

Actuator Control for the NASA-JSC Valkyrie Humanoid Robot: A Decoupled Dynamics Approach for Torque Control of Series Elastic Robots

Nicholas Paine

Electrical and Computer Engineering Department
The University of Texas
Austin, TX 78712 USA
npaine@utexas.edu

Joshua S. Mehling

NASA Johnson Space Center
Houston, Texas, TX 77058 USA
joshua.s.mehling@nasa.gov

James Holley

NASA Johnson Space Center
Houston, Texas, TX 77058 USA
james.j.holley@nasa.gov

Nicolaus A. Radford

NASA Johnson Space Center
Houston, Texas, TX 77058 USA
nicolaus.a.radford@nasa.gov

Gwendolyn Johnson

Mechanical Engineering Department
The University of Texas
Austin, TX 78712 USA
gwendolynbrook@gmail.com

Chien-Liang Fok

Mechanical Engineering Department
The University of Texas
Austin, TX 78712 USA
liangfok@utexas.edu

Luis Sentis

Mechanical Engineering Department
The University of Texas
Austin, TX 78712 USA
lsentis@austin.utexas.edu

Abstract

This paper discusses the actuator-level control of Valkyrie, a new humanoid robot designed by NASA's Johnson Space Center in collaboration with several external partners. We focus on several topics pertaining to Valkyrie's series elastic actuators including control architecture, controller design, and implementation in hardware. A decentralized approach is taken in controlling Valkyrie's many series elastic degrees of freedom. By conceptually decoupling actuator dynamics from robot limb dynamics, we simplify the problem of controlling a highly complex system and streamline the controller development process compared to other approaches. This hierarchical control abstraction is realized by leveraging disturbance observers in the robot's joint-level torque controllers. We apply a novel analysis technique to understand the ability of a disturbance observer to attenuate the effects of unmodeled dynamics. The performance of our control approach is demonstrated in two ways. First, we characterize torque tracking performance of a single Valkyrie actuator in terms of controllable torque resolution, tracking error, bandwidth, and power consumption. Second, we perform tests on Valkyrie's arm, a serial chain of actuators, and demonstrate its ability to accurately track torques with our decentralized control approach.

1 Introduction

As an entry into the 2013 DRC Trials, NASA-JSC formed a team with several external partners and led the development of Valkyrie (see Figure 1). Valkyrie was designed to perform tasks required both for responding to disasters, such as the Fukushima Daiichi nuclear disaster in Japan (Nagatani et al., 2013), and for advancing human spaceflight by one day assisting human explorers in extraterrestrial settings such as Mars. While many areas of expertise were required to produce a new humanoid robot in less than 12 months, this paper focuses on one area, namely the methods used to control Valkyrie’s series elastic actuators (SEAs). In the following discussion, we present our overall control approach, an outline of Valkyrie’s multi-joint control architecture, the actuator controller design and its implementation in hardware.

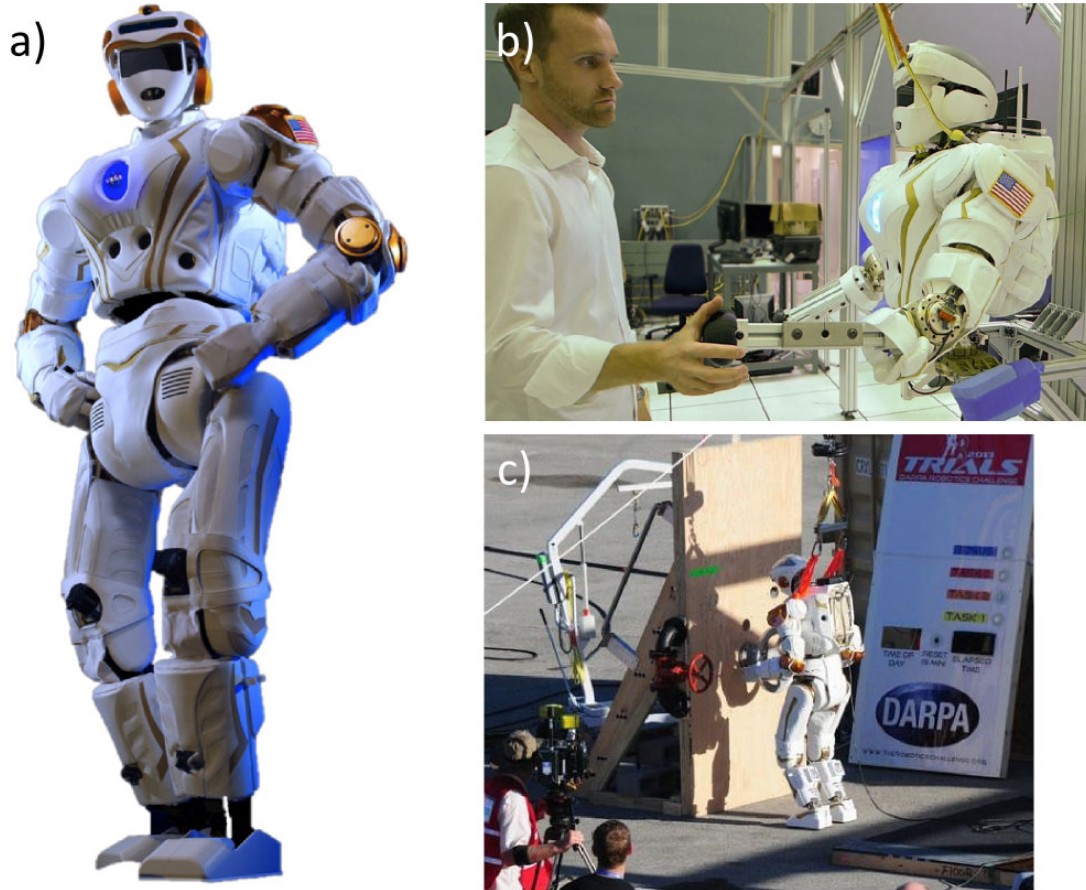


Figure 1: **a) NASA-JSC’s Valkyrie Humanoid Robot.** Valkyrie has 44 actuated degrees of freedom and a suite of sensors including stereo vision, laser range finders, sonar depth perception, and tactile feedback. **b) Safe Human Interaction.** Series elastic actuators and compliant torque control enable safe interaction with humans and with unexpected environmental collisions. **c) DRC Competition.** Valkyrie was one of 16 entries in the 2013 DARPA Robotics Challenge Trials. In this image Valkyrie is performing a valve turn task.

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

For Valkyrie, the actuation control requirements were driven by the needs of the holistic robot control approach and by the mechanical design of its actuators. Because the DRC tasks required movement and ma-

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

1.1 Multi-Joint Series Elastic Control Architecture Background

How to best address compliance within the context of a full body coordinated control architecture is an open question. Early investigations into elastic joint robots treated each motor as an ideal torque source and controlled both actuator and limb dynamics with a centralized algorithm (Spong, 1987). Signal latencies and the complexities of modeling the higher order dynamics inherent in a multi-Degree-Of-Freedom (DOF), series elastic robot are difficult to overcome. Nevertheless, approaches based on this method are widely used today and prove quite effective (Albu-Schaffer et al., 2007; Ott et al., 2008; Li et al., 2012). It has been suggested, however, that robots, particularly those designed to mimic the broad versatility of humans and animals, could benefit from a more embedded, collocated control of their actuators (Pratt et al., 2004). Physically, or at least conceptually, decentralizing joint-level control allows for the assumption of a rigid body model at the high level, actuated by ideal joint torque sources at the low level. This approach is taken in (Kim et al., 2012; Sentis et al., 2013; Hutter et al., 2013). The success of such decentralized control architectures directly depends on the performance capabilities of joint-level torque controllers, which must closely match an ideal torque source model.

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

1.2 Joint-Level SEA Control Background

Many different joint-level control architectures exist for torque control of series elastic actuators. Some measure spring force and control motor force using some variant of PID control structures (P, PD, etc.) (Pratt and Williamson, 1995; Sensinger and Weir, 2006; Hurst et al., 2010; Ragonese et al., 2011; Garcia et al., 2011). If friction and backlash are large, a single-loop PID force controller may become unstable before the desired force tracking is achieved. To remedy this issue, an inner position or velocity control loop may be used with an outer force control loop as proposed by (Robinson, 2000) and (Pratt et al., 2004). This idea has been adopted and carried on by many others, translating force control into a position or velocity tracking problem (Wyeth, 2006; Vallery et al., 2007; Lagoda et al., 2010; Thorson and Caldwell, 2011; Taylor, 2011). Other work shows how steady state tracking and disturbance rejection can be significantly improved through the use of disturbance observers (Kong et al., 2009; Kong et al., 2012; Paine et al., 2014). Building on these approaches, the Valkyrie SEA force control architecture presented in this paper is closely related to previous work on high-performance SEAs (Paine et al., 2014). It differs, however, in that PD feedback is used to shape the dynamic response of the actuator, eliminating the need for an inverse-dynamics feedforward term and improving the controller’s phase margin.

1.3 Summary of Our Approach and Paper Outline

This paper offers new results in two areas. First, we document the actuator control approach used in the field by the Valkyrie DRC robot. We provide details on actuator control performance using quantitative metrics and experimental results. We also demonstrate how these actuator controllers are used in tasks performed by Valkyrie prior to and during the DRC Trials. Second, we present generic methods for controlling the torque output of robots driven by series elastic actuators. These methods include both the conceptual decoupling of multi-joint versus single-joint dynamics and the joint-level series elastic torque controller that makes this conceptual decoupling feasible. In the process of characterizing this torque controller, we introduce a new method for analyzing the disturbance rejection capability of disturbance observers, which we apply to the case of a series elastic actuator with variable load inertia.

We begin by outlining Valkyrie’s multi-joint control approach and the requirements it places on the single-joint torque controllers. We then take a detailed look at the torque control implementation applied to Valkyrie’s various series elastic actuators in Section 3, providing both helpful tuning techniques and quantitative torque control metrics based on single-joint experimental results. Section 4 then uses a SEA plant with finite load inertia to gain insight into the disturbance rejection capabilities of disturbance observer control structures. Given this knowledge, in Section 5 we apply our proposed distributed torque controllers to Valkyrie’s arm, a multi-link chain of SEAs, demonstrating the achieved joint torque tracking performance in a real-world scenario.

2 The Role of Actuator-Level Control in Valkyrie

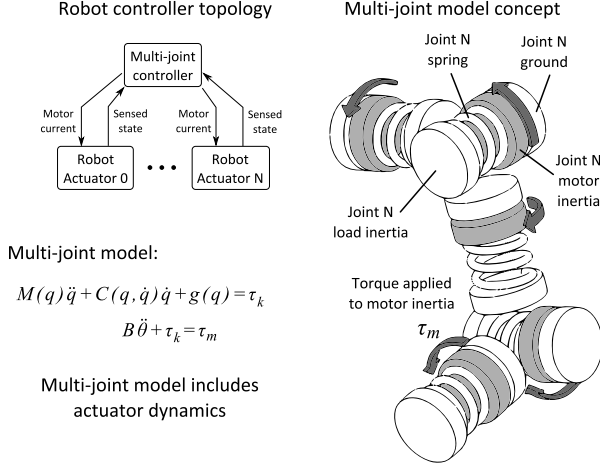
The selection and design of actuator-level control in a series elastic robot depends on the holistic robot control strategy. In this section, we give a brief overview of our holistic decentralized control approach and provide examples of the broad utility provided by joint-level torque controllers.

2.1 A Decentralized Control Approach

Valkyrie uses a decentralized control approach as shown in Figure 2b. This approach was chosen over a centralized approach (Figure 2a) for several reasons. First, using hierarchical control abstractions reduced the complexity of the robot model from a flexible-joint to a rigid-joint model. This simplification allowed our team to leverage existing whole-body control techniques which assumes a rigid-joint model actuated by ideal torque sources (Sentis et al., 2013). Second, a powerful, distributed embedded control element, using NASA-JSCs proprietary ‘Robonet’ high speed serial bus interface and ‘Turbodriver’ motor controllers, is co-located at each joint on Valkyrie. Because of this embedded processing capability, our single-joint controllers are able to utilize highly effective dynamic-model-based control schemes. Third, reducing the coupling between central and peripheral systems decreases overall communication latency due to the reduced number of signals required by the multi-joint controller, meaning they may be updated at a faster rate. Fourth, a decentralized control approach naturally lends itself to an incremental testing methodology, which aides development and debugging.

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

a) Centralized control approach for SEA-driven robots



b) Decentralized control approach for SEA-driven robots

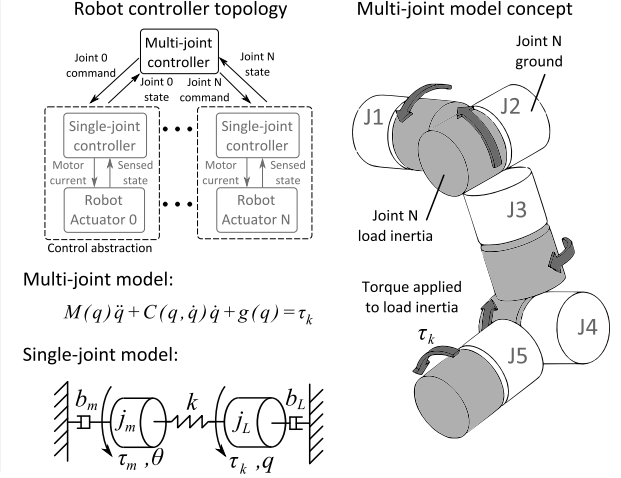
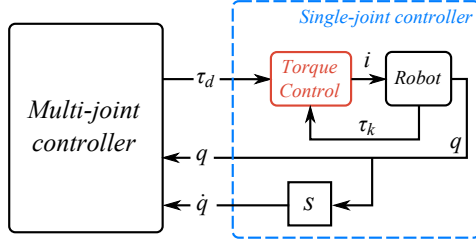


Figure 2: Two different approaches for controlling SEA-driven robots. **a) Centralized Approach.** In this approach, a multi-joint controller models both rigid-body dynamics and actuator dynamics. The multi-joint controller takes joint states as an input and outputs desired motor currents or positions. **b) Decentralized Approach (used on Valkyrie).** In this approach, actuator-level dynamics are abstracted from the multi-joint controller. The multi-joint controller models the robot as rigid bodies actuated by joint torques. The multi-joint controller generates desired joint torques which are then passed to a subordinate set of single-joint controllers. The single-joint controllers model the actuator dynamics and enforce the received torque commands.

a) Joint-level torque control mode



b) Joint-level impedance control mode

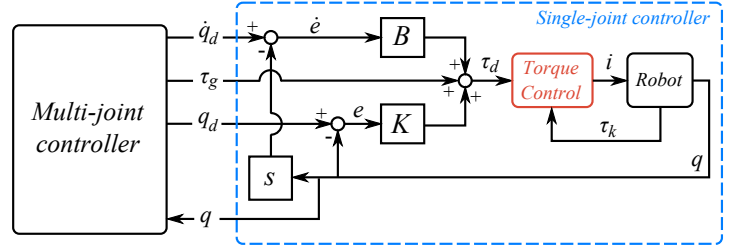


Figure 3: Two actuator control modes supported by Valkyrie. **a) Torque Control Mode.** This joint-level control mode matches the decentralized torque abstraction model shown in Figure 2b where the multi-joint controller sends desired torques to the single-joint controller. **b) Impedance Control Mode.** Desired joint position (q_d), velocity (\dot{q}_d), stiffness (K), damping (B) and gravity compensation torque (τ_g) are sent to the single-joint controller. An impedance control law is enforced at the single-joint level to minimize latency in the control loop, resulting in higher possible stiffness and damping gains. **Note:** Both control modes rely on an accurate torque feedback controller.

2.2 Primary Actuator Control Modes

During the early stages of the project, the control system design of Valkyrie closely matched the decentralized model shown in Figure 2b. The desired joint torque signal created a natural interface between multi-joint and single-joint control domains. Figure 2b represents the robot control system from the multi-joint controller's perspective. If we instead consider a single-joint controller's perspective, the same control system could be represented by Figure 3a. The single-joint controller receives a desired joint torque (τ_d) and uses a torque feedback controller (discussed in Section 3) to enforce this command.

Later in the project, a higher importance was placed on rejecting disturbances caused by multi-joint model uncertainty. As a result, high joint impedance became a new control target. When high impedance controllers were implemented at the multi-joint level, latencies from single- to multi-joint communication would limit the degree to which impedance could be increased before incurring control loop instability. As a solution to this issue and following work by (Pratt et al., 2004), we created another control paradigm at the single-joint level that we dubbed “impedance control mode”, which allowed the position and velocity feedback to occur locally on the single-joint controller (see Figure 3b). This control paradigm increases achievable joint impedance due to a significant reduction of control loop latency compared to those incurred in “torque control mode”.

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

3 Implementation of Torque Feedback Control Using Series Elastic Actuators

As discussed in Section 2, torque control plays an important role in the control of Valkyrie. In this section, we introduce Valkyrie’s series elastic actuators and give details on how they are controlled to accurately track torques.

Valkyrie’s torso, legs, and arms house a combination of rotary and linear SEAs (Figure 4) driven by brushless DC motors¹. The rotary SEAs use harmonic drives, while custom designed torsion springs² act as the compliant element. The spring’s deflection is measured to sense joint torques. The linear SEAs use roller screw drivetrains and commercial-off-the-shelf die springs as the compliant element. The linear actuators have redundant force feedback in the form of sensed deflection of the die springs and load cells in the actuator output linkage. All of these actuators are driven by NASA’s ‘Turbodriver’ motor controllers which perform the joint-level control discussed below.

Figure 5 shows the torque controller implemented in Valkyrie. The control plant is an SEA with a locked output, as shown in Figure 5b. The inner PD compensator is tuned to produce the desired frequency response based on this locked-output assumption. A disturbance observer (DOB) is then used to reject deviations from this nominal locked-output model and maintain torque tracking accuracy. Further discussions on DOB disturbance rejection for non-locked-output scenarios will be discussed in Section 4.

From Figure 5b the relation between torque applied to the spring (τ_m) and spring deflection (θ) is a second order dynamic system with j_m representing effective motor inertia felt by the spring, b_m representing effective motor-side damping felt by the spring, and k representing spring stiffness:

$$\frac{\theta(s)}{\tau_m(s)} = \frac{1}{j_ms^2 + b_ms + k}. \quad (1)$$

Defining the following variables to be: τ_k : spring torque, i : motor current, N : motor speed reduction, k_τ : motor torque constant, η : drivetrain efficiency, we can apply the following relations: $\tau_k = k\theta$ (Hooke’s law) and $\tau_m = i\beta$ where $\beta = Nk_\tau\eta$. The control plant P from motor current to spring torque is then found to be

¹During the DRC Trials 2013, the compliant element of the leg actuators was removed and replaced with rigid material. This change was made several weeks before the DRC in an attempt to improve Valkyrie’s locomotion performance.

²Spring stiffness for each joint was chosen based on a fixed desired spring deflection at peak joint torque.

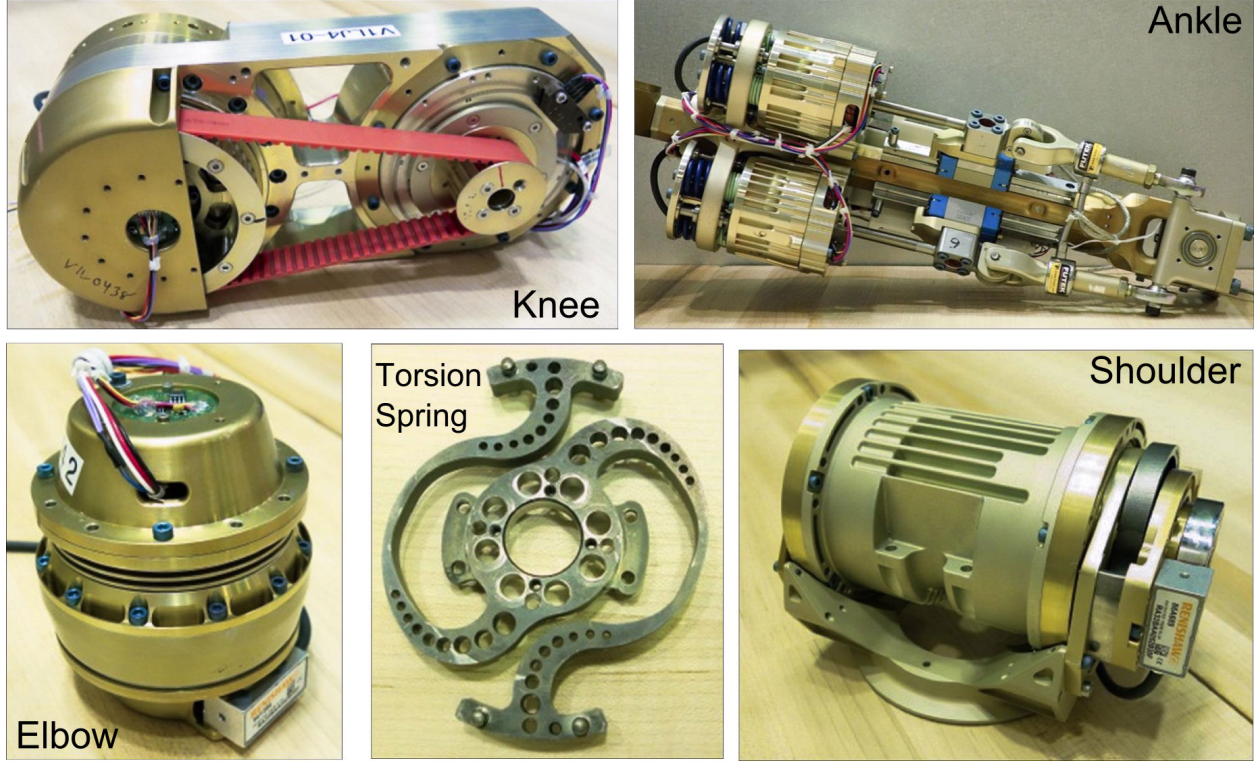
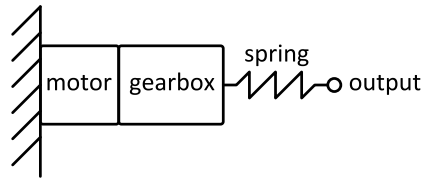
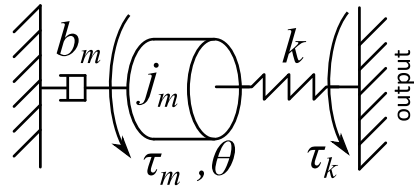


Figure 4: Valkyrie's series elastic actuators. For the rotary actuators, spring deflection corresponding to joint torque is sensed using a Renishaw optical sensor. For the linear actuators, load cells are included in addition to the spring deflection sensor to provide redundant force feedback. The loadcells are placed closer to the joint output giving them better dynamic sensing performance but suffer from a higher noise floor than the spring deflection sensor due to their analog signal properties.

a) Series Elastic Actuator (SEA)



b) Locked output SEA model (plant)



c) Torque control diagram

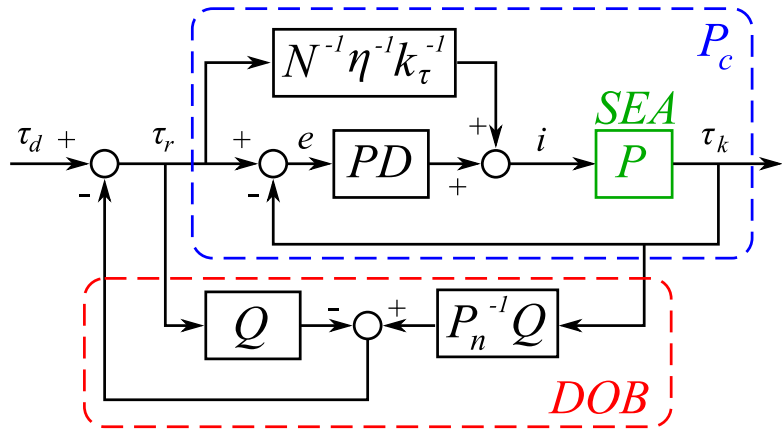


Figure 5: **a)** Schematic representation of a series elastic actuator. **b)** Plant model (P) used for control design. Note that a locked-output assumption is used. **c)** Diagram of the torque feedback controller used on Valkyrie's series elastic actuators. The PD compensator is used to shape the dynamics of the torque response while the DOB is used to improve disturbance rejection, especially at low frequencies. This control structure is based on the controller presented in (Paine et al., 2014).

$$P(s) = \frac{\tau_k(s)}{i(s)} = \frac{k\theta(s)}{\tau_m(s)} \frac{\tau_m(s)}{i(s)} = \frac{Nk_\tau\eta k}{j_ms^2 + b_ms + k} = \frac{\beta k}{j_ms^2 + b_ms + k}. \quad (2)$$

Assuming k is calibrated beforehand, all of the parameters in (2) can be found using system identification techniques with the actuator output locked.

As depicted in Figure 5c, the closed-loop torque-tracking transfer function³ (P_c) is composed of a feedforward term ($N^{-1}\eta^{-1}k_\tau^{-1}$) and a feedback term (PD). The feedforward term is used to scale desired actuator torques into approximate actuator output torques to minimize control effort from the feedback term. The feedback term is represented by the following transfer function:

$$PD(s) = \frac{i(s)}{e(s)} = k_d s + k_p. \quad (3)$$

For simplicity, we do not model the effects of a low-pass filter that is applied to the derivative term. In practice, the cutoff frequency for the $k_d s$ term in (3) is chosen to be sufficiently higher than the actuator's maximum closed-loop system bandwidth (200 Hz versus 70 Hz).

We can then solve for P_c as follows:

$$P_c(s) = \frac{\tau_k(s)}{\tau_r(s)} = \frac{P \cdot B^{-1} + P \cdot PD}{1 + P \cdot PD} = \frac{(k\beta k_d)s + k(1 + \beta k_p)}{j_ms^2 + (b_m + k\beta k_d)s + k(1 + \beta k_p)}. \quad (4)$$

3.1 Simplification of feedback gain selection

When faced with highly parameterized feedback controllers, it is often unclear how each parameter should be chosen. Clearly, a trade-off between poor performance (low gains) and poor stability margin (high gains) must be found. In practice, gain tuning is often performed manually to quickly locate parameters that balance these two trade-offs. From (3) we see that two parameters must be found (k_p and k_d). Instead of choosing k_p and k_d directly, it can be more insightful to consider an equivalent pair of feedback parameters: bandwidth and damping ratio. This is because a desired “shape” of the frequency response can be chosen by selecting a desired damping ratio, and then bandwidth may be chosen independently to satisfy the trade-off between performance and stability. Here, we derive the equations mapping damping ratio to k_p and k_d .

Notice that the characteristic polynomial of (4) is a second order system that can be represented in terms of an effective mass (\hat{M}), spring, (\hat{K}) and damper (\hat{B}):

$$\hat{M} = j_m \quad (5)$$

$$\hat{B} = (b_m + k\beta k_d) \quad (6)$$

$$\hat{K} = k(1 + \beta k_p). \quad (7)$$

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

³Note that in this paper we label the open-loop plant as P while closed-loop transfer functions use the format P_x , x being a unique subscript.

$$\zeta = \frac{\hat{B}}{2\sqrt{\hat{M}\hat{K}}} \quad (8)$$

Combining (5), (6), (7) and (8) forms a relation between desired closed-loop damping ratio (ζ_d) and k_p , k_d :

$$\zeta_d = \frac{(b_m + k\beta k_d)}{2\sqrt{j_m k(1 + \beta k_p)}}. \quad (9)$$

Solving (9) for k_d yields

$$k_d = \frac{2\zeta_d \sqrt{j_m k(1 + \beta k_p)} - b_m}{k\beta}. \quad (10)$$

Using (10), a desired damping ratio can be chosen and k_d can then be automatically calculated for a given k_p . In tuning Valkyrie's torque controller, k_p was used to represent bandwidth, therefore simplifying the tuning of the PD compensator to a single degree of freedom. Using this one parameter, the trade-offs between performance and stability could easily be changed on-the-fly while ensuring a dynamic response with the desired damping ratio. On Valkyrie, we chose a desired damping ratio of 0.9 to produce a flat torque response with little overshoot ($\zeta_d = 1.0$ is critically damped, $\zeta_d = 0.7$ is underdamped with minimum settling time).

3.2 Disturbance Observer

The use of a disturbance observer (DOB) applied to an inner PID/PD control loop has been shown to significantly enhance the torque tracking capability of SEAs (Kong et al., 2012; Paine et al., 2014). DOBs have several useful properties for our specific application. First, they preserve and enforce a dynamic plant model through the use of a nominal model. This means that the shape of the closed-loop frequency response of P_c will not be altered by adding a DOB. The DOB will try to maintain this characteristic response in the presence of either 1) external disturbances or 2) plant model variations. The latter characteristic is of central importance to our approach and is discussed in detail in Section 4. Secondly, DOBs excel at removing steady state error, and therefore effectively serve as an integral feedback term. This characteristic is useful in minimizing controlled torque resolution as discussed in Section 3.3.

A DOB applied to P_c is shown in Figure 5c. P_n^{-1} represents the inverse nominal closed-loop model (using the locked-output constraint). Q is a low-pass filter which is used both to make P_n^{-1} proper and to tune the frequency (f_q) up to which disturbances are rejected. In our implementation, the Q filter takes the form of a second order butterworth filter

$$Q(s) = \frac{1}{(s/\omega_q)^2 + 1.4142(s/\omega_q) + 1} \quad (11)$$

where $\omega_q = 2\pi f_q$.

In the tuning of Valkyrie's disturbance observers, f_q values in the range of 20 Hz to 70 Hz were found to adequately reject unmodeled disturbances while maintaining high control loop stability.

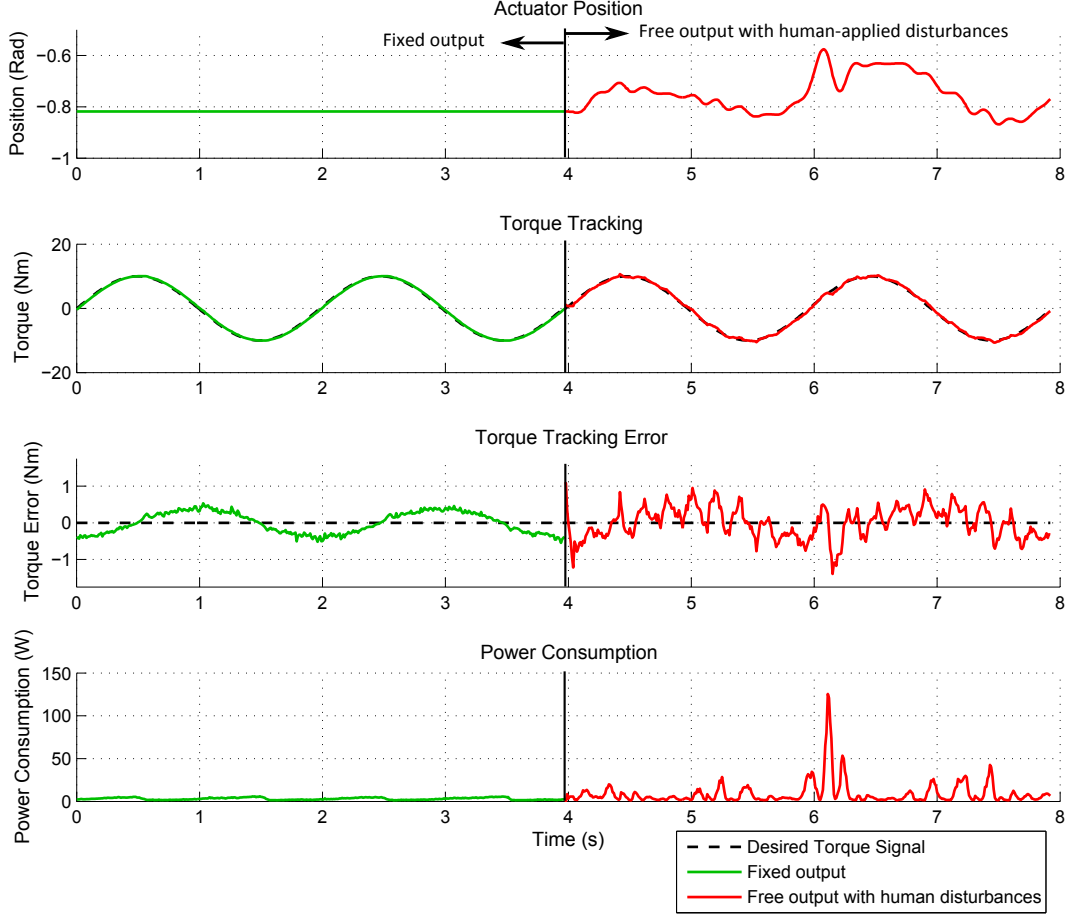


Figure 6: Torque tracking and disturbance rejection of Valkyrie’s torque control implementation. A 10 Nm sinusoidal reference torque is tracked with intentional disturbances (right half) and without intentional disturbances (left half). Tracking error peaks around 1 Nm.

3.3 Quantifying Torque Control Performance

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

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

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

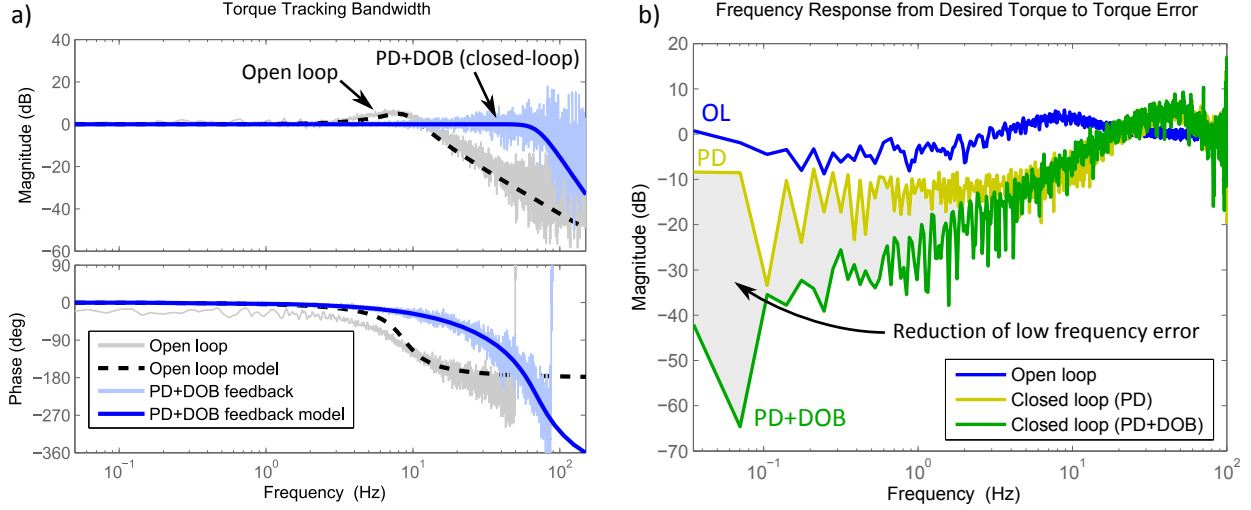


Figure 7: **a)** Frequency responses of a Valkyrie elbow actuator with a fixed output. The dashed line represents the open-loop plant response while the solid blue line represents torque tracking with the closed-loop controller shown in Figure 5c. The closed-loop response significantly increases torque tracking bandwidth and removes the resonant peak seen in the open-loop actuator response. **b)** Torque tracking error versus signal frequency. Adding PD feedback reduces tracking error at frequencies within the actuator bandwidth. The full controller (PD+DOB) further reduces error in low frequencies compared to the PD controller.

As a second measure of torque tracking performance, we performed a test where the actuator was commanded to track a sinusoidal reference torque. We then created disturbances by manually applying loads to the actuator output. By measuring the maximum torque tracking error, we obtained a rough estimate of nominal torque tracking performance. Figure 6 shows these results. As can be seen, torque tracking error remains relatively constant despite the added disturbances. Maximum torque tracking error both with and without disturbances was found to be approximately 1.0 Nm.

Torque bandwidth is an important metric that establishes the maximum signal frequency an actuator is able to accurately track. Figure 7a shows the torque tracking bandwidth of Valkyrie’s elbow actuator using a fixed-output constraint (matching Figure 5b). Using PD feedback, we extend the bandwidth⁴ of the SEA by a factor of 5.3 compared to the passive bandwidth of the SEA (70 Hz compared to 13 Hz). The torque error plot (Figure 7b) illustrates the effect of two different controllers on the torque tracking accuracy. The “PD feedback” line demonstrates the maximum performance a proportional controller can achieve before becoming unstable. The “PD+DOB” line, representing the full torque controller implemented in Valkyrie (that of Figure 5c), clearly illustrates torque tracking improvements in the low frequency range. This improved low-frequency torque tracking benefits the controlled torque resolution metric.

As a final metric, we consider the power consumption of a Valkyrie SEA in torque control mode. Based on the previous discussion of feedback gain selection in Section 3.1, at a first glance, the primary factors in determining feedback gains appear to be performance and stability. A study of actuator energetics quickly demonstrates that a third factor, power consumption, is of critical importance as well. Efficiency is especially important for Valkyrie due to its reliance on battery power.

Figure 8 shows the power consumption of Valkyrie’s elbow actuator for three different scenarios. In each scenario, a chirp signal is generated as a torque reference. The three scenarios are differentiated by the

⁴Bandwidth (BW) of a second order system is defined to be the point in the magnitude response where the value reaches -3 dB. This value can be calculated given the system’s natural frequency (ω_n) and damping ratio (ζ):

$$BW = \omega_n \cdot \sqrt{(1 - 2\zeta^2) + \sqrt{4\zeta^4 - 4\zeta^2 + 2}}.$$

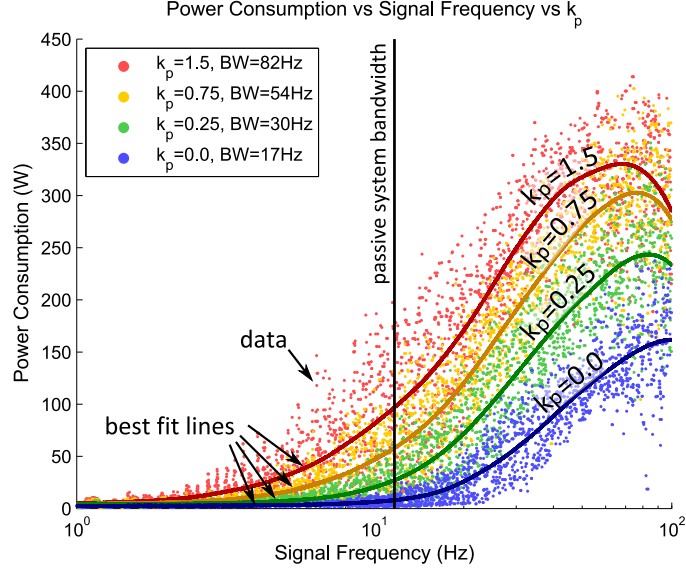


Figure 8: Power consumption of a Valkyrie knee actuator. A desired torque chirp signal was tracked using a PD controller (P_c in Figure 5c) for different values of k_p , thus altering the closed-loop system bandwidth (BW). Power was measured using knowledge of motor current and motor velocity. As is shown, while the actuator is physically capable of tracking torque signals with frequency content above the bandwidth of the passive system (11 Hz), doing so requires large power consumption and thus increases motor heating. Therefore, a balance must be established between torque tracking performance and power consumption to avoid overheating the actuator.

Table 1: Valkyrie Torque Control Performance Metrics

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

selection of k_p in the torque feedback loop. As is shown in Figure 8, torque tracking bandwidth plays a significant role in determining the efficiency of a torque controlled SEA. Increasing torque tracking bandwidth from 30 Hz to 82 Hz results in an average increase of power consumption at the actuator’s passive bandwidth frequency by a factor of 3.6 (97W versus 27W).

These power consumption considerations must be balanced with the needs of the multi-joint controller upstream of the torque controlled SEA. If this centralized controller requires high bandwidth torque tracking, the passive actuator cutoff frequency must be large enough to reduce the discrepancy between passive and active torque bandwidth, as this discrepancy determines peak power consumption (Belanger, 1995). Because the spring of an SEA heavily influences the passive actuator frequency, its selection should consider these energy-related issues. A large amount of work remains to better understand these relations, but the trends observed here are useful in directing future study.

Table 1 summarizes the torque control performance metrics of Valkyrie’s elbow actuator discussed in this section. Valkyrie’s other series elastic joints share a similar performance profile.

4 Decoupling Dynamics of Multi-Joint SEA-Driven Robots

The goal of our torque control approach is to make each SEA appear to the multi-joint controller as an ideal torque source, or at least a low-pass filtered torque source. That is, we want to avoid modeling internal actuator dynamics at the multi-joint level, and instead only model the effects of the rigid body system. Prior work in this area has demonstrated that effective decoupling of the fast actuator-level dynamics from the slower multi-body dynamics is indeed possible (Ott et al., 2003). However, in this work torque errors remain large (around 50Nm) and thus do not adequately abstract an actuator as an ideal torque source.

One option is to employ a decentralized joint-level control approach, which has been shown to work well for rigid-joint position controlled robots (Nakao et al., 1987; Godler et al., 1999). In these approaches subordinate joint-level controllers are co-located at each joint and use feedback to compensate for multi-body dynamics. Interestingly, the authors found that the largest disturbances to affect their joint controllers were caused by variation in the apparent load inertia due to changing robot pose. (Nakao et al., 1987) found that a Disturbance Observer (DOB) could fully reject the undesirable behavior resulting from these model variations that were unknown to their low level controllers. In this section, we study the effectiveness of a similar approach applied to the control of Valkyrie’s actuators. In contrast to the rigid actuator position controllers of (Nakao et al., 1987; Godler et al., 1999), our work focuses instead on the torque-tracking of a SEA in a multi-joint environment. As a result, the control plant models, their variation due to changing load inertia, and the DOB’s ability to cope with this variation all differ from these previous studies.

Outside of the decentralized control community, several methods exist for controlling the torque output of a SEA with no knowledge of the load inertia (Vallery et al., 2007; Kong et al., 2009; Kong et al., 2012; A. Schepelmann and Geyer, 2012). However, their disregard for the magnitude of the load inertia also limits their use; no guidance or condition is provided for the case where load inertia becomes too small and inevitably deteriorates tracking performance or causes instability, as it is known to do (Pratt and Williamson, 1995). In fact, few studies have considered the effects of load inertia on the torque-control performance of SEAs (Wyeth, 2006) and instead usually qualify that their controllers assume “sufficiently large” load inertia. How large is “sufficiently large?” The methods presented in this section can be used to identify this minimum load inertia, or more generally, the maximum tolerable deviation from the control plant’s nominal dynamics, given the DOB-based controller presented previously in Section 3.

We begin the study of a joint-level DOB’s ability to reject the effects of multi-body dynamics by modeling the effects of variable actuator load inertia on the controller presented in Section 3. As previously discussed, the controller in Section 3 is tuned for the case where the load inertia is locked and is not altered thereafter. That is, no adaptive control techniques or gain scheduling based on robot pose is used. Instead, a DOB is used to attenuate disturbances due to control plant variation and maintain the desired closed-loop dynamics. This feature greatly benefits the tuning procedure for the many actuators used in Valkyrie. Each actuator can be separately tuned on a bench with minimal parameter tuning required once the actuator is assembled into a multi-joint system.

4.1 Augmented Control Plant Model

To carry out our analysis, we must augment the locked-output plant model (2) by including the effects of load inertia. The modeling approach we take follows work presented in (Kong et al., 2009) and is extended to map motor torques to spring torques.

A multi-input model for an SEA is shown in Figure 9, with τ_m representing motor torque, θ_m motor angle, j_m motor inertia, b_m motor damping, k spring stiffness, τ_L external torque applied at the load, θ_L load angle, j_L load inertia, and b_L load damping. The state equations for the system are

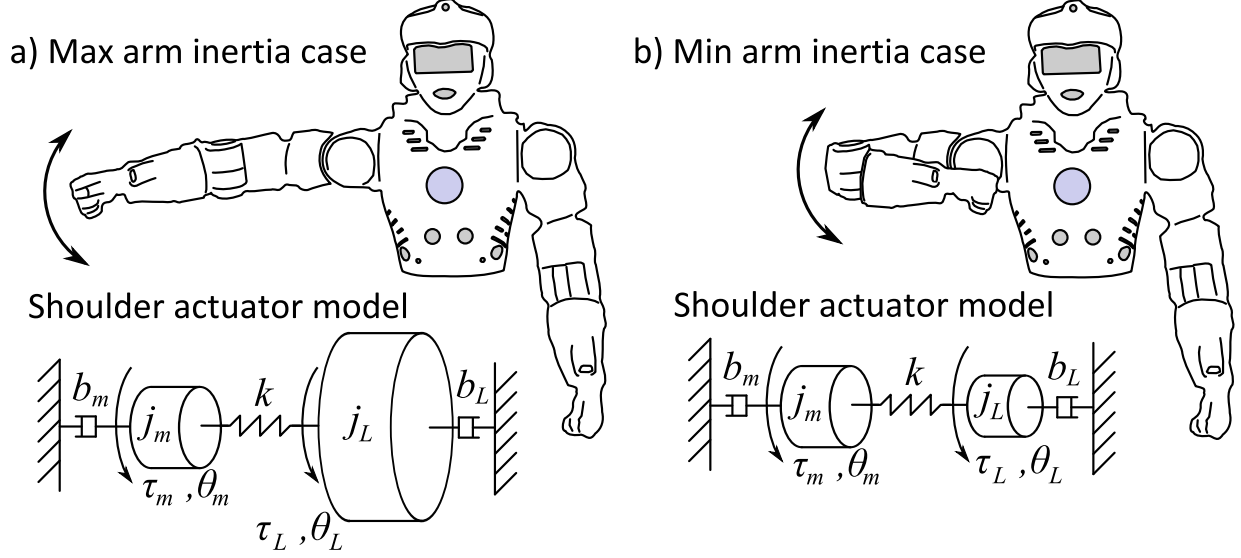


Figure 9: An example of a variable load inertia scenario using the abduction/adduction shoulder joint of Valkyrie. Approximate actuator-level models are shown for each case, a large load inertia corresponding to the out-stretched case and a smaller load inertia corresponding to the bent-elbow case.

$$j_m \ddot{\theta}_m + b_m \dot{\theta}_m + k(\theta_m - \theta_L) = \tau_m \quad (12)$$

$$j_L \ddot{\theta}_L + b_L \dot{\theta}_L + k(\theta_L - \theta_m) = \tau_L. \quad (13)$$

Assuming no external torque input, combining (12) and (13), and representing spring deflection as $\theta_d = \theta_m - \theta_L$, yields a relation (α) between motor angle and spring deflection

$$\alpha(s) = \frac{\theta_d(s)}{\theta_m(s)} = \frac{j_L s^2 + b_L s}{j_L s^2 + b_L s + k}. \quad (14)$$

Using Hooke's law and (14), spring torque (τ_k) is then represented by

$$\tau_k(s) = k\theta_d(s) = k\alpha(s)\theta_m(s). \quad (15)$$

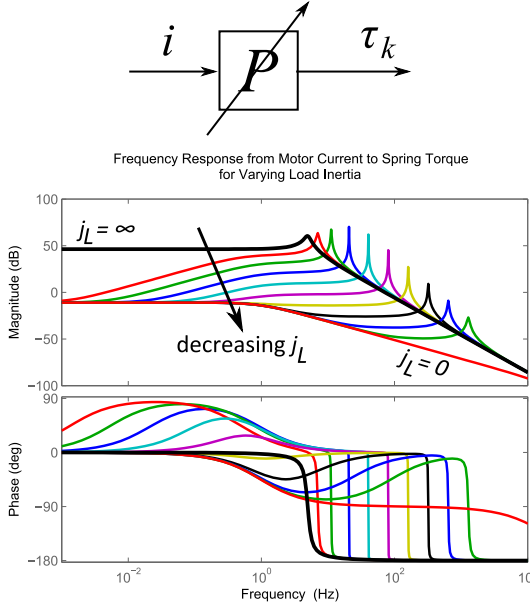
Combining (12) and (15) yields a transfer function from motor torque to motor angle

$$\frac{\theta_m(s)}{\tau_m(s)} = \frac{1}{j_m s^2 + b_m s + \alpha(s)k} \quad (16)$$

which, combined with (14), provides the transfer function from motor current to spring torque for any load inertia

$$P(s) = \frac{\tau_k(s)}{i(s)} = \frac{\tau_m(s)}{i(s)} \frac{\theta_m(s)}{\tau_m(s)} \frac{\theta_d(s)}{\theta_m(s)} k = \frac{\beta \alpha(s) k}{j_m s^2 + b_m s + \alpha(s)k}. \quad (17)$$

a) Open-loop frequency response with variable SEA inertia



b) Closed-loop frequency response with variable SEA inertia

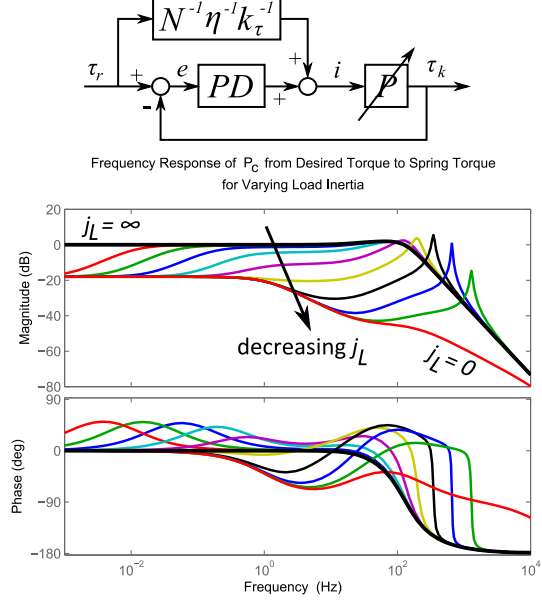


Figure 10: **a) Series of SEA control plants (17) for varying load inertia.** Reducing load inertia is shown to increase resonant frequency and alter effective level of damping in the response. **b) Closed-loop (P_c in Figure 5c) torque response for a series of SEA plants with decreasing j_L .** For larger values of j_L , the bandwidth of the system remains fairly constant while only low frequency behavior changes.

Equation (17) represents a version of the fixed-output control plant model (2) augmented with the effects of finite load inertia.

4.2 Sensitivity Analysis to Variations of Load Inertia

Plotting (17) for a wide range of load inertias provides a visualization and intuition towards understanding how an SEA will respond to different load inertias (see Figure 10a). When j_L is infinite (bold black line), the system exhibits the familiar second order underdamped response, assuming underdamped system parameters. As j_L decreases, low frequency behavior changes noticeably and the resonant peak moves to higher frequencies. In the extreme case, when $j_L = 0$, the load inertia no longer carries energy and the system becomes the first order mass-damper system defined by j_m and b_m .

Similarly, we may simulate the dynamics of the closed-loop response (P_c from Figure 5c) for varying plant load inertias (see Figure 10b). In this case, we see that despite significant variation in load inertia, the high frequency dynamics remain relatively constant. The main variation in the closed-loop response occurs at low frequencies, except for very small values of j_L .

4.3 DOB Disturbance Rejection

Our objective is to leverage the disturbance rejection capability of a DOB to create a controller that is insensitive to variation in load inertia. The result shown in Figure 10b demonstrates the qualitative behavior of finite load inertia on the inner PD compensator, P_c . In this section we seek to understand a DOB's disturbance rejection capability and compare this capability with the disturbances shown in Figure 10b. If the DOB can successfully remove the disturbances of Figure 10b, we will have achieved our goals of creating

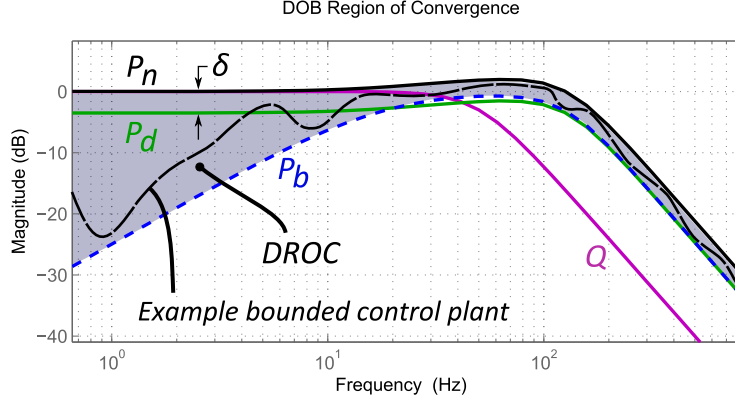


Figure 11: Disturbance observer Region Of Convergence (DROC) for a DOB with a given Q filter. A control plant P_c will converge to an error less than or equal to δ if P_c is contained within the shaded area between P_n and P_b .

an inertia-independent torque controller and thus decentralize torque control of Valkyrie's SEAs.

The structure of a DOB can be seen referring back to Figure 5. Given a nominal model (P_n) of a control plant (P_c), the measured plant output (τ_k) is passed through the nominal model inverse to produce an estimate (τ_e) of the control plant input (τ_r). A lumped disturbance value (d) is calculated by subtracting τ_r from τ_e and is compensated for by subtracting d from the desired value (τ_d). A low pass filter (Q) is used both for tuning purposes and to make the inverse nominal plant model proper.

The transfer function for the DOB shown in Figure 5 is

$$P_{DOB} = \frac{\tau_k}{\tau_d} = \frac{P_c}{1 + Q(P_c P_n^{-1} - 1)} \quad (18)$$

The Q filter in (18) affects the range of disturbances that are rejected by the DOB. Consider the case where $Q = 1$. In such a case, deviations from P_n are rejected at all frequencies, resulting in $P_{DOB} = P_n$. In practice, actuator limitations and sensor noise place a limit on the upper bound of the cutoff frequency for Q . Nonetheless, this basic example demonstrates that large disturbances, or large deviations of the control plant, may be compensated for by using sufficiently high Q filter cutoff frequencies.

In the remainder of this section we analyze the relation between the disturbance rejection capability of a DOB and the variations in closed-loop plant behavior of P_c caused by varying load inertia (Figure 10b). The result of this analysis is a criteria that guarantees bounded tracking error of the DOB. (Kong and Tomizuka, 2013) provides a similar criteria for determining stability bounds of DOB-based systems with multiplicative uncertainties, but does not give the same error bound guarantee presented here.

The underlying idea behind our method of characterizing DOB disturbance rejection is as follows (see Figures 5 and 11). A DOB attenuates deviations of the control plant, P_c , from some nominal model, P_n . The discrepancy between P_n and the DOB-compensated system (P_{DOB}) is the tracking error of the full closed-loop system. Because the DOB only *attenuates* deviations, tracking error can never be reduced to zero if finite deviations of P_c versus P_n exist. However, if a minimum acceptable tracking error is specified (δ), the worst-case deviation of the control plant from the nominal plant that satisfies the tracking error can be found (P_b). Therefore, any P_c contained within the area bounded by the best case scenario (P_n , zero deviation) and the worst case scenario (P_b , δ deviation) satisfies the minimum tracking error requirement. We refer to the area between P_n and P_b as the Disturbance observer Region Of Convergence (DROC).

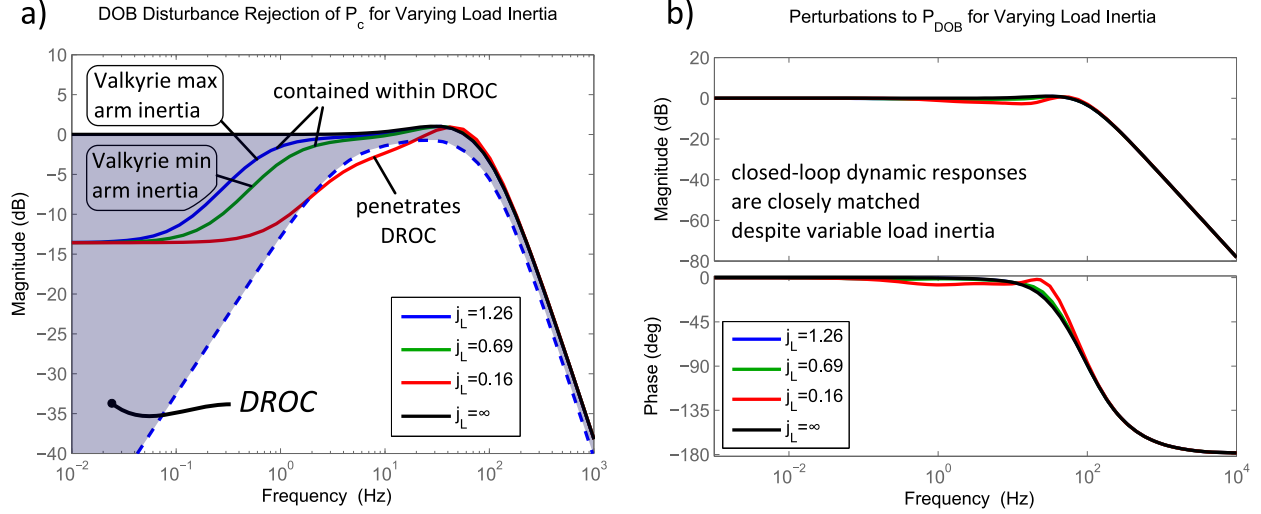


Figure 12: These figures demonstrate how a disturbance observer can eliminate the effects of plant deviations caused by changes to actuator load inertia. **a)** Closed-loop frequency response of P_c for different load inertias. The dashed line represents 20% error (< 2 dB) compared to the nominal high load inertia model while the shaded area represents the DOB region of convergence. Both the maximum and minimum arm inertia lines are contained within the DROC, meaning a DOB can be used to attenuate the disturbance due to these effects. **b)** Closed-loop frequency response from desired torque (τ_d) to measured spring torque (τ_k) for the full torque controller as shown in Figure 5c with a finite actuator load inertia. A series of responses to plants with varying inertias is shown representing the full inertial operating range of the Valkyrie shoulder abduction/adduction actuator ($j_L = 1.26 \rightarrow 0.69 \text{ Kg} \cdot \text{m}^2$). Note the insensitivity of the system to load inertia variation.

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

$$P_d = P_n(1 - \delta) \quad (19)$$

where P_d is the desired maximum allowable deviation from the nominal plant. Setting $P_{DOB} = P_d$ and solving (18) for P_c yields

$$P_b = P_c = \frac{P_d(1 - Q)}{1 - QP_dP_n^{-1}}. \quad (20)$$

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

4.4 DOB Rejection of Disturbances Due to Variable Load Inertia

Analyzing the results shown in Figure 10b using our DROC method gauges how a DOB will reject disturbances due to load inertia variation and thus maintain the desired ideal torque source abstraction. In this analysis, we model a Valkyrie shoulder actuator for the two load inertia cases shown in Figure 9 using the control parameters shown in Table 2. We also simulate a smaller load inertia ($0.16 \text{ Kg} \cdot \text{m}^2$), which is included to demonstrate a case which violates the DROC. Figures 12 and 13 show these results.

Figure 12a illustrates how the closed-loop responses for both values of arm inertia remain within the DROC.

Table 2: Variable Inertia Simulation Parameters

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

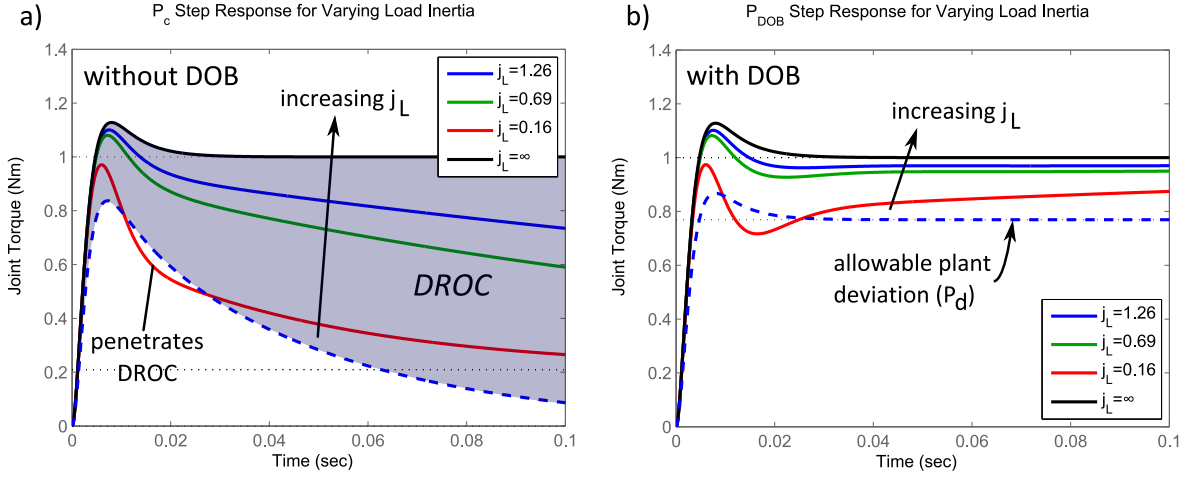


Figure 13: Time domain representations of Figures 12a and 12b. In **a)** the step response of only the inner portion of the proposed force controller (P_c) to various load inertias is shown. Again, for these control parameters, all values of load inertia except $j_L = 0.16 \text{ Kg} \cdot \text{m}^2$ remain within the DROC. In **b)** step responses are shown for the same SEA plant parameters as in **a)**, but a DOB is now applied. The DOB is shown to dramatically improve low-frequency performance, maintaining the full range of Valkyrie’s arm inertia within the allowable plant deviation.

Because both responses remain within the DROC, a DOB applied to either closed-loop response is guaranteed to bring the system response to within δ ($< 2 \text{ dB}$ in this case) of the nominal fixed-output plant model. We can visualize this result in Figure 12b where the perturbations to the full system transfer functions due to load inertia variation are imperceptible for the two arm inertia cases and only become perceptible in the scenario where $j_L = 0.16 \text{ Kg} \cdot \text{m}^2$. These results also show that the stability of the system is minimally influenced over the full range of Valkyrie’s arm inertia due to the relationship between the system bode plot and stability (gain and phase) margins. Because of this relationship, the DROC is also a useful tool for assessing the variation of stability for systems with DOBs.

Figure 13 shows the data of Figures 12a and 12b in the form of time domain step responses. Here, it is again shown how the DOB is able to compensate for large deviations of the control plant, provided these deviations primarily occur at frequencies below the Q filter cutoff frequency.

5 Experimental Results on Valkyrie

In this section, we put the disturbance rejection capability of the DOB to the test in the form of two experiments using Valkyrie’s arm, a serial chain of four SEAs. From Section 4 we have learned that the DOB is capable of rejecting disturbances at frequencies below its Q filter cutoff frequency. In this section, the complexity of disturbances experienced by the DOB are increased beyond the variable inertia case considered in Section 4.4. Now, the controller must attenuate disturbances from the full dynamics of a serial chain of

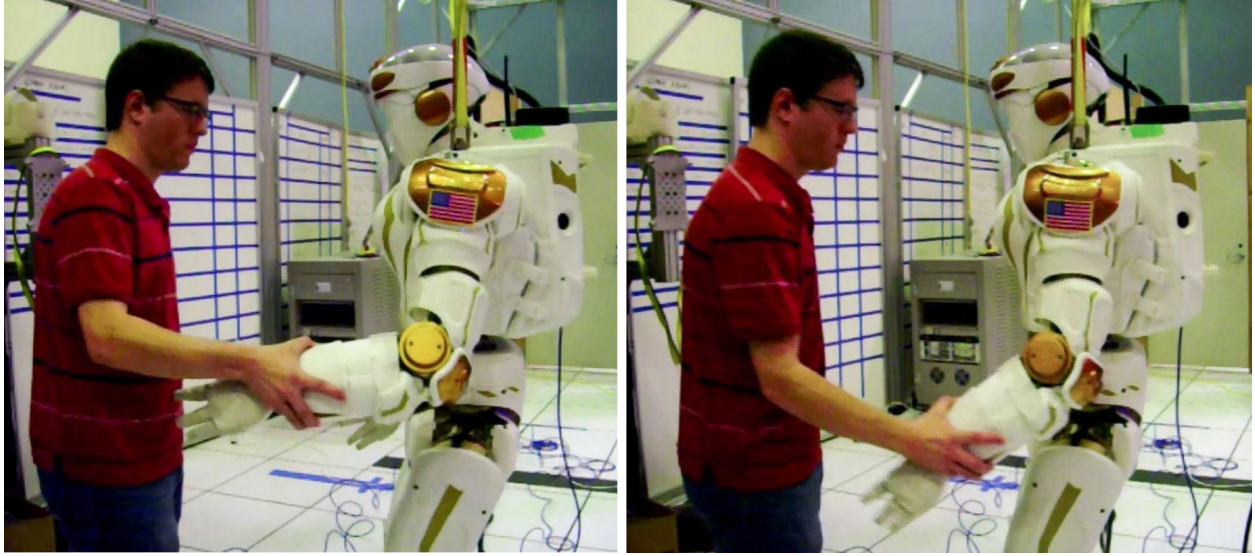


Figure 14: **Experiment 1: Human Interaction.** Valkyrie arm joints J1 through J4 are placed in torque control mode and are actively tracking a constant desired torque signal. A human applies motions to the arm causing all four joints to move.

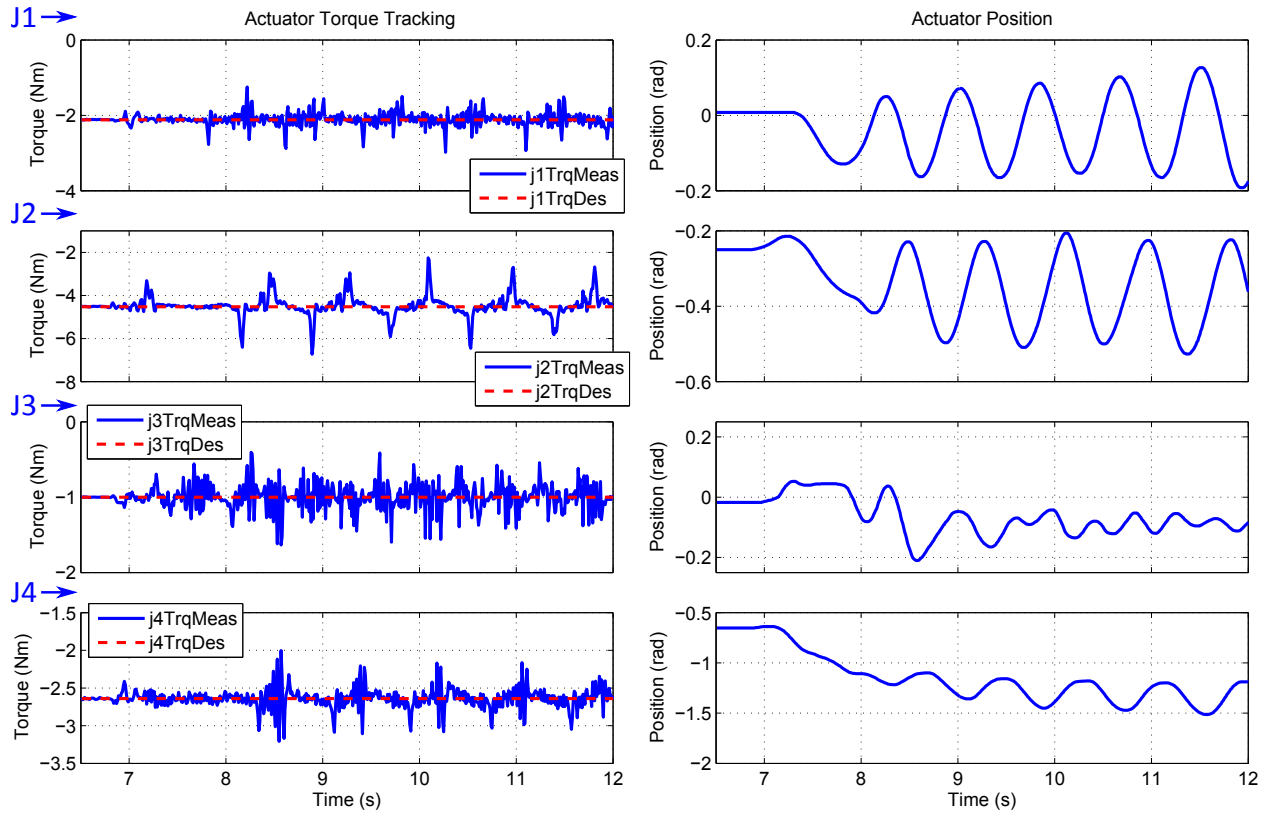


Figure 15: **Data from Experiment 1: Human Interaction.** Joint motion is caused by the human who is interacting with Valkyrie's arm. The peak tracking error for each joint is: J1: 0.85 Nm, J2: 2.2 Nm, J3: 0.63 Nm, J4: 0.64 Nm.

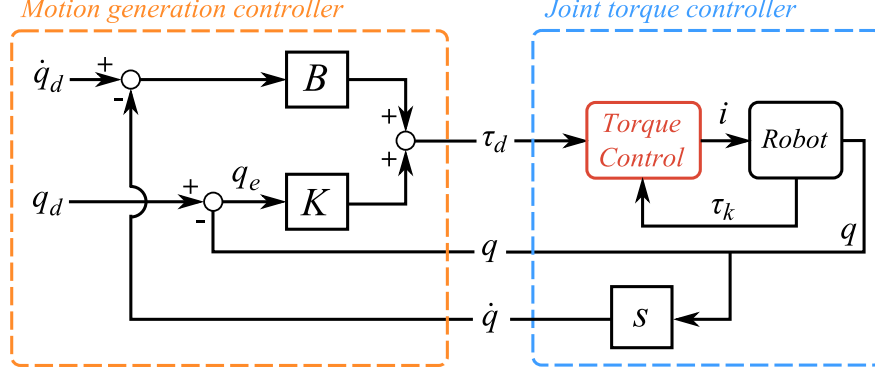


Figure 16: Motion generation controller used to test joint torque control in Experiment 2. The parameter K represents joint stiffness while B represents joint damping.

Table 3: Valkyrie Control Parameters used in Experiments 1 and 2

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

four SEAs along with the disturbances introduced by control action of each of their motors.

The goal of the two experiments shown here is to assess the torque tracking capability of the distributed joint-level torque controllers, which possess no knowledge of neighboring joints. The only models used in these tests are the nominal high-output-impedance models used by the DOB as discussed in Section 3. The control parameters used in these experiments are listed in Table 3.

In the first experiment, joints one through four⁵ are placed into torque control mode and are commanded to track a constant torque value. A human then grasps Valkyrie’s forearm and applies motions such that all four joints move (see Figure 14). Figure 15 shows the data from this experiment. Low torque tracking errors are maintained, despite the interaction forces and motions from the human.

In the second experiment, joints one through four are again placed into torque control mode. In this experiment, however, desired joint torques are generated by the controller shown in Figure 16. This is a simple proportional-derivative controller which creates an apparent joint stiffness (K) and damping (B) based on a desired joint position (q_d) and velocity (\dot{q}_d). We use this controller to generate motion in joints one through four; we are less concerned with the position tracking error (q_e) in this particular experiment. Figure 17 shows the basic arm motion generated during Experiment 2. Correspondingly, Figure 18 shows the data from Experiment 2.

These two experiments demonstrate that the disturbance attenuation properties of the joint-torque controller are able to suppress the dynamics of neighboring joints, producing accurate torque tracking. However, by visual inspection of Figures 15 and 18 it is difficult to evaluate the achieved tracking performance based on torque error magnitude alone. A helpful metric in this scenario is the maximum torque tracking error relative to the rated maximum joint torque. Using this metric, we find that the torque error relative to the rated joint torque of each joint in Experiment 1 is: J1: 0.4%, J2: 0.76%, J3: 0.97%, and J4: 0.98%. For

⁵Refer back to Figure 2b for joint naming conventions.

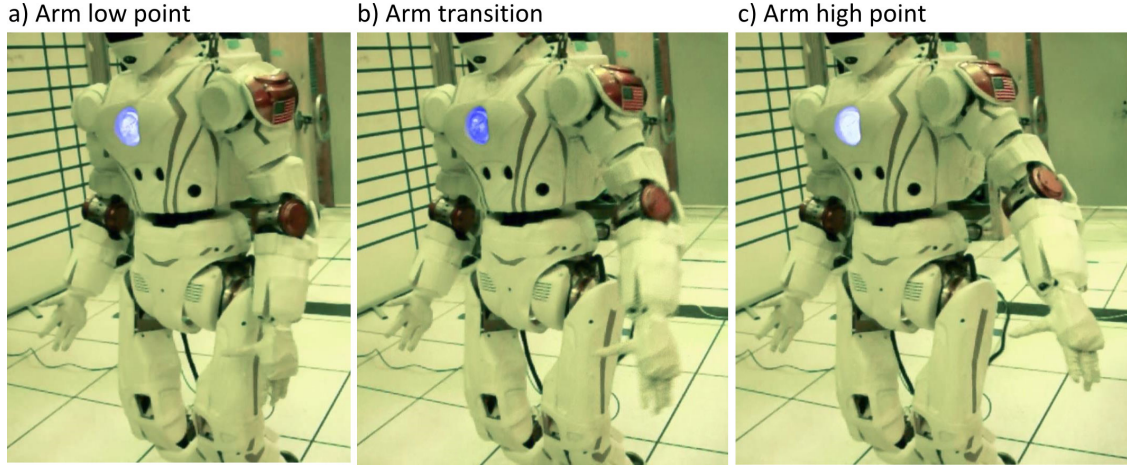


Figure 17: **Experiment 2: Coordinated Motion.** Valkyrie arm joints J1 through J4 are placed in torque control mode and are supplied desired torque signals according to the controller shown in Figure 16. As a result, each joint roughly tracks a sinusoidal position.

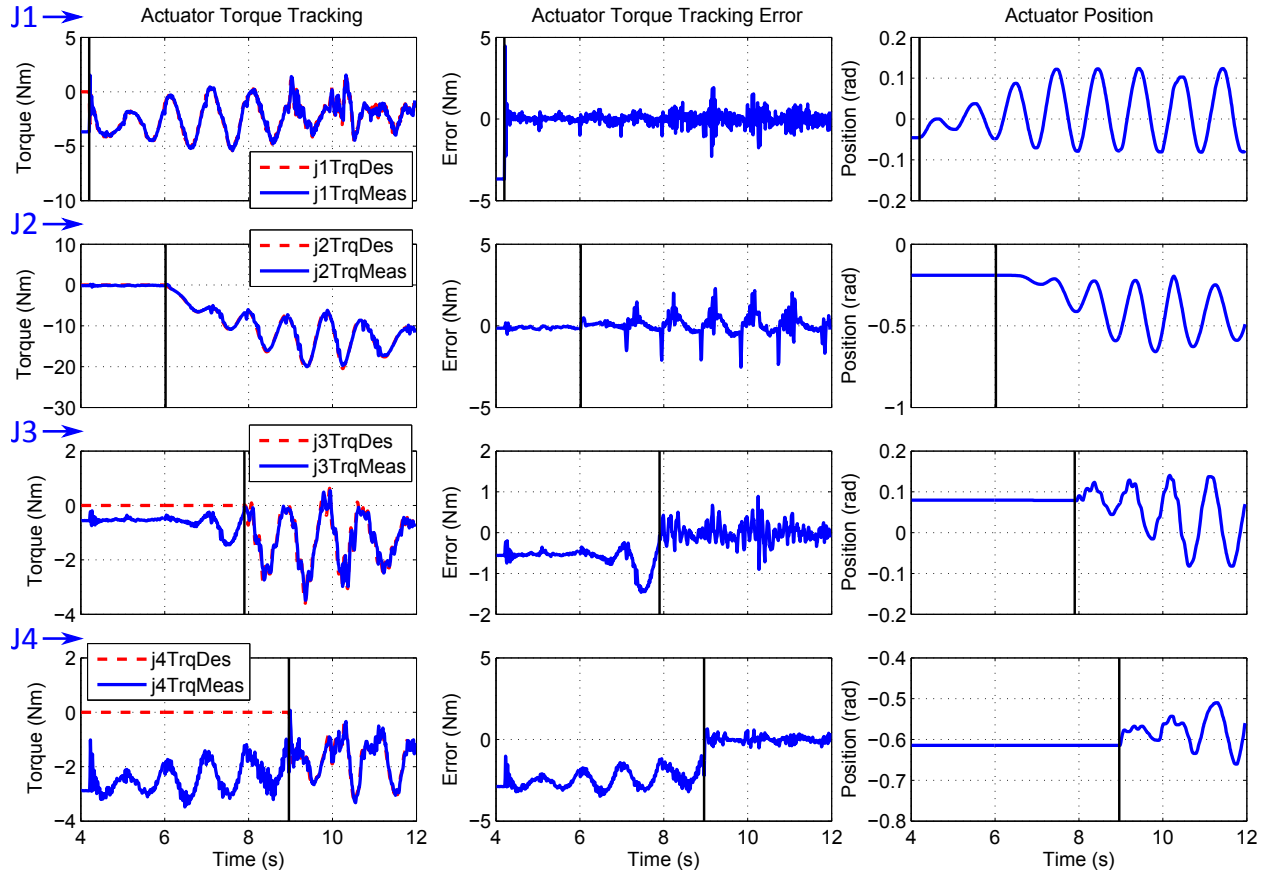


Figure 18: **Data from Experiment 2: Coordinated Motion.** The purpose of this experiment is to demonstrate the torque tracking accuracy of a serial chain of SEAs controlled using the methods presented in this paper. The vertical black lines in the figure represent the time where each joint begins tracking the desired torque signal. The peak tracking error for each joint is: J1: 2.2 Nm, J2: 2.5 Nm, J3: 0.9 Nm, J4: 0.65 Nm.

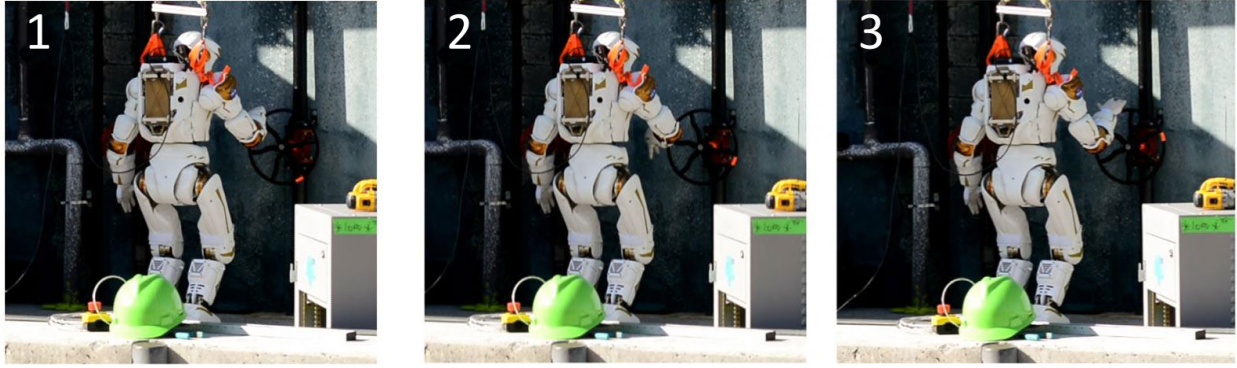


Figure 19: **Valve Turn at the DRC Trials.** Valkyrie using the actuator control methods described in this paper to turn a valve during the DRC Trials, December 2013 in Miami, Florida, USA.

Experiment 2 these numbers are: J1: 1.16%, J2: 0.87%, J3: 1.38%, and J4: 1%. When considering that the accuracy of many *sensors* fair no better than this in terms of full-scale accuracy⁶, we consider our results to be quite strong. Based on this metric, our results also fare well compared to leading research in the field of accurate torque control of SEAs with unmodeled disturbances. To date and to the authors' best knowledge, the leading results in literature achieve tracking accuracies of 1.6% of full scale (Kong et al., 2009) and 15% of full scale (Kong et al., 2012).

The proposed actuator control methodology has also been used by Valkyrie to perform useful tasks in the field, such as the valve turn task during the DRC Trials 2013 (Figure 19).

6 Conclusions and Discussion

The challenge tasks created for the DARPA Robotics Challenge set an unprecedented bar for the application of advanced robotics. In response to this challenge, the NASA-JSC Valkyrie team designed a unique humanoid robot. The combination of human-inspired dexterity, physical compliance, and an impressive suite of sensory modes distinguished Valkyrie from the other DRC entries.

Valkyrie's use of series elastic actuators inspired adoption of a decentralized torque control architecture. To realize such an architecture, excellent disturbance rejection was required from the joint-level controllers to provide actuators that can be abstracted as near-ideal torque sources, regardless of kinematic configuration. In addition to proposing control methods aimed at this goal, we have provided a tool, the DROC, which may be used to indicate in which conditions a DOB-controlled SEA meets this "near-ideal torque source" specification.

We characterized Valkyrie's SEA torque control performance with several empirical tests. In the first set of tests, we characterized performance of a single actuator, obtaining measures of controllable torque resolution (0.002 Nm), nominal torque tracking error (1 Nm), and torque tracking bandwidth (up to 70 Hz). These numbers fare well in comparison to other torque-controlled SEAs which achieve bandwidths of 5-25 Hz (Pratt and Williamson, 1995), 16 Hz (Vallery et al., 2007), 10 Hz (Kong et al., 2009) and 19 Hz (Sensinger and Weir, 2005). In addition, actuator power consumption was considered. We demonstrated its critical relation to torque tracking bandwidth, where a 30 Hz increase in bandwidth produced a 3.6x increase in power consumption. The performance of our control approach was also shown in two tests using Valkyrie's arm, a serial chain of SEAs. In these tests, Valkyrie's actuators accurately tracked torques to within 1.38% of their rated torque values, an improvement compared to leading results in literature (1.6% and 15%) (Kong et al.,

⁶<http://www.futek.com/files/pdf/Product%20Drawings/lcm300.pdf>

2009; Kong et al., 2012).

Several unanswered questions remain for future study. First, a side effect of nearly-ideal torque source actuators is a significant reduction in effective joint friction. This low-friction environment leads to a heavy reliance on rigid body model accuracy. Accurate gravity compensation, for instance, is crucial to employ in such systems as the torque-controlled actuators are easily backdriven by the weight of the robot. Therefore, methods for validating and improving the rigid body model in hardware should be adopted when applying the near-ideal torque source approach presented here. Second, when the rigid body model is inaccurate, as was the case during much of Valkyrie’s development, controllers must be employed that can compensate for these model discrepancies. Such controllers must rely more on position feedback effort than on the feedforward effort produced from the model-based portion of the multi-joint controller (such as gravity compensation). Maximizing the output impedance of these controllers^{7,8} can be beneficial when there are rigid body model discrepancies, a topic not well understood for the case of SEAs. This area of research deserves further study.

Acknowledgments

The Valkyrie robot was the product of many people’s efforts, most of whom are not listed as authors of this paper. The work presented in this paper represents only a small portion of the efforts undertaken by the Valkyrie team as a whole.

References

- A. Schepelmann, M. T. and Geyer, H. (2012). Development of a testbed for robotic neuromuscular controllers. In *Proceedings of Robotics: Science and Systems*.
- Albu-Schaffer, A., Ott, C., and Hirzinger, G. (2007). A unified passivity-based control framework for position, torque and impedance control of flexible joint robots. *Int. J. Robot. Res.*, 26(1):23–29.
- Ambrose, R. O., Burridge, R. R., Bluethmann, W., Aldridge, H., Askew, R. S., Rehnmark, F., Diftler, M., Magruder, D., and Lovchik, C. (2000). Robonaut: NASA’s space humanoid. *IEEE Intelligent Systems*, 15(4):57–63.
- Bélanger, P. R. (1995). *Control Engineering: A Modern Approach*. Saunders College Publishing, Orlando, Florida.
- Bluethmann, W., Ambrose, R., Diftler, M., Askew, S., Huber, E., Goza, M., Rehnmark, F., Lovchik, C., and Magruder, D. (2003). Robonaut: A robot designed to work with humans in space. *Autonomous Robots*, 14(2-3):179–197.
- Bridgwater, L. B., Ihrke, C., Diftler, M. A., Abdallah, M. E., Radford, N. A., Rogers, J., Yayathi, S., Askew, R. S., and Linn, D. M. (2012). The robonaut 2 hand-designed to do work with tools. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3425–3430. IEEE.
- Diftler, M. A., Mehling, J., Abdallah, M. E., Radford, N. A., Bridgwater, L. B., Sanders, A. M., Askew, R. S., Linn, D. M., Yamokoski, J. D., Permenter, F., et al. (2011). Robonaut 2-the first humanoid robot in space. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pages 2178–2183. IEEE.
- Garcia, E., Arevalo, J., Sanchez, F., Sarria, J., and Gonzalez-de Santos, P. (2011). Design and development of a biomimetic leg using hybrid actuators. In *Intelligent Robots and Systems (IROS), IEEE/RSJ International Conference on*, pages 1507 –1512.

⁷Perhaps even beyond the bounds of passivity (Ott et al., 2008)

⁸The “impedance mode” in Figure 3 was designed for this purpose

- Godler, I., Inoue, M., Ninomiya, T., and Yamashita, T. (1999). Robustness comparison of control schemes with disturbance observer and with acceleration control loop. In *Industrial Electronics, ISIE Proceedings of the IEEE International Symposium on*, volume 3, pages 1035–1040 vol.3.
- Hurst, J., Chestnutt, J., and Rizzi, A. (2010). The actuator with mechanically adjustable series compliance. *Robotics, IEEE Transactions on*, 26(4):597–606.
- Hutter, M., Remy, C., Hoepflinger, M., and Siegwart, R. (2013). Efficient and versatile locomotion with highly compliant legs. *Mechatronics, IEEE/ASME Transactions on*, 18(2):449–458.
- Kim, J., Kwak, H., Lee, H., Seo, K., Lim, B., Lee, M., Lee, J., and Roh, K. (2012). Balancing control of biped robot. In *IEEE international Conference on Systems, Man, and Cybernetics*.
- Kong, K., Bae, J., and Tomizuka, M. (2009). Control of rotary series elastic actuator for ideal force-mode actuation in human-robot interaction applications. *Mechatronics, IEEE/ASME Transactions on*, 14(1):105–118.
- Kong, K., Bae, J., and Tomizuka, M. (2012). A compact rotary series elastic actuator for human assistive systems. *Mechatronics, IEEE/ASME Transactions on*, 17(2):288–297.
- Kong, K. and Tomizuka, M. (2013). Nominal model manipulation for enhancement of stability robustness for disturbance observer-based control systems. *International Journal of Control, Automation, and Systems*, 11(1):12–20.
- Lagoda, C., Schouten, A., Stienen, A., Hekman, E., and van der Kooij, H. (2010). Design of an electric series elastic actuated joint for robotic gait rehabilitation training. In *Biomedical Robotics and Biomechatronics (BioRob), 3rd IEEE RAS and EMBS International Conference on*, pages 21–26.
- Li, Z., Tsagarakis, N., and Caldwell, D. (2012). A passivity based admittance control for stabilizing the compliant humanoid coman. In *Humanoid Robots (Humanoids), 12th IEEE-RAS International Conference on*, pages 43–49.
- Nagatani, K., Kiribayashi, S., Okada, Y., Otake, K., Yoshida, K., Tadokoro, S., Nishimura, T., Yoshida, T., Koyanagi, E., Fukushima, M., and Kawatsuma, S. (2013). Emergency response to the nuclear accident at the Fukushima Daiichi Nuclear Power Plants using mobile rescue robots. *Journal of Field Robotics*, 30(1):44–63.
- Nakao, M., Ohnishi, K., and Miyachi, K. (1987). A robust decentralized joint control based on interference estimation. In *Robotics and Automation. Proceedings. IEEE International Conference on*, volume 4, pages 326–331.
- Ott, C., Albu-Schaffer, A., Kugi, A., and Hirzinger, G. (2003). Decoupling based cartesian impedance control of flexible joint robots. In *Robotics and Automation, Proceedings. IEEE International Conference on*, volume 3, pages 3101–3107 vol.3.
- Ott, C., Kugi, A., and Hirzinger, G. (2008). On the passivity-based impedance control of flexible joint robots. *Robotics, IEEE Transactions on*, 24(2):416–429.
- Paine, N., Oh, S., and Sentis, L. (2014). Design and control considerations for high-performance series elastic actuators. *Mechatronics, IEEE/ASME Transactions on*, 19(3):1080–1091.
- Paluska, D. and Herr, H. (2006). Series elasticity and actuator power output. In *Robotics and Automation, ICRA Proceedings IEEE International Conference on*, pages 1830–1833.
- Pratt, G. and Williamson, M. (1995). Series elastic actuators. In *Intelligent Robots and Systems. 'Human Robot Interaction and Cooperative Robots', Proceedings. IEEE/RSJ International Conference on*, volume 1, pages 399–406 vol.1.

- Pratt, G., Willisson, P., Bolton, C., and Hofman, A. (2004). Late motor processing in low-impedance robots: impedance control of series-elastic actuators. In *American Control Conference, Proceedings of the*, volume 4, pages 3245–3251 vol.4.
- Pratt, J., Krupp, B., and Morse, C. (2002). Series elastic actuators for high fidelity force control. *Industrial Robot: An International Journal*, 29(3):234 – 241.
- Ragonesi, D., Agrawal, S., Sample, W., and Rahman, T. (2011). Series elastic actuator control of a powered exoskeleton. In *Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE*, pages 3515–3518.
- Reiland, M. J., Hargrave, B., Platt, R., Abdallah, M. E., and Permenter, F. N. (2013). Architecture for robust force and impedance control of series elastic actuators. US Patent 8,525,460.
- Robinson, D. W. (2000). *Design and Analysis of Series Elasticity in Closed-loop Actuator Force Control*. PhD thesis, MASSACHUSETTS INSTITUTE OF TECHNOLOGY.
- Sensingier, J. W. and Weir, R. F. (2005). Design and analysis of a non-backdrivable series elastic actuator. In *Rehabilitation Robotics, ICORR. 9th International Conference on*, pages 390–393.
- Sensingier, J. W. and Weir, R. F. (2006). Unconstrained impedance control using a compact series elastic actuator. In *Mechatronic and Embedded Systems and Applications, Proceedings of the 2nd IEEE/ASME International Conference on*, pages 1–6.
- Sentis, L., Petersen, J., and Philippsen, R. (2013). Implementation and stability analysis of prioritized whole-body compliant controllers on a wheeled humanoid robot in uneven terrains. *Autonomous Robots*, 35(4):301–319.
- Spong, M. (1987). Modeling and control of elastic joint robots. *J. Dyn. Sys., Meas., Control.*, 109(4):310–318.
- Taylor, M. D. (2011). A compact series elastic actuator for bipedal robots with human-like dynamic performance. Masters thesis, Carnegie Mellon University.
- Thorson, I. and Caldwell, D. (2011). A nonlinear series elastic actuator for highly dynamic motions. In *Intelligent Robots and Systems (IROS), IEEE/RSJ International Conference on*, pages 390–394.
- Vallery, H., Ekkelenkamp, R., van der Kooij, H., and Buss, M. (2007). Passive and accurate torque control of series elastic actuators. In *Intelligent Robots and Systems, IROS IEEE/RSJ International Conference on*, pages 3534–3538.
- Wyeth, G. (2006). Control issues for velocity sourced series elastic actuators. In *Australian Conference on Robotics and Automation*.