy consumption," in *International Conference on Rehabilitation Robotics*, 2013, pp. 1–6.

[58] G. A. Pratt, P. Willisson, C. Bolton, and A. Hofman, "Late motor processing in low-impedance robots: impedance control of series-elastic actuators," in *Proceeding of the 2004 American Control Conference*, 2004, pp. 3245–3251.

[59] K. Kong, J. Bae, and M. Tomizuka, "Control of rotary series elastic actuator for ideal force-mode actuation in human-robot interaction applications," *IEEE/ASME Trans. Mechatronics*, vol. 14, no. 1, pp. 105–118, 2009.

[60] L. M. S. do Amaral, B. Jardim, and A. A. G. Siqueira, "Robust force and impedance control of series elastic actuators," *ABCM Symp. Ser. Mechatronics*, vol. 5, pp. 212–221, 2012.

[61] M. Grun, R. Muller, and U. Konigorski, "Model based control of series elastic actuators," in *IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics*, 2012, pp. 538–543.

[62] J. W. Hurst and A. A. Rizzi, "Series compliance for an efficient running gait," *IEEE Robotics & Automation Magazine*, vol. 15, no. 3, pp. 42–51, 2008.

[63] J. W. Sensinger, L. E. Burkart, G. A. Pratt, and R. F. Weir, “Effect of compliance location in series elastic actuators,” *Robotica*, vol. 31, no. 8, pp. 1313–1318, 2013.

[64] M. Lauria, M.-A. Legault, M.-A. Lavoie, and F. Michaud, “Differential elastic actuator for robotic interaction tasks,” in *IEEE International Conference on Robotics and Automation*, 2008, pp. 3606–3611.

[65] A. Edsinger-Gonzales and J. Weber, “Domo: A force sensing humanoid robot for manipulation research,” *Int. Conf. Humanoid Robot.*, vol. 1, pp. 273 – 291, 2004.

[66] D. Paluska and H. Herr, “The effect of series elasticity on actuator power and work output: implications for robotic and prosthetic joint design,” *Rob. Auton. Syst.*, vol. 54, no. 8, pp. 667–673, 2006.

[67] K. Kemper, H. R. Vejdani, B. Piercy, and J. Hurst, “Optimal passive dynamics for physical interaction: catching a mass,” *Actuators*, vol. 2, pp. 45–58, 2013.

[68] H. R. Vejdani and J. W. Hurst, “Optimal passive dynamics for physical interaction: throwing a mass,” in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 796–801.

[69] V. L. Orekhov, C. S. Knabe, M. A. Hopkins, and D. W. Hong, “An unlumped model for linear series elastic actuators with ball screw drives,” in *IEEE International Conference on Robotics and Automation*, 2015.

[70] B. Vanderborght, R. Van Ham, D. Lefeber, T. G. Sugar, and K. W. Hollander, “Comparison of mechanical design and energy consumption of adaptable, passive-compliant actuators,” *Int. J. Rob. Res.*, vol. 28, no. 1, pp. 90–103, 2009.

[71] B. Vanderborght, B. Verrelst, R. Van Ham, M. Van Damme, P. Beyl, and D. Lefeber, “Development of a compliance controller to reduce energy consumption for bipedal robots,” *Auton. Robots*, vol. 24, no. 4, pp. 419–434, 2008.

[72] M. Grimmer and A. Seyfarth, “Stiffness adjustment of a series elastic actuator in an ankle-foot prosthesis for walking and running: the trade-off between energy and peak power pptimization,” in *IEEE International Conference on Robotics and Automation*, 2011, pp. 1439–1444.

[73] M. Grimmer and A. Seyfarth, “Stiffness adjustment of a series elastic actuator in a knee prosthesis for walking and running: the trade-off between energy and peak power optimization,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011, pp. 1811–1816.

[74] M. Grimmer, M. Eslamy, and A. Seyfarth, “Energetic and peak power advantages of series elastic actuators in an actuated prosthetic leg for walking and running,” *Actuators*, vol. 3, no. 1, pp. 1–19, 2014.

[75] N. G. Tsagarakis, S. Morfey, G. M. Cerda, Z. Li, and D. G. Caldwell, “COMpliant huMANoid COMAN: optimal joint stiffness tuning for modal frequency control,” in *IEEE International Conference on Robotics and Automation*, 2013, pp. 673–678.

[76] S. Schmitt and M. Günther, “Human leg impact: energy dissipation of wobbling masses,” *Arch. Appl. Mech.*, vol. 81, no. 7, pp. 887–897, 2011.