

PASSIVE CONTROL AND OPTIMAL REGENERATION IN DC MOTORS

By

Léo Gustavo Vailati

Dissertation

Submitted to the Faculty of the
Graduate School of Vanderbilt University
in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Electrical Engineering

August 11, 2023

Nashville, Tennessee

Approved:

Michael Goldfarb, Ph.D.

David Braun, Ph.D

Eric Barth, Ph.D.

Nilanjan Sarkar, Ph.D

Richard Alan Peters, Ph.D

Table of Contents

List of Tables	iv
List of Figures	v
1 Introduction	1
1.1 Chapter 2: Fault Monitoring in Pneumatic Systems	1
1.2 Chapters 3 and 4: Controlling Brushless Motors to Implement Passive Braking Behaviors.....	3
1.3 Chapter 5: Optimizing Energy Regeneration in a Microprocessor-Controlled Knee Prosthesis...	5
2 A Method for Mass Flow and Displacement Estimation in a Pneumatic Actuation System using Valve-based Pressure Sensing.....	11
2.1 Introduction	11
2.2 Adapted Mass Flow Model	12
2.3 Experimental Validation.....	17
2.4 Results.....	23
2.5 Discussion.....	27
2.6 Conclusions	29
3 On Using a Brushless DC Motor as a Passive Torque-Controllable Brake	30
3.1 Introduction	30
3.2 Modeling Motor Behavior.....	31
3.3 Controlling a Motor as a Brake	33
3.4 Other Control Considerations	39
3.5 Experimental Implementation	42
3.6 Discussion.....	45
3.7 Conclusion.....	46
4 On the Benefits and Limitations of Modulated Damping with Passive Motor Control	47
4.1 Introduction	47
4.2 Passive and Conventional BLDC Motor Control.....	48
4.3 Application of passive switching to damping emulation	51
4.4 Implementation	55
4.5 Experimental Results and Discussion.....	59
4.6 Discussion.....	63
4.7 Conclusion.....	65
5 A method for optimized energy regeneration in a knee prosthesis.....	66
5.1 Introduction	66

5.2	Prosthesis Hardware	68
5.3	Optimizing Energy Regeneration	71
5.4	Experimental Evaluation	79
5.5	Discussion of Results	86
5.6	Conclusion	88
5.7	Appendix I: Optimizing Energy Regeneration in Dynamical Systems	88
5.8	Appendix II: Applying the optimal regeneration framework to the prosthetic system	93
6	Additional Contributions	96
6.1	Controllers for Powered Knee Prosthesis	96
6.2	Understanding Stumble Recovery Strategy Selection	96
7	References	98

List of Tables

Table 2.1. Instrumentation for Validation Experiments	17
Table 2.2. Polynomial Fit Parameters	20
Table 2.3. Sigmoid surface parameters	21
Table 2.4. Error in estimation of mass use.....	25
Table 2.5. Error in estimation of rod position.....	26
Table 3.1. Forced and unforced behaviors for the switching model	35
Table 3.2. Parameter values	36
Table 4.1. Parameters of the experimental apparatus.....	56
Table 5.1. Association of ECT hardware configuration with phase of walking.....	69
Table 5.2. Parameter values for ECT friction, transmission ratio, and the motor	71
Table 5.3. Parameter values for the optimized control laws.....	80
Table 5.4. Values of optimization hyperparameters for each phase of swing.	94

List of Figures

Fig. 1.1. Valve with pressure sensor and pressure sensor assembly	2
Fig. 1.2. Comparing motors and brakes in	4
Fig. 1.3. Angle, angular velocity, torque, and power of the knee joint in level walking at 0.8 m/s	6
Fig. 1.4. Varying regions of energy regeneration and motor plugging with transmission ratio	8
Fig. 2.1. Diagram of directional control valve	13
Fig. 2.2. Effect of pressure ratio on mass flow through an orifice.....	14
Fig. 2.3. Mass flow measurement system.....	18
Fig. 2.4. Mass flow with static flow restriction valve and with directional control valve.....	19
Fig. 2.5. Mass flow and adjusted mass flow.....	20
Fig. 2.6. Activation function given by sigmoid surface fit	21
Fig. 2.7. Four actuator configurations tested with Actuator 1	22
Fig. 2.8. Mass flow estimation.	23
Fig. 2.9. Total variation of mass estimate.	24
Fig. 2.10. Accumulated variation of downstream volume.....	26
Fig. 2.11. Estimation of displacement.....	27
Fig. 3.1. Bond graph model of a motor.	31
Fig. 3.2. Achievable behavior (continuous) for motor under passive switching.....	32
Fig. 3.3. H-bridge motor circuit and three equivalent circuits under passive switching	34
Fig. 3.4. Motor damping behavior as a function of duty cycle and velocity	37
Fig. 3.5. Controllability and energy generation metrics.....	41
Fig. 3.6. Testing setup with two motors, encoder, control electronics, and CAN interface device.	42
Fig. 3.7. Schematic representation of the control system	43
Fig. 3.8. Model validation results.....	44
Fig. 3.9. First two motor passive motor damping control experiments	45
Fig. 3.10. Third passive motor damping control experiment.....	46
Fig. 4.1. Simplified motor model within an H-bridge circuit and switching configurations	49
Fig. 4.2. Schematic of damping control structure.....	52

Fig. 4.3. Root locus of feedback damping control with conventional motor control techniques	53
Fig. 4.4. Simulation of conventional and passive control compared to an ideal damper	54
Fig. 4.5. Second set of simulated results.....	55
Fig. 4.6. Photo of the experimental apparatus	56
Fig. 4.7. Implementation of feedforward controller for passive switching	59
Fig. 4.8. Vibration metric as measured by FFT analysis of the tangential accelerometer sensor	60
Fig. 4.9. Results of first asymmetric damping experiment	61
Fig. 4.10. Results of the second asymmetric damping experiment.....	62
Fig. 4.11. Current tracking and controller output.....	63
Fig. 4.12. Achievable and unachievable behaviors with the passive control approach	64
Fig. 5.1. Photographs of the ECT prototype and plot of the motor-to-joint transmission ratio.....	68
Fig. 5.2. Summary of walking gait energetics and ECT control structure.	71
Fig. 5.3. Schematic diagram of the leg with variables and parameters of interest	73
Fig. 5.4. Optimization results for swing flexion.....	75
Fig. 5.5. Optimization results for swing extension.....	76
Fig. 5.6. Exemplification of baseline and optimized behaviors.....	77
Fig. 5.7. Amputee participant on the treadmill and schematic of the controller structures examined.....	79
Fig. 5.8. Average-stride results for fixed treadmill walking speeds	81
Fig. 5.9. Results of the set of experiments with varying walking speed	82
Fig. 5.10. Joint and motor mechanical energy and regeneration per phase of gait.....	83
Fig. 5.11. Per-stride overall regeneration efficiency of the system	84
Fig. 5.12. Average-stride electrical current signals for the motor, solenoids, and battery	85
Fig. 5.13. Breakdown of battery energy per stride with varying treadmill walking speeds	85
Fig. 5.14. Bond graph of a control system with a DC motor capable of energy regeneration	89
Fig. 5.15. Exemplification of optimized behaviors	92
Fig. 5.16. Examples of optimized behaviors with kinematic goals.	93

1 Introduction

This dissertation covers three main topics: fault monitoring in pneumatic systems, passive control of brushless DC motors, and optimization of energy regeneration in a motorized prosthesis. The first topic of fault monitoring in pneumatic systems is based on work that was developed in 2017 and 2018 and led to the publication of a journal paper which has been adapted into Chapter 2 of this document. The second topic of this dissertation, on passive motor control, was a subsequent research topic motivated by the development of the ECT (electronically controllable transmission) knee prosthesis in the Goldfab lab at the CREATE center. The author participated in the development of the ECT prototype, being responsible for the design, construction, and testing of the prosthesis embedded system and servo control, as well as contributing with controller development. Research on passive motor control explores the benefits of using a motor to implement braking tasks (such as higher power capability and energy regeneration) without compromising on the guaranteed control stability of a passive interface. This line of research culminated in two journal publications (the first one published, the second one accepted and in production) which are reproduced in Chapters 3 and 4. Finally, the last topic of this dissertation is on optimization of energy regeneration efficiency in knee prostheses, where the ECT is used for implementation and experimental evaluation of the method. Chapter 5 contains a draft of the journal paper on the topic which will be submitted for review in the following months.

The following subsections provide motivation and a general introduction to the three topics covered in each of the chapters of this dissertation. Chapter 6 contains a brief description of additional collaborative contributions by the author during graduate work, which have also resulted in second- or third-author publications.

1.1 Chapter 2: Fault Monitoring in Pneumatic Systems

The objective of the author's first research topic was to study the viability and develop methods that permit the detection of leak and motion faults in a pneumatic system, using only pressure sensors at the valve level. The system in consideration consists of an on/off directional control valve that configures an actuator (e.g., pneumatic cylinder) into two discrete positions: fully retracted or fully extended. This system represents most pneumatic applications in industry (servo control applications are dominated by electric actuators).

Pressure sensors have become low-cost and compact due to the popularization of vehicular tire pressure monitoring systems, enabling the installation of sensors within control valves. Limiting sensors to within the body of the valve simplifies the application of a fault monitoring system because the valve is already connected to the control apparatus (which would benefit from knowledge of these faults). If the fault monitoring system operates entirely within the valve, then the installation of such system would be as simple as replacing an old valve with a new model.

The author participated in the development of a pneumatic valve system with an embedded pressure sensor immediately prior to becoming a graduate student at Vanderbilt. A photo of the valve (with cover removed) and pressure sensing unit is presented in Fig. 1.1. The pressure sensor in the valve prototype is connected to one of the downstream ports, and was used to indicate leakage faults in addition to pressure faults (i.e., supply line pressure below or above a predefined value).



Fig. 1.1. Valve with exposed embedded electronics and pressure sensor in (a) and pressure sensor assembly in (b).

Detecting leak faults is important because air leaks incur significant operational costs due to elevated energy consumption and more frequent maintenance of air compressors. Leak faults are normally detected by monitoring air consumption with a flow meter located at each actuator supply line; flow meters are costly sensors, and the highly-distributed nature of actuators make them difficult to install in a factory setting. In the presented valve, leak detection was implemented by combining a pressure sensor with a unique valve configuration in which the two downstream ports were momentarily configured in fluid communication, creating a pressurized closed fluid circuit. If a significant pressure decay was observed during that period, a leak was indicated. Even though the approach was successful and novel in the industry, limitations remained. For instance, the relatively short window of time used for leak detection limited its sensitivity. Additionally, it was later determined that two sensors were necessary to detect motion faults.

Motion faults are defined as the inability of the actuator to reach either full extension or full retraction positions. This type of fault can be caused by obstacles in the trajectory of the actuator end effector or excessive load. This type of fault would normally be detected by installing position sensors at the actuator. Even though these sensors are not typically costly, running leads from the actuator unit to the control apparatus can be difficult if the actuator is not easily accessible.

The results of this line of research are presented in the paper “A Method for Mass Flow and Displacement Estimation in a Pneumatic Actuation System using Valve-based Pressure Sensing,” published in the IEEE Transactions on Mechatronics [1], and reproduced as Chapter 2 in this thesis. We show that it is possible to detect motion faults and accurately estimate air consumption with pressure sensors installed at the valve inlet and outlet ports. Characterization was limited to the charging process (forward stroke). Mass flow through the valve was estimated with a dynamic valve model. By assuming that the downstream system has a homogeneous pressure distribution, readings of the downstream pressure sensor and the estimated mass flow were combined to estimate downstream volume variation (which is proportional to actuator velocity). Therefore, air usage could be determined by integrating the mass flow estimate and actuator motion was determined by integrating the volume variation estimate. Integration drift was reduced by resetting the integrator at every stroke.

The dynamic valve model was parameterized for two commercially available valves. The fault detection methods were tested with the two valves and two cylinders. For one of the cylinders, multiple motion and load conditions were tested: free motion, impeded motion by a constant opposing force, obstacle at 50% of the trajectory, and actuator blocked. Only the free motion condition was tested with the second cylinder. In total, ten configurations of valve, cylinder, and load or obstacle were tested. Relative errors

for both air usage and displacement estimation were in the order of 10%. The case of impeded motion demonstrated higher estimation errors and the case of cylinder blocked demonstrated lower errors in general. The trend indicated that the approach was sensitive to long estimation times (due to integration drift) and performed less well for low-amplitude mass flow due to low signal-to-noise ratio.

1.2 Chapters 3 and 4: Controlling Brushless Motors to Implement Passive Braking Behaviors

Brushless DC (BLDC) motors are widely used, versatile machine components that have seen significant performance gains in the last decade [2]. As is well known, electric motors can operate in four power quadrants, adding or removing energy from the mechanical system in which they are inserted. In some control applications, BLDC motors are used to implement strictly passive behaviors, such as viscous damping, holding torques, or other forms of controllable motion-opposing behaviors. When implementing such energetically passive control tasks, the behavior of a BLDC motor is analogous to that of a physical embodiment of the element that it is meant to emulate, such as a viscous damper.

In a more general case, a BLDC motor implementing strictly passive control tasks can be viewed as an analog to a torque-controllable brake, such as a magnetorheological particle brake. The analogy between BLDC motors and particle brakes is well suited since electrical current will dictate braking torque in both devices, so arbitrary passive behaviors can be implemented in either case by measuring the system state and controlling electrical current according to a torque control law. For example, to emulate viscous damping, one would set the reference torque equal to the desired damping coefficient times the negative of the measured system velocity. However, while the torque produced by a brake is always intrinsically passive (i.e., the device cannot inadvertently add energy to the system) that is not the case with a torque-controlled motor using conventional switching and commutation techniques. For instance, a motor controlled in the manner described will momentarily deliver active torques when velocity changes sign due to control lag. That phenomenon, which is known in the haptics literature as “energy leaks”, can cause behavioral inaccuracies, vibration, and even control instability. That is in contrast with the strict passivity of brakes, which decouples behavioral accuracy from stability; that is, the system is always stable regardless of the control system performance.

The fundamental distinction between the two types of actuators is that electric motors are energy transducers connecting electrical and mechanical domains, while brakes are modulated resistors that dissipate mechanical energy directly as heat. That is, although torque-controllable brakes have an electrical port which is used to control torque, there is no power bond between that port and the mechanical domain; all mechanical power is dissipated as heat. That distinction confers higher power densities to motors relative to brakes because a portion of the mechanical power removed from the system can be directed away from the actuator as regenerated electrical power (e.g., to recharge a battery). In fact, a brief survey of commercially available BLDC motors and magnetorheological particle brakes showed that brakes have lower power capability relative to similarly sized BLDC motors, likely as a result of heat management and power transduction.

Energy regeneration in DC motors is, however, highly dependent on the desired operating point of the system. It can be shown that the ability of a motor to generate electrical power is limited to a maximum impedance (ratio of torque to velocity) which corresponds to the behavior obtained by short-circuiting the motor leads. For higher magnitudes of mechanical impedance, a motor will demand electric power to sink mechanical power; a generally undesirable condition referred to as reverse current braking or motor plugging [3], [4]. Those characteristics lead to regions in a torque-velocity plane in which a motor performs

advantageously and other regions where a brake would lead to less electrical power use and less heat to achieve the same behavior. A schematic representation of those regions is presented in Fig. 1.2. Note that the distinction in behavior is structural, emerging from the shape of the curves irrespective of device parameters.

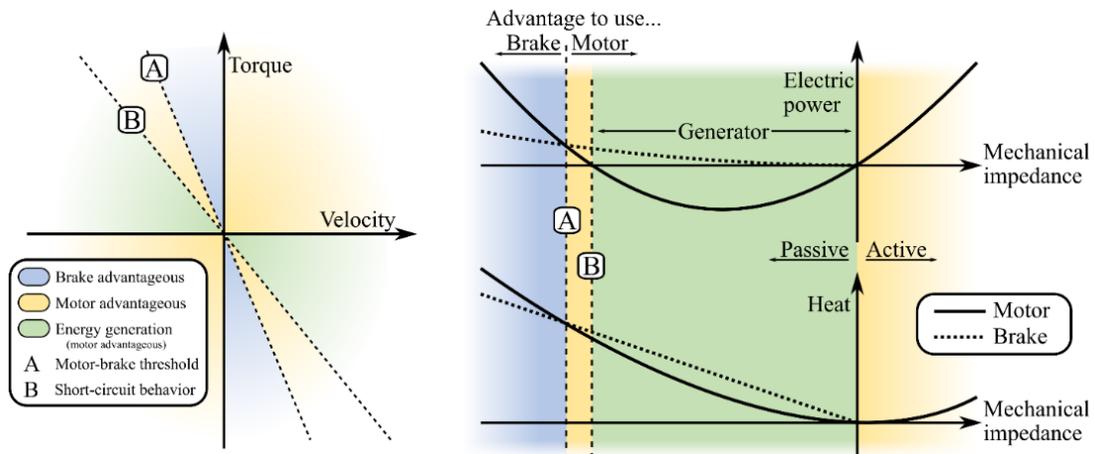


Fig. 1.2. Representation of regions in the torque-velocity plane in which a motor is advantageous over a brake in terms of electric power and heat generation. The same regions are represented as a function of mechanical impedance (ratio of torque to velocity) with exemplary curves for motor and brake electrical and heat power (valid for any non-zero velocity). Regions with blue background indicate that the use of a brake is advantageous in terms of electric power consumption and heat; yellow and green backgrounds both indicate that a motor is advantageous instead. The green background differs from yellow to indicate energy generation by the motor.

In summary, the use of a properly sized motor to implement braking tasks brings about advantages in terms of versatility, device packaging, heat management, and energy regeneration. The only situation in which a brake would outperform a motor is if very high impedances are desired (such as holding torques) or if passivity is required, which guarantees control stability.

Within that context, the works presented in Chapters 3 and 4 introduce a control and switching scheme for BLDC motors that allows a motor to operate passively like a brake, while preserving its energy transduction characteristics and therefore its advantage in power density. The method consists of shorting the motor leads in a PWM fashion and is implemented with the same hardware as typical motor control. The switching scheme constrains the motor to generator behavior (i.e., the green areas in Fig. 1.2) which generally comprises the majority of the region in which using a motor is advantageous.

Chapter 3 was adapted from a paper published in the ASME Journal of Dynamic Systems, Measurement, and Control [5]. The work compares the power capabilities of motors and brakes with a model-based analysis of power flow in both devices. Then, the passive switching scheme is presented, and its behavior is modeled and analyzed; the model is later validated with experimental results. A feedforward controller is constructed from the system model by numerically inverting it in real-time. The feedforward controller is experimentally shown to outperform the transient response of a feedback current controller. Conversely, the feedback controller presents better tracking in steady state. Ultimately, the combination of feedback and feedforward controllers results in a control system with the desirable characteristics of both methods.

The effects of choice of system parameters, namely PWM frequency and battery voltage, were studied with respect to the performance of energy generation and controllability. It was found that a reduced PWM frequency leads to improved control sensitivity (i.e., linearity) which was a counterintuitive result. The same effect can be observed with a reduction in battery voltage. Generation performance is reduced in both cases, indicating a tradeoff between controllability and generation.

Chapter 4, which was adapted from a subsequent paper to be published in the ASME Journal of Dynamic Systems, Measurement, and Control, compares the passive switching scheme to conventional motor control methods, which as discussed above, are prone to instability and energy leaks. Theoretical, simulated, and experimental results show that the passive method outperforms conventional methods in damping emulation tasks in terms of behavioral accuracy, stability, and vibration. The feedforward controller was reimplemented for the work in Chapter 4, this time being entirely implemented in the servo embedded system. As a result, the system no longer required connection to a computer, which noticeably increased the system performance by mitigating the effects of the communication delay between embedded system and computer.

Additionally, a slight modification to the control structure was made to permit passive control of asymmetric or unilateral impedance. That is, mechanical impedance can be independently controlled for different directions of motion. As a potential application, asymmetric behavior may be of use in the swing phase of a knee prostheses to quickly react to perturbations such as scuff or stumble. Unidirectional behavior may be useful in the stance phase to bias the leg towards extension. Passive asymmetric behavior may be an additional advantage of motors relative to brakes. Two feedforward controllers are used in the implementation of asymmetric behavior, one for each direction. The controllers are selected based on the direction of rotation reported by the motor encoder (a single encoder count is required to change behaviors). As is presented in Chapter 4, the passive asymmetric implementation significantly outperforms the conventional, feedback-based approach.

1.3 Chapter 5: Optimizing Energy Regeneration in a Microprocessor-Controlled Knee Prosthesis

Harnessing power from human locomotion has been a research goal in lower-limb prosthetics for several decades, particularly at the knee joint [6]. Energy harnessing offers performance advantages to electronically controlled prosthetic devices, allowing them to operate for longer periods of time, or alternatively to be designed with smaller and lighter batteries. The knee is considered the best candidate for energy generation out of the three lower-limb joints because it demonstrates a net negative energetic profile [7] and produces the greatest amount of negative work compared to ankle and hip joints [8] in level walking per stride. For a prosthetic knee device, the energy generation potential lies generally between late stance and terminal swing, where the knee joint power at moderate walking speeds is effectively nonpositive (see Fig. 1.3). The remaining phases of walking gait, comprising mostly of leg loading and support require greater amounts of torque to prevent buckling under the weight of the user and are approximately net zero in terms of joint energy.

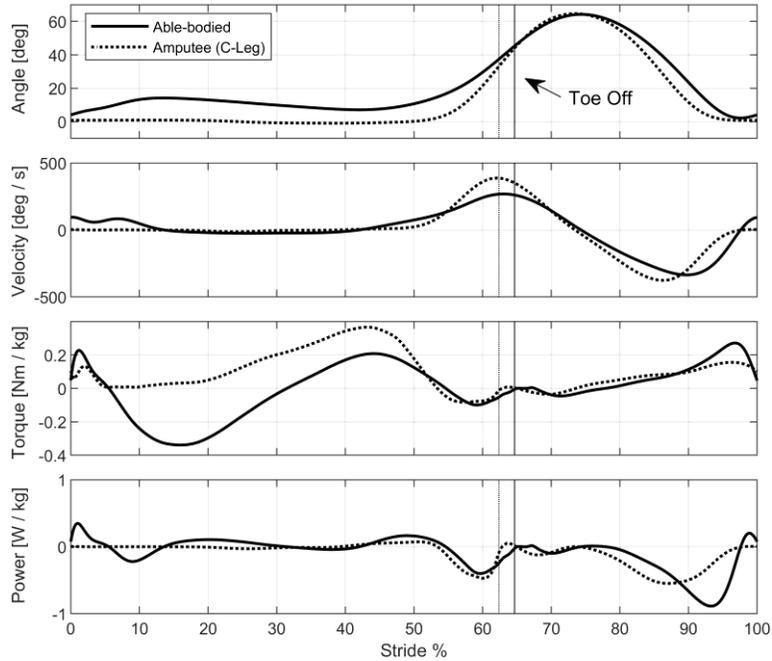


Fig. 1.3. Angle, angular velocity, torque, and power of the knee joint in level walking at 0.8 m/s. The solid lines show an average stride of several able-bodied participants and the dotted lines show an average stride of a single amputee participant wearing the Ottobock C-Leg.

A summary of the different actuator structures that have been applied to microprocessor-controlled knee prostheses (MPK) over the decades is presented in the section below, including commercially available devices and research prototypes. As will be presented, the different actuator constructions (with passive and active elements) lead to tradeoffs of function, control complexity, and regeneration potential. Therefore, we provide a literature review on the topic in order to discuss and justify the selection of an actuator structure in which the energy output can be maximized in a meaningful way.

A variable transmission ratio structure is ultimately selected for this work. The subsequent section briefly introduces the development of optimized energy harnessing controllers for the electronically controlled transmission (ECT) knee prosthesis prototype, which is the topic of the work adapted into Chapter 5. The prototype was designed and constructed by Steven Culver in Dr. Goldfarb’s lab. My involvement in the project consisted of developing the embedded system and motor control architecture for the device, as well as assisting Steven with the development of the first controllers for walking and stair ambulation.

1.3.1 Selection of Actuator Structure

Despite the significant energetic potential and the historical research effort, the leading commercial devices that are considered the standard of care for transfemoral amputation do not implement any form of energy harnessing in their actuators. For instance, the Ottobock C-Leg uses a hydraulic system modulated by control valves to restrict the motion of the knee joint and provide a comfortable walking gait with appropriate support to the user [9], [10], and the Ossur Rheo Knee uses a magnetorheological brake to the same end [9]. Note that the nonpositive energy profile of the knee joint in level walking allows for the employment of strictly dissipative actuators. Ultimately, these devices rely on the user to initiate motion and rely on batteries to power their embedded systems, sensors, and actuators. Both devices have been shown to improve the users’ gait biomechanics, metabolic cost, stability, and lead to

higher device confidence relative to implementations that do not employ electronic modulation (i.e., wherein the mechanisms that modulate the dissipative behavior are constructed mechanically) [11].

A major limitation of dissipative devices is their inability to produce power for tasks that would benefit from it such as stair ascent, sit-to-stand transitions, and level walking at low speeds. Users either avoid such tasks or employ compensatory actions which lead to pathological gait. To that end, knee prostheses that employ electric motors to actively control motion have been developed to enable amputees to perform such tasks with healthy gait patterns. Although not usually a stated goal, active devices have the potential for energy harnessing by operating the motor as a generator when the user backdrives the joint.

However, motorized devices struggle to adequately cover the wide range of behavioral impedance (ratio of torque to velocity) which is characteristic of walking gait; the stance phase is characterized as high-impedance while the swing phase is low-impedance. The dissipative actuators in the commercial devices mentioned previously can assume wide ranges of impedance and quickly vary that impedance. In the case of electric motors, operation at high efficiency and energy regeneration can only occur within a limited range of mechanical impedance [3], [4], [12]. More specifically, if a drive consisting of motor and transmission is sized such it can produce the elevated levels of torque for stance (such as [13], [14]), then its output impedance will typically be elevated and the range of impedances for which it will be able to regenerate may lie above the impedance levels associated with swing. That is, active devices that are appropriate for stance phase effectively require friction compensation in swing, leading to power expenditure. The natural motion coordination that arises from the inertial linkage between thigh and shank [15] is lost, and the device must initiate and control motion during swing by estimating the user's intention, leading to coordination and control challenges. Conversely, if a drive is sized for the requirements of the swing phase, then it will not be able to produce the elevated torques required for stance.

A demonstration of the tradeoff between optimizing the motor drive for high and low impedance levels is shown in Fig. 1.4. A lower transmission ratio allows the motor to generate energy during most of the swing phase, but stance behavior requires substantial energy cost; a higher transmission ratio allows the motor to operate efficiently in stance, but no regeneration is possible in swing. Note, from Fig. 1.3, that there is little power to recover in stance when amputee gait is considered; most power generation in amputee gait occurs in late swing.

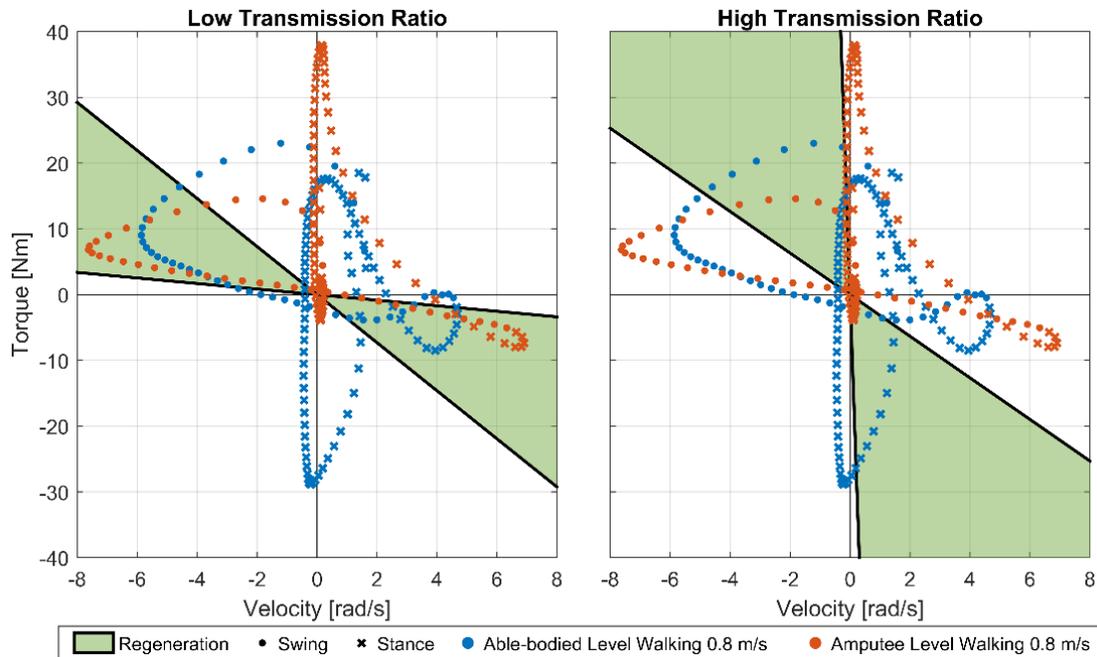


Fig. 1.4. Demonstration of varying regions of energy regeneration (green) and motor plugging (no background) with transmission ratio.

Motorized knee prostheses can therefore be classified based on which phase of gait their drives are sized to operate. Examples of devices appropriate for the stance phase have been shown to reduce metabolic cost of stair ascent (see [16] with the Vanderbilt Powered Leg [13]) and reduce loading on the contralateral leg in step-up and sit-to-stand tasks (see [17], [18] with the Ossur Power Knee). Partially as a result of the control challenge mentioned above, several studies report mixed results between strictly dissipative and stance-capable motorized devices and that, given the choice, participants oftentimes prefer the less capable dissipative devices as their daily use prosthesis [11], [18]. As expected, these devices do not report energy regeneration in walking tasks, instead relying heavily on batteries. Conversely, devices sized for the swing phase can achieve compelling results with energy harnessing [19]–[21], but the lack of stance support ultimately limits the applicability of the device and leads to important user safety concerns with the possibility of knee buckling.

Within this context, several research devices employ hybrid actuators that combine a torque-dense dissipative element to achieve safe stance support and a motorized drive appropriate for low impedance (swing). The use of a motor mechanically in parallel with a brake was originally proposed in the 1980's [22]–[24], [3] with the explicit goal of optimizing energy harnessing in level walking gait. The concept consists of controlling the motor to track an optimal regeneration setpoint while controlling the brake to produce the difference between the motor output and the desired behavior. This approach had some success, reporting a regeneration efficiency of up to 30% for benchtop, fixed-impedance trials [23], although the devices were never tried on users with amputation. Ultimately, one of the conclusions was that friction in the drive was too high to allow effective energy harnessing during walking, in part due to the presence of the brake [23], [24].

In summary, all four actuator configurations (brake alone, stance-appropriate motor, swing-appropriate motor, and swing-appropriate motor in parallel with a brake) contain unique benefits and limitations in

terms of function, energetics, and control. Crucially, energy harnessing seems to only be viable in the case of swing-appropriate motor alone (i.e., without the added impedance of a drive with high transmission ratio or a brake in parallel).

Finally, a fifth controller structure consists of the use of a single electric motor with a mechanism to vary the transmission ratio of the drive, allowing for the selection of high transmission ratios appropriate for stance phase functions and low transmission ratios appropriate for the swing phase. With that actuation structure, the prosthesis can provide adequate stance support, apply power when needed in either phase of gait, and assume low impedance in swing phase (enabling a natural coordination of movement and swing initiation by the user). Since this structure does not involve the addition of other actuators in parallel, the intrinsic impedance of the system can remain as low as the swing-optimized motor structure.

Recent examples of the variable transmission-ratio structure exist in the literature. The knee prosthesis presented in [25] uses an actively variable transmission (AVT) mechanism to modulate the geometry of the actuator crank slide, varying the transmission ratio. However, the AVT cannot be operated under load, so the transmission ratio is only varied between ambulation modes (e.g., from walking to stair ascent). Therefore, for walking tasks, the system suffers from the same dichotomy of stance or swing-optimized systems. Nevertheless, the authors report a significant amount of regeneration per step (around 3 J for the fast walking speed of 1.3 m/s) but also mention that energy consumption may increase in the future given that in the presented experiments the participant did not properly load the prosthesis in stance.

From the same research group, the robotic leg prosthesis presented in [26] uses passive variation of the AVT mechanism in the knee joint. Instead of actively varying the transmission ratio, the joint transmission ratio is modulated passively with the action of a spring that deflects when the joint is under torque. Differently than with their previous device, with this new structure, the transmission ratio varies between stance and swing. Although net energy regeneration is not reported at the knee for the default controller, a passive controller is offered as an alternative which requires additional user exertion but can achieve net-zero energy walking (that is, the prosthesis can operate indefinitely because all its internal systems can be powered from energy that is harnessed from the user).

Lastly, the Vanderbilt electronically-controlled transmission knee prosthesis introduced in [27], [28] uses a clutch mechanism to quickly alternate between high and low gear settings. The ECT already contains a set of controllers covering various ambulation tasks. Notably, swing control for walking tasks is achieved entirely passively for moderate and elevated walking speeds (i.e., with the motor operating as a brake as per the works in Chapters 2 and 3). However, the controllers developed initially focused on function, not on energy harnessing, so the amount of energy regenerated per stride was limited (around 2 J).

1.3.2 Deriving Controllers to Optimize Energy Regeneration in the ECT Prosthesis

Since the ECT device has the actuator structure deemed most appropriate for energy harnessing, we seek to improve the energy regeneration efficiency with optimal control techniques while maintaining the benefits of the actuator structure outlined above. Additionally, we do not want to derive controllers that achieve higher levels of regeneration at the expense of additional user effort, as was proposed in [26]. Namely, the knee should regenerate, without requiring any additional effort from the user, and without entailing any notable changes in kinetic or kinematic behavior of the knee.

For each phase of swing (flexion and extension), a constrained-optimization problem is posed and solved using a dynamical model of the lower leg, nominal and experimentally determined parameters of the ECT

prototype, desired kinematic outcomes (e.g., nominal peak swing flexion, terminal swing velocity, etc.), and exogenous inputs to the system (approximate motion of the thigh given by biomechanical datasets with amputee and able-bodied data). The optimization produces simulated motor current commands in time that minimize the error to desired terminal configurations, while maximizing regeneration. The authors are unaware of similar optimization approaches being used for any energy harnessing problem in any field, so the method is laid out in detail in Appendix I (in section 5.7). In order to preserve the desirable natural coordination of swing motion, the optimal results are not directly applied to the prosthesis. Instead, they inform motor damping commands implemented as functions of the real-time state of the leg with parameters to allow for user-specific adjustments.

Trials with an amputee participant who was already familiar with the original ECT controller were conducted. Parameters for the proposed control laws were determined experimentally. Embedded system data was collected with the participant walking on a treadmill with both fixed and varying velocity settings, between 0.8 and 1.2 m/s. We report a significant increase in energy regeneration, about 1 J per stride, in all walking speeds, while gait kinematics and joint mechanical power remained almost unaltered (relative to the original, baseline controller). The user reported that the control laws “felt the same” and could hardly tell them apart when asked. Altogether, regeneration efficiency (ratio of energy regenerated by the motor to total joint mechanical energy per stride) increased from 10-15% (increasing with walking speed) to about 20% (for any walking speed). The additional regeneration was also measured at the battery terminals.

An experimental control law for late stance (referred to as the stance flexion phase) was evaluated as well. The experimental strategy consisted of increasing motor resistance during that phase of walking to extract more power. Note that the foot is still in contact with the ground, so the hip controls motion with elevated torque capability. Ultimately, that strategy had little to no effect on both regeneration and gait kinematics and was not considered worth the additional control complexity and potential increase in user effort.

2 A Method for Mass Flow and Displacement Estimation in a Pneumatic Actuation System using Valve-based Pressure Sensing

This chapter has been published as a journal paper in IEEE/ASME Transactions on Mechatronics. Modifications to the text were kept to a minimum when adapting the journal paper into chapter format. The full reference to this work as published is presented below (referenced as [1] elsewhere in this document).

L. G. Vailati and M. Goldfarb, "A Method for Mass Flow and Displacement Estimation in a Pneumatic Actuation System Using Valve-Based Pressure Sensing," in IEEE/ASME Transactions on Mechatronics, vol. 26, no. 1, pp. 235-245, Feb. 2021, doi: 10.1109/TMECH.2020.3011348.

2.1 Introduction

Pneumatic actuation systems are widely used in industry because of their simplicity, cleanliness, low cost, flexibility, compliance, and high power-to-weight ratio [29]. One of the most common applications of pneumatic systems involves the use of on/off directional control valves to configure a pneumatic actuator (typically a pneumatic cylinder) into one of two discrete positions – fully extended or fully retracted. A manufacturing facility can employ tens or hundreds of these valve/actuator systems. Each valve/actuator system is herein referred to as an axis of actuation.

In many applications involving pneumatic actuation systems, a customer could benefit from knowledge of air consumption in each actuator axis (which may provide an indication of a leak), and also by knowing if there is a malfunction in the movement of the actuator (i.e., if there exists an actuator fault). Each of these measures can be obtained by dedicated sensors; however, air flow measurement is expensive, particularly if one is required on each axis of actuation. Actuator position measurement is less costly; however, position sensors need to be placed on the actuator, which is often remotely located relative to the measurement and control (e.g., PLC) unit, and requires wire leads to be run between the actuator and control unit, which is cumbersome and can reduce the robustness of the actuation system. The valve, conversely, is typically located in close proximity to the PLC unit, since the control connection between valve and PLC is a fundamental connection in the pneumatic control system.

In order to provide estimates of air consumption and actuator motion along each axis, the authors propose a method that employs pressure sensing at the directional control valve, rather than requiring mass flow sensing and motion sensing at the actuator. Specifically, with the recent widespread adoption of tire pressure monitoring systems (TPMS), semiconductor pressure sensors with pressure ranges appropriate for pneumatic actuation systems are compact and have become extremely low cost (i.e., a few dollars per sensor). The authors propose to employ a system model that can be identified offline to provide a low-cost and well-packaged option for estimating mass flow and actuator motion. Since the pressure sensors would be located at the valve, which is connected to the control unit, the approach does not require substantial changes in wire routing to implement.

The commercially-available pneumatic monitoring system Intellisense, introduced by the Bimba Manufacturing Company in 2014, is capable of detecting air leaks, stroke time, and end of travel, and monitors pressure and temperature at the actuator level [30]. However, the system has the limitation of requiring the installation of pressure sensors in very close proximity to the ports of the actuator; in this

work, the sensors are located at the valve level. Moreover, the Intellisense system does not estimate mass flow or air usage and cannot provide piston position information.

Despite their operational simplicity, precisely modelling the behavior of pneumatic systems, particularly the mass flow dynamics, is challenging due to compressibility of air, non-linear air flow dynamics, complexity in flow geometry, friction and stiction at the valve and actuator levels, and dead-band and dead time associated with valve function [31]–[33]. Several researchers have employed variations on the well-established model of gas flow through an orifice for modeling and/or control of pneumatic servo-actuation applications [31], [32], [34]–[38]. Of these prior works, only [31], [32] experimentally validated the mass flow estimates, and only [31], [37] utilized the aforementioned model with on/off spool valves. The estimation of transient air pressure at the actuator chambers based on actuator displacement was studied by Gulati and Barth in [39] and Turkseven and Ueda in [40]. The estimation of actuator position was investigated by Frédéric Abry et. al. in [41] with compelling results, but their method contemplated estimation of position only in the case of zero actuator velocity and requires the use of proportional valves on both ports of a double-acting actuator to modulate pressure during the estimation process.

In this work, a different variation of a mass flow model is presented, specifically one that describes the flow of air through a two-position directional control valve with significant switching dynamics, typical of the aforementioned on/off control industrial applications. Experimental results are presented that validate the ability of the model to predict the dynamic mass flow rate across a range of conditions, including a range of supply pressure and system geometry. Following validation of the mass flow estimation, the ability to estimate actuator displacement is experimentally investigated, also across a range of conditions including supply pressure, system geometry, actuator load, and movement obstruction. Relative to [41], the method proposed here is simpler and is applicable for real-time use in all load conditions but is not intended for proportional position control applications.

2.2 Adapted Mass Flow Model

2.2.1 Orifice Flow Model Adapted for On/Off Spool Valve

A widely used model to determine mass flow of a compressible fluid through a known orifice (derivation can be found in [29]) that is commonly utilized in modeling pneumatic valves is given by the equations,

$$\dot{m} = \begin{cases} k a p_u & \text{if } \frac{p_d}{p_u} < r_{th} \\ k b p_u \left(\frac{p_d}{p_u}\right)^{\frac{1}{\gamma}} \sqrt{1 - \left(\frac{p_d}{p_u}\right)^{\frac{\gamma-1}{\gamma}}} & \text{otherwise} \end{cases} \quad (2.1)$$

where p_u and p_d are, respectively, the upstream and downstream absolute pressures and \dot{m} is the mass flow through the valve. A diagram that shows a schematic of a directional control valve of the type contemplated in this work is presented in Fig. 2.1; the variables of interest are labelled in the diagram. The terms a , b , k , and r_{th} are given in (2.2)-(2.5):

$$a = \sqrt{\frac{\gamma}{R} \left(\frac{2}{\gamma+1}\right)^{\frac{\gamma+1}{\gamma-1}}} \quad (2.2)$$

$$b = \sqrt{\frac{2\gamma}{R(\gamma - 1)}} \quad (2.3)$$

$$r_{th} = \left(\frac{2}{\gamma + 1}\right)^{\frac{\gamma}{\gamma - 1}} \quad (2.4)$$

$$k = \frac{A_v C_f}{\sqrt{T}} \quad (2.5)$$

In all equations, γ is the ratio of the specific heat at constant pressure to the specific heat at constant volume of the fluid and R is the gas constant. The values of R and γ for dry air are approximately 287 J/(kg K) and 1.4, respectively. Note that a , b , and the threshold variable r_{th} depend only on these well-known terms. The threshold r_{th} marks the change in dynamics between choked and unchoked flow regimes, which is approximately 0.528 for dry air. The coefficient k is comprised of the terms that are difficult to determine analytically in a typical directional control valve application, and includes: the absolute temperature T of the fluid at the orifice, the effective orifice area A_v through which fluid is flowing, and the experimentally determined discharge coefficient C_f . When modeling a pneumatic valve, A_v is a dynamic term (i.e., a function of time) associated with spool displacement dynamics and internal channel geometry of the valve (i.e., the effective orifice area changes as a function of time as the spool travels within the valve). Note that since k is a function of A_v , it will also be a function of time. The type of valve used in this work is pilot operated and draws pilot pressure from the same line that feeds the actuator (see Fig. 2.1). For that reason, the evolution of A_v (and, therefore, k) in time will also depend on upstream pressure. Consequently, for a given valve, k is an empirical function of time and upstream pressure, which can be determined experimentally as described below. Note that the constituent components of k could be separately identified (i.e., the spool dynamics could be separated from the geometrical changes along the length of the spool and corresponding effects on areas and discharge coefficient). Treating k as a single variable with dependence on time and pressure, however, reduces the degrees of freedom in the model to the minimum necessary to describe the behavior of interest.

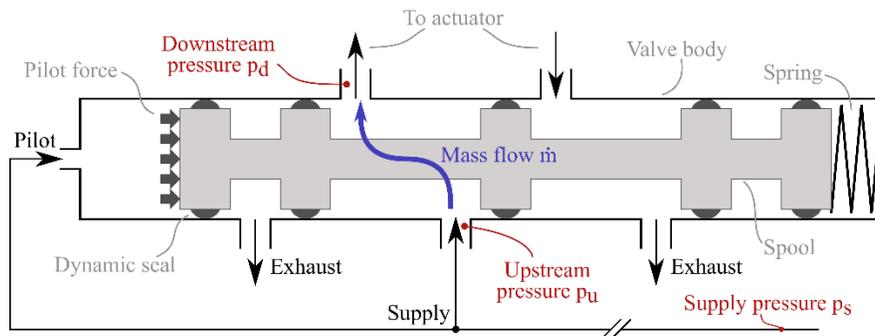


Fig. 2.1. Diagram of directional control valve.

From (2.1), the mass flow rate model can be rewritten as a product of k , the upstream pressure p_u , and a dimensionless multiplier \dot{m}_u which incorporates the nonlinear, pressure-dependent, portion of the model,

$$\dot{m} = k p_u \dot{m}_u \quad (2.6)$$

where \dot{m}_u is given by (2.7), in which the variable $r = p_d/p_u$ is the ratio of downstream and upstream pressures. The definition of \dot{m}_u is similar to the reduced flow function developed in [32] and the normalized flow rate function in [38]. The behavior of \dot{m}_u is presented in Fig. 2.2.

$$\dot{m}_u(r) = \begin{cases} a & \text{if } r < r_{th} \\ b r^{\frac{1}{\gamma}} \sqrt{1 - r^{\frac{\gamma-1}{\gamma}}} & \text{otherwise} \end{cases} \quad (2.7)$$

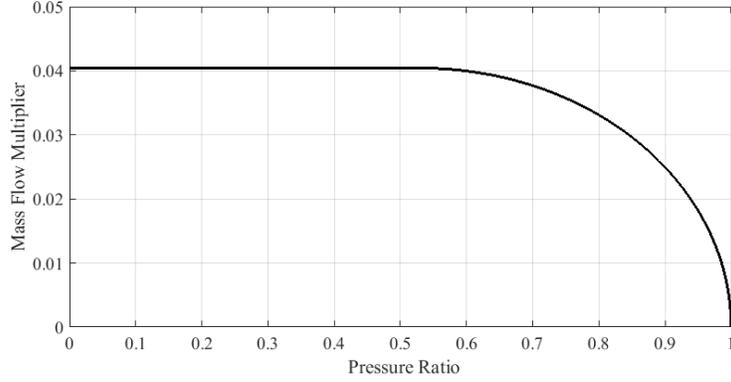


Fig. 2.2. Effect of pressure ratio on mass flow through an orifice.

This model can be further simplified by introducing what will be referred to here as the dynamic discharge coefficient (DDC), defined as $k_d = k p_u$. The final mass flow formulation demonstrates a separation between the effects of flow dynamics (given by \dot{m}_u) and the valve behavior (given by the DDC, k_d),

$$\dot{m} = k_d \dot{m}_u \quad (2.8)$$

2.2.2 Method for Finding the Dynamic Discharge Coefficient

The DDC is modeled by the empirical model with following form:

$$k_d \cong f(p_u) g(p_s, t) \quad (2.9)$$

where the function f represents the static flow behavior and effectively captures all static terms in conjunction with the upstream pressure p_u , and g is an activation function that is time and supply pressure p_s dependent, and represents the dynamic behavior of the system during the valve switching phase. As such, the function $g(p_s, t)$ evaluates to one when the valve is open and zero when the valve is closed. Intuitively, function g captures the dynamic behavior of the valve orifice area, while function f captures the static behavior of the valve as if no dynamics were present. Note from (2.5) and (2.9) that mass flow depends linearly on orifice area yet nonlinearly on upstream pressure, which motivates the separation of functions f and g .

2.2.2.1 Modelling the steady-state behavior with function $f(p_u)$

The first part of the model consists of determining the behavior of the system without the effect of valve dynamics; that is, treating the valve as a static orifice through which air flows. That behavior can be characterized by measuring mass flow with a mass flow meter in line with the valve in its fully open state. In order to obtain measurements with $g(p_s, t) = 1$, one can either lock the valve spool to eliminate the activation dynamics completely, or, in a simpler approach that does not require modifications to the valve,

take measurements of \dot{m} when it is known that the valve spool has reached the end of its trajectory within the valve body. Function f can then be estimated based on,

$$f(p_u) \cong \left. \frac{\dot{m}}{\dot{m}_u} \right|_{g(p_s, t)=1} \quad (2.10)$$

The upstream pressure term p_u was incorporated into the determination of f to allow for greater flexibility in modelling the nonlinear behavior of air flow through the complex internal geometry of directional control valves. Based on the measured behavior and a process of cross-validation, a fifth-order polynomial fit was chosen to model the variation in flow as a function of upstream pressure,

$$f(p_u) = \alpha_5 p_u^5 + \alpha_4 p_u^4 + \alpha_3 p_u^3 + \alpha_2 p_u^2 + \alpha_1 p_u + \alpha_0 \quad (2.11)$$

where values for the α_i coefficients are found with a least-squares linear regression routine. A model for mass flow that is valid after the valve dynamics have settled (steady state) is obtained with the product of the polynomial and \dot{m}_u ,

$$\hat{m}_{(\text{steady state})} = f(p_u) \dot{m}_u \quad (2.12)$$

2.2.2.2 Modelling the valve activation with function $g(p_s, t)$

The model must also take into account the spool movement and its effect on the valve orifice area and geometry. The contribution of the valve dynamics to the mass flow behavior can be characterized as an initial pure delay that holds the mass flow at zero as the spool travels across the spool land, followed by a quick rise in mass flow when the effective orifice area starts to increase as the spool moves beyond the land. Once fully displaced, the valve dynamics have no more influence on the mass flow behavior, and the system follows the steady state model in (2.12). In addition to the spool land, the pure delay phase is attributed to transmission delays through the internal pilot channels and spool stiction (both of which depend on supply pressure). The following activation phase in which the spool starts to move also depends on supply pressure, but to a lesser extent.

Note that supply pressure p_s is used here instead of upstream pressure p_u . The two pressures are the same when there is no flow through the valve. Once the spool starts to move and the valve opens, flow causes upstream pressure to drop below supply pressure. That drop, although significant to \dot{m}_u and to $f(p_u)$, does not affect the valve delay or activation behavior. To model the valve activation behavior, a sigmoid-like surface is employed that takes into consideration both time (main axis of the sigmoid) and supply pressure,

$$g(t, p_s) = \frac{1}{1 + \exp(-\beta_1 [t - \beta_2 (p_s - \beta_3)^{-1} - \beta_4])} \quad (2.13)$$

where the parameters β_i in the sigmoid model are found empirically using the measured mass flow \dot{m} normalized using both $f(p_u)$ from the previous result, and $\dot{m}_u(r)$,

$$g(t, p_s) \cong \frac{\dot{m}}{f(p_u) \dot{m}_u} \quad (2.14)$$

Then, a nonlinear optimization routine can be employed to find values for the parameters β_i .

2.2.3 Estimation of Actuator Displacement

Upon establishment of the functions that comprise the DDC, the mass flow model (2.1)-(2.8) can be employed to estimate downstream volume variation, which happens when the downstream actuator chambers expand or contract. By using the derivation in (2.15)-(2.18) one can obtain the time derivative of volume based on mass, mass flow, and pressure by applying the law of ideal gases [42]. The derivation relies on a few important assumptions: 1) that the charging process is isothermal, which provides a good approximation according to [43], [44]; 2) that the temperature of the air already in the cylinder is approximately equal to the temperature of air flowing through the valve; 3) that the pressure at the actuator chamber and lines can be approximated by the pressure at the valve outlet; and 4) that there are no air leaks downstream of the valve. Given these assumptions, the rate of change of volume is given by:

$$p_d V_d = mRT \quad (2.15)$$

$$V_d = RT \frac{m}{p_d} = RT \frac{m_0 + \Delta m}{p_d} \quad (2.16)$$

$$\dot{V}_d = RT \frac{\dot{m}p_d - \Delta m \dot{p}_d - m_0 \dot{p}_d}{p_d^2} \quad (2.17)$$

$$\Delta m = \int_0^t \dot{m}(\tau) d\tau \quad (2.18)$$

In the equations, p_d is the pressure at the downstream port of the valve, V_d represents the variable volume of the downstream system (consisting of lines and one of the chambers of the actuator), \dot{m} is the mass flow through the valve, Δm is the cumulative integral of \dot{m} , as in (2.18), and m_0 is the mass of air present in the downstream system at the moment of activation, $t = 0$, when the lines and dead volume of the actuator are charged to atmospheric pressure due to the preceding exhaust cycle (before the valve was activated).

The initial mass m_0 can be estimated employing the reasonable assumption that the actuator will not start moving immediately after the valve switches (i.e., volume variation should only start after a brief period of deadtime, t_d). This assumption is justified based on the delay of the pneumatic transmission lines, and stiction in the actuator. The value of the activation function, g , was used to dynamically determine when to estimate initial mass. Specifically, \hat{m}_0 is estimated when $g(t, p_s) \approx 0.75$. During the dead time, the downstream volume is constant, $\dot{V}_d = 0$ (isochoric behavior) and m_0 can be estimated using the relationship in (2.19), which is valid for $t < t_d$.

$$\hat{m}_0 = \frac{\dot{m}p_d - \Delta m \dot{p}_d}{\dot{p}_d} \Big|_{0.74 \leq g(t, p_s) \leq 0.76} \quad (2.19)$$

The estimation of m_0 occurs at an early stage in the actuation cycle, which allows its use to improve the estimation of volumetric variation during motion, resulting in the final form in (2.20).

$$\hat{V}_d = RT \frac{(\dot{m}p_d - \Delta m \dot{p}_d - \hat{m}_0 \dot{p}_d)}{p_d^2} \quad (2.20)$$

Note that the effect of m_0 in the estimation becomes negligible as the pressure converges (\dot{p}_d goes to zero) and rises (p_d^2 term in the denominator grows in the charging cycle). Note that, as implemented, \hat{V}_d is held at zero during the deadtime, which is consistent with the previously described assumptions.

If the actuator geometry is known, \hat{V}_d can be used to infer piston velocity. For instance, if an axis of actuation is composed of a single linear actuator with internal bore area A_d , the actuator's piston velocity v_d can be estimated by,

$$v_d = \frac{\hat{V}_d}{A_d} \quad (2.21)$$

Likewise, the piston displacement can be estimated with a cumulative integral,

$$\Delta d_d = \int_0^t v_d(\tau) d\tau \quad (2.22)$$

This integration permits estimation of motion along the axis of actuation without knowledge of the system volume or direct sensing of motion or pressure at the actuator level. The only information needed is the pressure at the valve ports (which are measured) and mass flow (which is estimated).

Even if actuator geometry is not known, one can perform a comparative identification of downstream variation of mass and volume between actuation cycles. For example, if the total mass used in an actuation cycle is less than the previous cycles, it can be inferred that the piston has not completed a full stroke (perhaps because of an obstacle in its trajectory). Likewise, variations in the magnitude of \hat{V}_d can indicate that the actuator's piston is moving at different speeds, likely because of a variation in load.

2.3 Experimental Validation

The mass flow model presented above was experimentally implemented and validated on two commercially-available two-position pilot-operated directional control valves (referred to as Valve A and Valve B), and each valve was tested with two double-acting actuators, Actuator 1 and Actuator 2 (i.e., four combinations of valve and actuator were tested). Note that the actuator volumes were different by approximately a factor of four. The manufacturer and model information for each valve, actuator, and sensor are given in Table 2.1.

Table 2.1. Instrumentation for Validation Experiments

Instrument	Model
Valve A	Parker HA2WXBG2G9000FA
Valve B	Norgren VS26G517DF313A
Actuator 1	SMC NCDGLA25-0300 (bore: 12.57 cm ² ; stroke: 7.62 cm)
Actuator 2	BIMBA SR-504-D-00MC (bore: 31.67 cm ² ; stroke: 10.16 cm)
Pressure sensor at tank, p_s	FESTO SDET-22T-D50-G14-U-M12
Pressure sensor at inlet, p_u	FESTO SED-16-10V/20mA 19564
Pressure sensor at outlet, p_d	FESTO SDET-22T-D50-G14-U-M12
Pressure sensor at volume, p_v	FESTO SPTW-P10R-G14-VD-M12
Displacement sensor	SpaceAge Control L021-00

A computer equipped with a data acquisition board NI PCI-6071E was used in conjunction with the software NI LabVIEW 2017 to collect data for experimental trials. Data was acquired with a sampling rate of 10 kHz, although a reduced sampling rate would be sufficient for subsequent applications of the model. The methods presented in this work were implemented in MathWorks MATLAB. Note that all computations employed in the estimator were strictly real-time (e.g., filtering, differentiation, etc.).

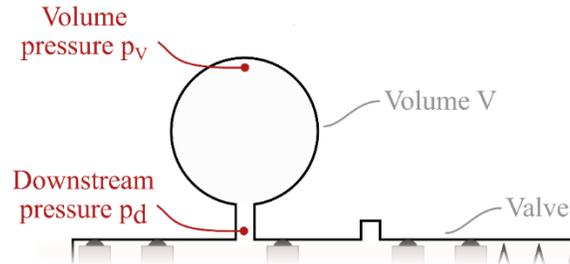


Fig. 2.3. Mass flow measurement system.

2.3.1 Test Setup for Mass Flow Model Parametrization

An initial experimental setup (which did not include any of the actuators) was employed to parameterize the mass flow model for each of the two valves. The setup employed four pressure transducers: one connected to a receiver tank, two connected respectively to the inlet and outlet ports of the valve, and the last one connected to a volume V (325 ml) installed downstream of the valve. Pressure at the volume was used in lieu of a mass flow meter, similar to the setup proposed in [33], [45], [46]. A schematic representation of the system is given in Fig. 2.3. This setup is typical of procedures used in standard flow rate testing of pneumatic devices [47]. Mass flow can be measured indirectly based on pressure p_V in the well-known volume V with a simple application of the law of ideal gases, assuming an isothermal charging process,

$$p_V V = m_V RT \quad (2.23)$$

where m_V is the mass of air in the volume. Mass flow \dot{m}_V can be obtained by differentiating both sides of the equality,

$$\dot{m}_V = \frac{\dot{p}_V V}{RT} \quad (2.24)$$

The distance between the pressure sensors in the ports of the valve and in the volume was approximately 0.3 m. A delay of one millisecond was added to the signals of the sensors at the valve level when comparing them to p_V to account for the propagation of pressure in the volume at the speed of sound (343 m/s).

To validate the experimental apparatus for measuring mass flow rate, a manual flow restriction valve was connected to the circuit in place of the directional control valve. The resulting mass flow to pressure ratio curves, corresponding to different valve opening and supply pressure configurations, are presented in Fig. 2.4 (a). Note that the behavior is very similar to the ideal curve in Fig. 2.2.

To obtain the data used to derive the mass flow model, a receiver volume of 19 liters was brought to approximately 9 bar and connected to the valve inlet. The connection between receiver and main supply was closed and the valve was repeatedly actuated until the pressure in the receiver tank was too low to pilot the valve. This allowed for a fine resolution in supply pressure configuration, since the volume of the

receiver is much greater than the downstream volume, i.e., each valve cycle causes a small decrease in receiver pressure. The pressure reading at the valve inlet immediately preceding each actuation cycle was used as nominal supply pressure for that cycle.

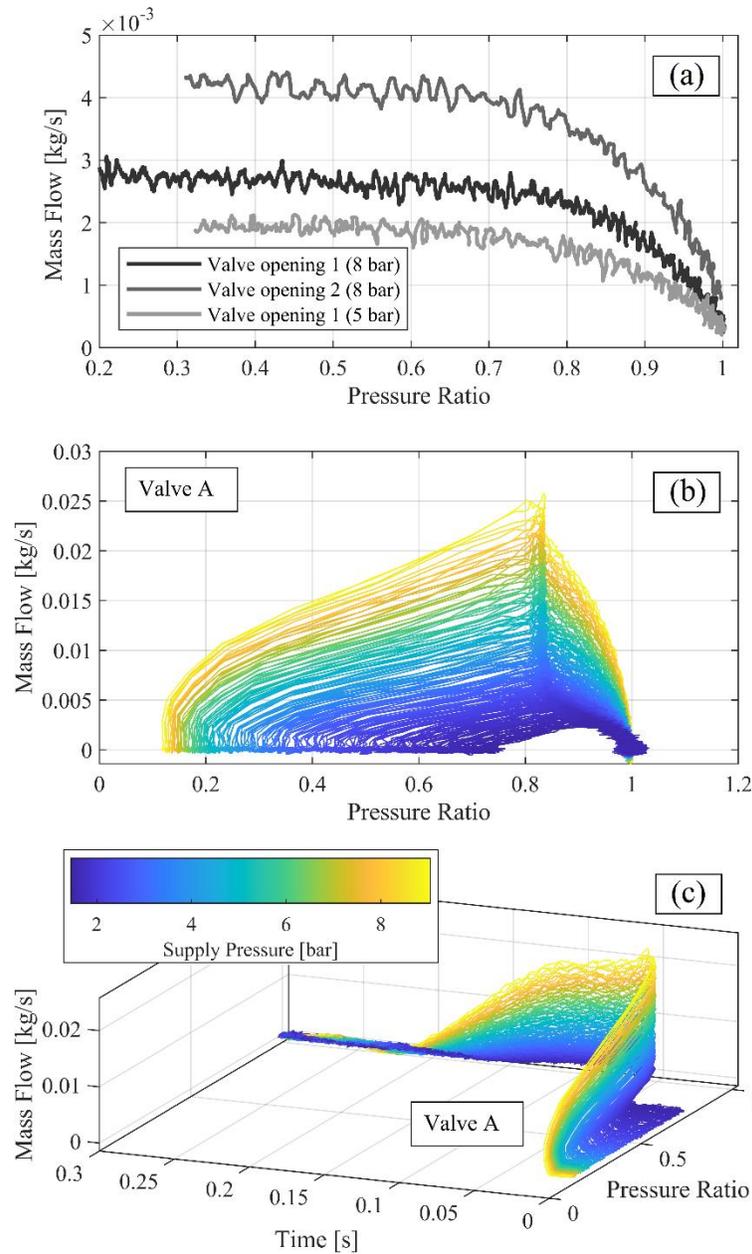


Fig. 2.4. Mass flow observed with static flow restriction valve in different opening and supply pressure configurations in (a); mass flow observed with directional control valve during charging cycle in (b) and (c).

The mass flow to pressure ratio curves, obtained by plotting \dot{m}_V as a function of r and p_s are presented in Fig. 2.4 (b) where different colors represent varying supply pressure configurations. The curves resemble the theoretical curve when pressure ratio is greater than around 0.8. By analyzing the same data with the addition of a time dimension, it becomes clear that the reason the curves do not match the theoretical model is due to the dynamics of the valve. This is shown in Fig. 2.4 (c). One can clearly see the

effect of the valve delay at $t = 0$ followed by a quick rise in mass flow that would indicate a fast spool displacement, increasing the effective orifice area in the valve channel. The same procedure was repeated for Valve B, with similar results.

2.3.2 Parameterizing the DDC Model

Recall that the DDC model includes both a steady-state and a dynamic function, f and g respectively. Based on the behavior observed in Fig. 2.4 (c), it can be reasonably assumed that the valve is in a steady-state condition when the mass flow is maximum for all configurations of supply pressure. The points of maximum flow rate lie at the division between the two regimes, when pressure ratio is small (i.e., \dot{m}_u is large) and effective orifice area is large. For every configuration of supply pressure, real mass flow \dot{m}_v and pressure ratio r are recorded at the point of maximum mass flow and then at time offsets of 20, 40, 60, and 80 ms from the maximum point. The steady-state flow behavior is then modelled according to the normalization in (2.10). In Fig. 2.5, the mass flow is plotted against the upstream pressure for all supply pressure configurations before (left) and after (right) the normalization by \dot{m}_u . Note that, for both valves, the normalization step is effective in making the points obtained at different times converge, indicating that flow is indeed following the structure of the model. The coefficients α_i that form the steady-state behavior function $f(p_u)$, obtained from a fifth-order polynomial fit routine, are given in Table 2.2.

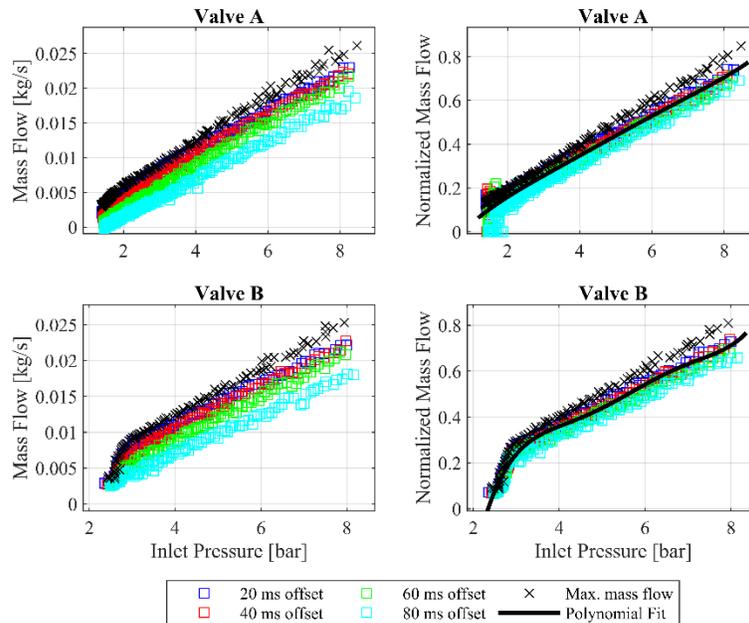


Fig. 2.5. Mass flow (left) and adjusted mass flow (right) for different configurations of inlet pressures with time offsets from the point of maximum mass flow.

Table 2.2. Polynomial Fit Parameters (inlet pressure in bar).

Parameter	Valve A	Valve B
α_0	-0.227	-6.412
α_1	0.202	6.374
α_2	3.276×10^{-3}	-2.389
α_3	-8.931×10^{-3}	$4,418 \times 10^{-1}$
α_4	1.520×10^{-3}	-3.976×10^{-2}
α_5	-7.677×10^{-5}	1.397×10^{-3}

The second step in deriving the mass flow model for this valve is to estimate the parameters β_i that form the dynamic behavior function $g(t, p_s)$. Mass flow \dot{m}_V is normalized using the previous result, $f(p_u)$, and \dot{m}_u , as in (2.14). The resulting points and the fitted sigmoid surface are presented in Fig. 2.6. The resulting parameters β_i are presented in Table 2.3. Note that in both cases the valve delay is close to constant for supply pressure configurations above 4 bar and becomes more variable below that point.

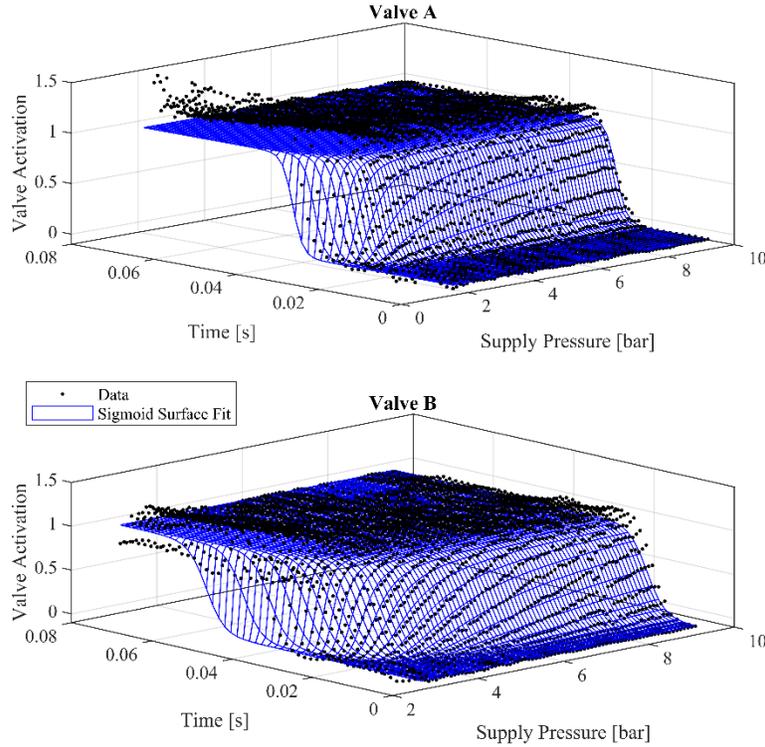


Fig. 2.6. Activation function given by sigmoid surface fit.

Table 2.3. Sigmoid surface parameters (time in seconds, supply pressure in bar).

Parameter	Valve A	Valve B
β_1	803.8	455.2
β_2	535.1	1630
β_3	1.247	2.306
β_4	2.159×10^{-2}	9.68×10^{-3}

2.3.3 Model Validation Experiments

Following parameterization of the valve models, the mass flow measurement apparatus was removed. Actuator 1 was installed using $\frac{1}{4}$ " flexible pneumatic lines (4.2 mm internal diameter), with length of 1.8 m. The displacement sensor was attached to the actuator to measure the true position y of the actuator's rod.

Four configurations were tested with Actuator 1: A) full movement with no load; B) full movement with constant force load; C) partial movement with no load; and D) actuator blocked. The four configurations

are represented schematically in Fig. 2.7. The constant force load was approximated by using a second pneumatic cylinder of smaller cross-sectional area positioned in series with the rod of the first cylinder and connected to the pressure supply line. The bore area of the smaller load-providing cylinder was 39% of the larger actuator cylinder. For the partial movement case, a hard stop was added at approximately 50% of the trajectory of the cylinder rod.

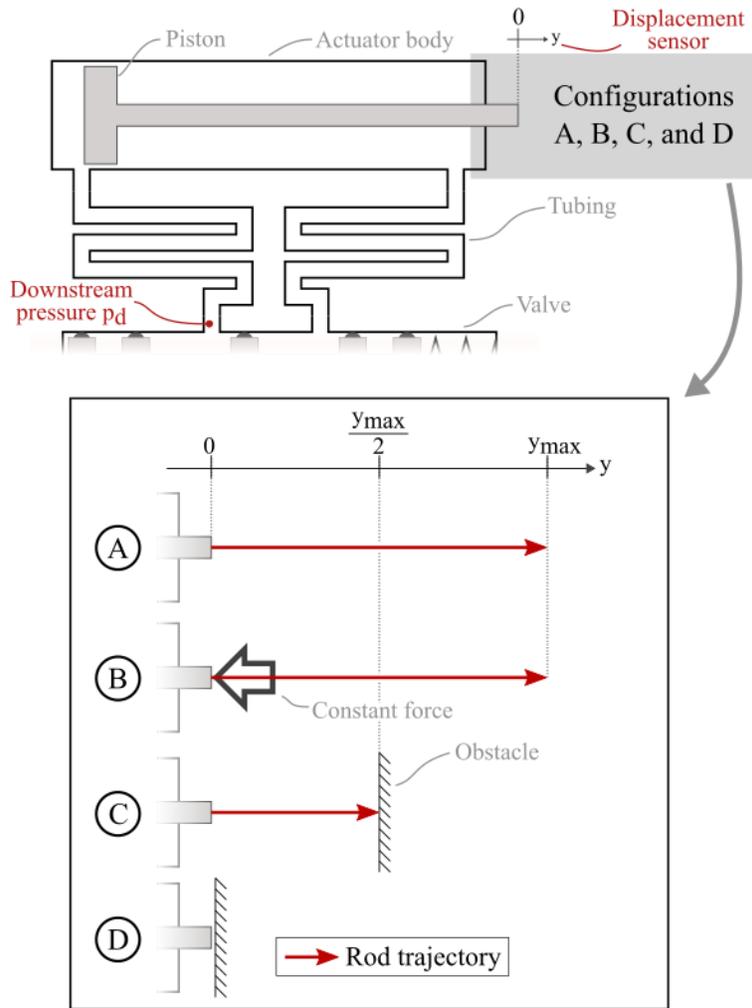


Fig. 2.7. Four actuator configurations tested with Actuator 1. Only configuration A was tested with Actuator 2.

For each of the four experimental conditions, a data collection procedure similar to that used for model parameterization was performed using the 19 l receiver tank: the receiver was charged to approximately 9 bar, connected to the valve subsystem, and isolated from the supply line. The actuation system was then allowed to cycle until the air pressure in the receiver was too low to pilot the valve. The procedure was carried out for both valves with Actuator 1 in the four load and obstacle configurations, totaling eight trials. Actuator 1 was then replaced by Actuator 2. In order to avoid an excessive number of permutations, only the first configuration (full movement with no load) was tested with Actuator 2, using both valves. The combination of Valves A and B with Actuators 1 and 2 are referred to as A1, A2, B1, and B2, respectively.

Pressure signals were calibrated prior to the experiments to ensure accurate measurement. The pressure and displacement sensors were filtered with a low-pass filter with time constant of 4 ms. Since the displacement estimator requires measurement of the derivative of the outlet pressure, a differentiating real-time filter was used, with the same time constant.

2.4 Results

2.4.1 Mass Flow Through Valve

The mass flow model was applied to the data obtained with the constant volume V and the result was then compared to the baseline mass flow, \dot{m}_V as in (2.24). Examples of the mass flow measured and estimated with the two valves and three settings of supply pressure are presented in Fig. 2.8. The importance of the valve activation term g is demonstrated by separating the contributions of the terms in the model presented in (2.9). Without the activation function, the mass flow model would behave as if the valve were instantly activated. It can be observed that the variable valve delay was appropriately captured, and the mass flow gain is consistent over the supply pressure range.

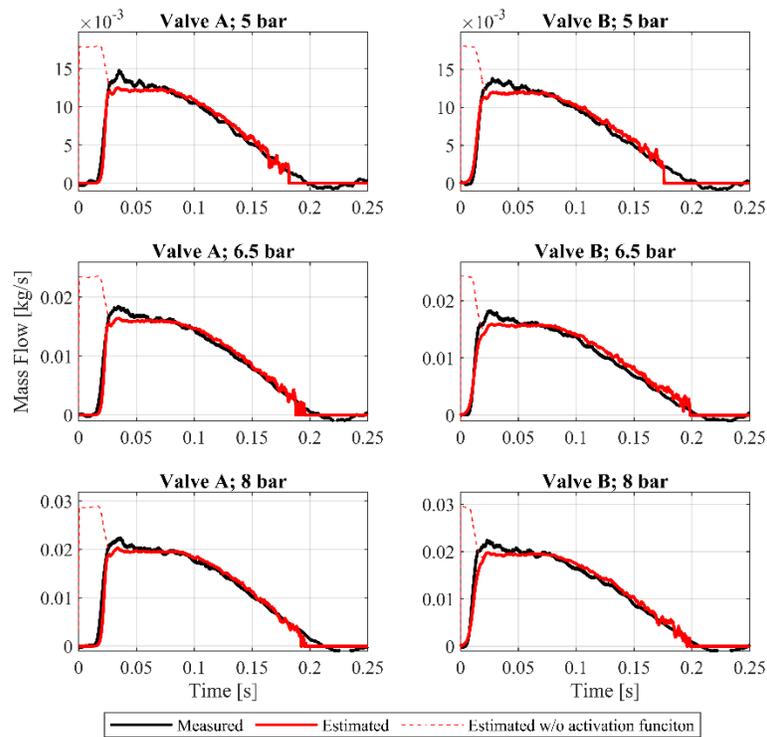


Fig. 2.8. Mass flow estimation.

To avoid the effects of sensor noise, a dead zone was added that effectively deactivates estimation whenever the difference between upstream and downstream pressures was below the threshold of 0.07 bar. The threshold was determined by qualitatively analyzing the noise characteristics of the pressure signals.

Absolute errors in estimation of mass flow are on average small, on the order of 1×10^{-3} kg/s, for all supply pressure cases. Because the true mass flow value equals zero in several points, the relative error was calculated for all configurations using the following formulation, for the first 300 ms of estimation,

$$e_{ps} = \frac{|\hat{m} - \dot{m}_V|}{\max_{p_s} \dot{m}_V} \quad (2.25)$$

where \hat{m} is the estimated mass flow obtained with the developed model and \dot{m}_V is the indirectly-measured mass flow into the volume, according to (2.24). This formulation yields an average relative error of 7.4% for valve A and 6.3% for valve B. Higher relative error was observed for lower supply pressure levels. When pressure is low, mass flow has smaller amplitude and is more affected by noise in the pressure reading.

2.4.2 Total Air Mass Consumption per Actuation Cycle

The estimation of mass flow \hat{m} allows for the indirect determination of mass of air used, according to (2.18). The estimation of $\Delta\hat{m}$ represents the mass added to the system that passed through the valve in one actuation cycle. The estimated mass variation in the four scenarios of the experiment with both actuators and both valves is presented in Fig. 2.9. Solid lines represent the expected mass of air given the total volume of the system in each motion and supply pressure configuration, discounting the mass initially in the system (i.e., lines and dead volume charged to atmospheric pressure).

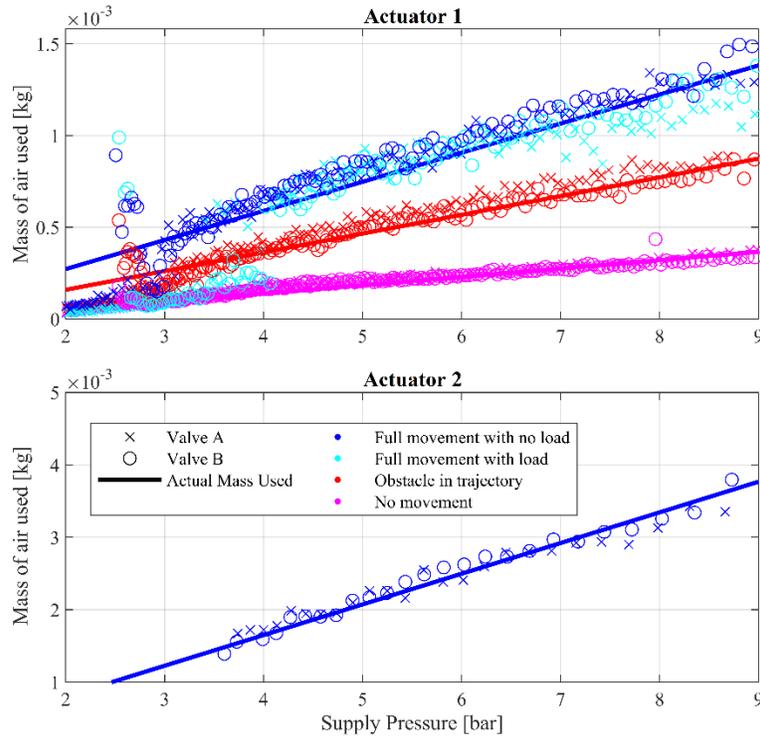


Fig. 2.9. Total variation of mass estimate.

In the experiments with Actuator 1, it can be observed that the two cases in which there is full movement (with and without load, in blue and cyan, respectively) match as expected, given that the total volume in both cases is the same. The three configurations of movement (no obstacle, obstacle at 50% of trajectory, and no movement) are easily separable for typical levels of supply pressure. Absolute and relative errors in mass estimation are presented in Table 2.4.

The load affects the velocity of the actuator rod, which reduces the amplitude of mass flow and causes mass to be underestimated in some cases. As previously mentioned, mass flow is assumed zero when the difference between the inlet and outlet pressure sensor reading becomes sufficiently small.

The experiment with Actuator 2 demonstrates that the difference in system geometry has little impact on the performance of the estimation process in relative terms. Errors were only calculated for supply pressure settings of above 4 bar because even though the actuators are still capable of operating, the quality of the results degrades substantially below that point.

Table 2.4. Error in estimation of mass use.

Configuration	Actuator	Average absolute error for supply pressure above 4 bar (mg)	
		Valve A	Valve B
No load, no obstacle	1	0.05	0.09
Load, no obstacle	1	0.11	0.05
No load, obstacle	1	0.07	0.03
No motion	1	0.01	0.01
No load, no obstacle	2	0.12	0.10

Configuration	Actuator	Average relative error for supply pressure above 4 bar (%)	
		Valve A	Valve B
No load, no obstacle	1	5.0	8.8
Load, no obstacle	1	9.9	5.1
No load, obstacle	1	11.0	4.5
No motion	1	5.2	5.1
No load, no obstacle	2	4.4	3.8

2.4.3 Estimation of Displacement

The four configurations of obstacle in the trajectory of the actuation axis and load are utilized in this section to test the performance of the actuator (i.e., rod) displacement estimator. By taking the cumulative integral of volume variation, one can determine the total variation in downstream volume at any time t . If the geometry of the actuator is known, volume can be translated into piston and rod displacement, using (2.22). In Fig. 2.10, the cumulative integral of the estimated volume variation is compared to measured actuator displacement in various example scenarios. Final displacement estimation for all contemplated configurations of valve, actuator, obstacle or load configuration, and supply pressure are presented in Fig. 2.11. Average absolute and relative errors in final displacement estimation are presented in Table 2.5. The relative errors are presented as a percentage of each actuator's stroke length. In the same way as the total mass variation results, error greatly increases when supply pressure is below 4 bar, indicating a limitation of the instrumentation. Therefore, those cases were not presented here.

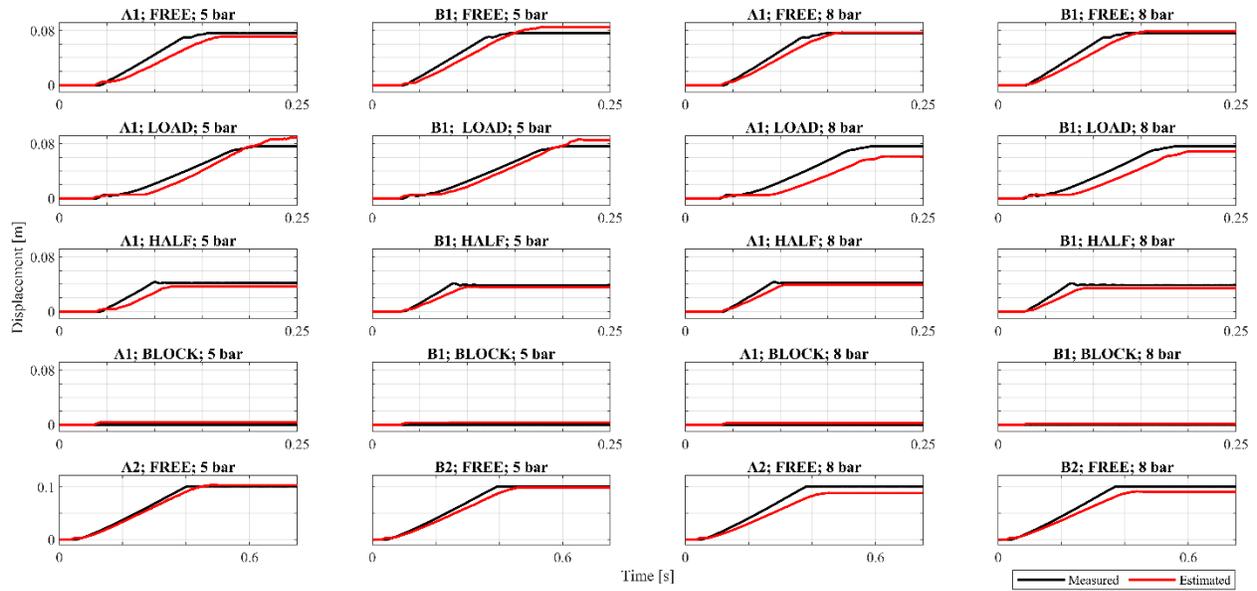


Fig. 2.10. Accumulated variation of downstream volume, which equates to estimated actuator displacement. The axes are titled according to the valve (A or B), downstream actuator (1 or 2), configuration (FREE, LOAD, HALF, and BLOCK, referring to the four scenarios of no load and no obstacle, constant load, obstacle at 50% of the trajectory, and actuator blocked, respectively), and supply pressure (5 or 8 bar).

A dead zone condition was added to the volume variation estimation (similarly to mass flow estimation) due to signal noise effects. Volume variation estimation is set to zero when mass flow is zero (i.e., mass flow dead zone is activated) and the pressure differential is below 50 bar/s. As an effect of the dead zone, it can be observed in Fig. 2.11 that displacement for the case of Actuator 1 with constant load is consistently underestimated. The lower velocity results in lower mass flow and slower pressure dynamics, which are partially lost in sensor noise.

Note that the relative error in final displacement for the case of obstacle at 50% of trajectory is smaller than the error encountered in the full motion cases. This lower error is to be expected because the system reaches steady-state (in which mass flow, pressure variation, and therefore volumetric variation are zero) in a shorter amount of time, which prevents the accumulation of estimation error.

Table 2.5. Error in estimation of rod position.

Configuration	Actuator	Average absolute error in final displacement for supply pressure above 4 bar (mm)	
		Valve A	Valve B
No load, no obstacle	1	4.65	7.81
Load, no obstacle	1	10.90	8.28
No load, obstacle	1	3.74	4.18
No motion	1	3.07	2.32
No load, no obstacle	2	7.57	4.84
Configuration	Actuator	Average relative error in final displacement for supply pressure above 4 bar (% of full stroke)	
No load, no obstacle	1	6.1	10.3
Load, no obstacle	1	14.3	10.9
No load, obstacle	1	4.9	5.4
No motion	1	4.0	3.0
No load, no obstacle	2	7.5	4.8

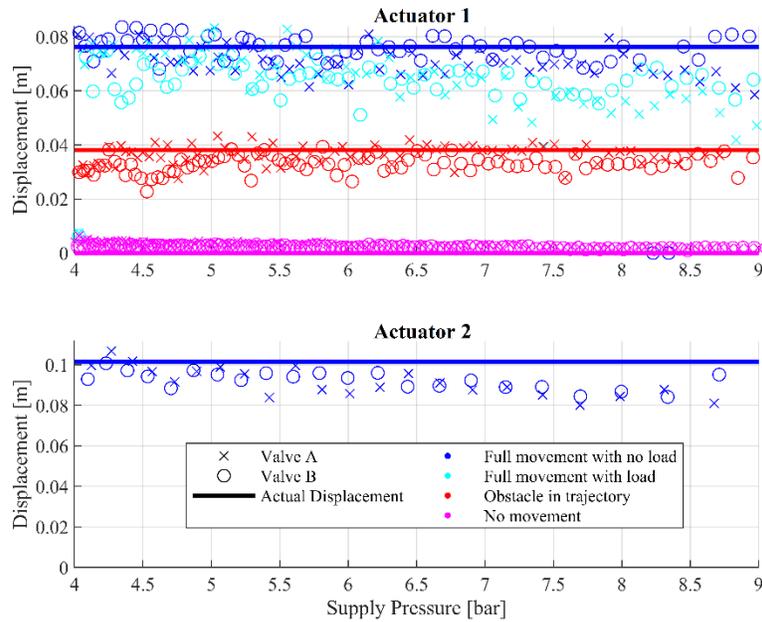


Fig. 2.11. Estimation of displacement.

2.5 Discussion

Based on the results presented in Table 2.5, the method presented here provided average real-time actuator position estimates of within 7.1% error, averaged across all experimental cases. For each load condition averaged across both valves, the method provided average position estimates within 8.2% error for the no load, no obstacle case in Actuator 1 and 6.5% in Actuator 2; 12.6% for the loaded case with no obstacle; 5.2% for the case of no load with obstructed movement at 50% stroke; and 3.5% for the case of obstructed movement at 0% stroke. As given in Table 2.4, the equivalent estimates for mass flow were within 6.3% across all cases, and were 6.9%, 4.1%, 7.5%, 7.8%, and 5.2%, respectively, for the load-specific cases. These real-time positional estimates are well within the accuracies suited to real-time fault detection in binary positioning systems, such as those typical in manufacturing. Additionally, the real-time mass flow estimates are provided with accuracies that would be useful in estimating air leakage, assuming the mass flow expectation could be initially benchmarked in the no-leak condition.

For purposes of implementation, the performance of the mass flow estimator is sensitive to accurate calibration of the pressure sensors at the valve ports. Mass flow estimation relies substantially on the pressure ratio; as such, a relative calibration procedure can be applied, consisting of finding gains that make the measurements match whenever it is safe to assume that the system is in steady state, in which case p_u should equal p_d . Note that even a slight difference in sensor gains will either make the estimate of mass flow converge to zero too fast (if $p_d > p_u$) or converge to a value greater than zero (in the opposite case).

The quality of the results degraded significantly for lower supply pressure configurations, when the signal-to-noise ratio in the pressure ratio signal is lower. That limits the applicability of the method, but the large majority of pneumatic systems operate using supply pressures above the experimentally found threshold of 4 bar (absolute) for which the methods demonstrate adequate results.

The model was applied to two valves which presented similar error characteristics, indicating that the level of accuracy observed is inherent to the method, not the valves. This is a favorable outcome that indicates the generality of the method in capturing the mass flow behavior of two-position spool valves. Tests with two actuators (the second with about four times the volume of the first) also demonstrated similar relative errors.

Variations in temperature were not accounted for in the models developed, and the isothermal assumption was made throughout the development of this work, even when the behavior of the system could have been modeled more precisely with, for instance, the adiabatic assumption. Literature shows, however, that the difference in displacement dynamics between models that use the isothermal assumption and models that use the opposite, adiabatic assumption is very small [43], [44]. More reliable estimates of mass flow could be obtained by adding a temperature sensor to the valve, in addition to the pair of pressure sensors proposed. Such a system would not require prior knowledge of air temperature (such as in the case of the experiments presented here). However, it is unlikely that temperature sensing at the valve would improve the estimation of displacement at the actuator level. Note that introducing a temperature sensor at the actuator would compromise the basic motivation for this work of keeping all sensing apparatus at the valve level. It is, therefore, inevitable that some estimation error will be introduced due to thermodynamic effects in the actuator.

The method presented here was tested using pneumatic transmission lines of 1.8 m (6 ft). Although experiments were not performed with longer lines to limit the number of experimental permutations, the modeling assumptions (of homogeneous downstream pressure) associated with line lengths employed here have been experimentally shown in other studies to be valid for line lengths up to 10 m (33 ft). Specifically, Turkseven and Ueda [48] show that explicitly modelling pressure variation associated with transmission lines becomes relevant only for line lengths above 5 m. Richer and Hurmuzlu [32] derive a simplified transmission line model that works well for typical line lengths (i.e., from 1 to 10 m) which could be used to improve the accuracy of the model if the geometry of the system and characteristics of the transmission lines are well known and if accounting for line dynamics becomes necessary. Recall that the approach presented here does not require any dimensional knowledge of the actuator in order to provide useful information. As such, requiring knowledge of line lengths, diameter, and their friction characteristics would diminish the applicability of the method, which was developed and intended to be realistically implementable in industrial applications, understanding that entering component dimensions during setup in such applications is unlikely to be realistic.

The setup also allows for the characterization of mass flow through the valve during the exhaust process by using the pressure in the outlet port of the valve as upstream pressure and the constant atmospheric pressure p_{atm} downstream. It would be possible to use those characteristics in combination with the charging characteristics to enrich the mass flow analysis. For example, one could compare the total mass accumulated in the actuator during charging with the total mass that was vented out during exhaust to obtain a more accurate air usage estimation.

The possibility of using a single pressure sensor, installed at the outlet port of the valve, was contemplated during the development of this work, although not presented here. A single pressure sensor can be used when assuming that the inlet pressure is constant and equal to the nominal supply pressure. Unfortunately, this assumption causes degradation of the mass flow estimation results, especially at the early stages of the actuation cycle, because the assumption of constant inlet pressure does not hold when

there is considerable mass flow through the valve (i.e., pressure drops between the supply and the valve due to pneumatic resistance).

Another point that was not discussed in this work is the potential for variations in spool movement dynamics over the lifetime of a valve due to changes in the spool friction characteristics resulting from seal wear. It could be of interest to employ a scheme that automatically updates the parameters of the activation function g either through a pre-programmed schedule or with a real-time estimation procedure. However, for supply pressure configurations of 4 bar and above, it is noted that most of the delay in both valves seems to be associated with air transmission delays between the pilot valve and the spool, which is inherent and immutable throughout the lifespan of the valve.

2.6 Conclusions

A method to estimate mass flow in internally piloted, spool operated, directional control valves was developed and experimentally tested with two commercially-available valves, two commercially-available actuators, and under differing load and movement conditions. The method requires measurement from two pressure sensors at the valve inlet and outlet ports and uses models that can be easily obtained offline (i.e., prior to valve installation). Estimates of air consumption can be immediately derived from the estimation of mass flow. Furthermore, the displacement estimator developed in this work allows for estimation of rod velocity and displacement, relying only on the aforementioned pressure readings, mass flow estimation, and basic knowledge of system geometry (e.g., actuator size). For the various cases experimentally tested, the method provided an average real-time estimate of actuator displacement within 7.1% error, and provided a real-time mass flow estimate within 6.3% error, relative to the full actuator stroke displacement.

These results demonstrate that the present method may be useful for identification of actuator motion and leak faults without the addition of any instrumentation at the actuator or knowledge of the actuator load, when the required accuracy is within the capability of the methods presented. Motion faults can be identified by variations of rod velocity or final displacement from nominal values, while leak faults would manifest as an increase in air usage. Since low-cost pressure sensing is now readily-available, the proposed method offers a low-cost and easy-to-implement means of fault detection in pneumatic actuation systems.

3 On Using a Brushless DC Motor as a Passive Torque-Controllable Brake

This chapter has been published as a journal paper in ASME's Journal of Dynamic Systems, Measurement, and Control. Modifications to the text were kept to a minimum when adapting the journal paper into chapter format. The only addition to the work is a more detailed derivation of the averaged PWM behavior in section 3.3.1. The full reference to this work as published is presented below (referenced as [5] elsewhere in this document).

L. G. Vailati and M. Goldfarb, "On Using a Brushless Motor as a Passive Torque-Controllable Brake," *J. Dyn. Syst. Meas. Control*, vol. 144, no. 9, Jun. 2022, doi: 10.1115/1.4054733.

3.1 Introduction

Several control applications require the use of an electrically-controlled proportional rotary brake, hereafter referred to as electrically-modulated brakes (EMBs), which offer strictly passive controllable resistance to motion. Among the EMBs employed in such applications are magnetic particle brakes, hysteresis brakes, and magnetorheological and electrorheological brakes. All such devices are energetically two-port devices – each comprising an electrical and mechanical port of interaction, wherein the rotary braking torque associated with the mechanical port is controlled generally by the current at the electrical port. However, these devices are not energy transducers, but rather modulated resistors (i.e., power is not directly transformed between domains; rather, the electrical input modulates a mechanical resistance to rotary motion). As such, power does not flow between the electrical and mechanical ports. These EMBs have traditionally been attractive for applications requiring modulated resistance because they provide high torque-to-weight and power-to-weight ratios; they provide guaranteed dissipation (i.e., are guaranteed to be strictly passive); they do not require measurement of velocity to impose a controllable torque; and when necessary they can provide a holding torque (i.e., torque without speed).

In recent years, the torque and power densities of brushless DC (BLDC) motors has increased substantially [49], and as a result, BLDC motors can now provide substantially higher power-to-weight ratios than EMBs. As such, employing a BLDC motor as a torque-controllable brake instead of an EMB in applications requiring a brake can offer performance advantages. As is well known, motors can be used with velocity feedback control to emulate passive impedance [50], [51]. However, as has been characterized by others [52]–[54], the resulting emulated behavior is not guaranteed to be passive, since energy may flow from the electrical domain into the mechanical domain as a result of measurement error and control lag inherent in all control systems. Moreover, emulating a resistance via feedback control can require substantial electrical power, depending on the desired impedance being emulated (condition known as reverse current braking or motor plugging) [3], [4], [24].

This chapter describes a control scheme for using a motor as a torque-controllable brake that provides high-fidelity torque control and guarantees strictly passive behavior. The approach requires no electrical power (other than signal power), but does require a means for removing energy from the system under consideration (i.e., a battery as is implemented here, as opposed to strict rheostatic braking as in [55], [56]). The approach constrains the motor to a passive-generator behavior, entirely preventing motor plugging at the expense of limiting its range of achievable impedance. Passive behavior in this configuration is not emulated: it emerges from the motor's constitutive characteristics. With this approach, a motor offers higher power density than a brake and a similar guarantee of passivity.

The control scheme is implemented via PWM shorting of the motor leads, which is known as the “three-switches” method in the regenerative braking literature [57]. Both the three-switches method, and similar single-switch method, have been adopted in recent works that seek to balance energy regeneration with factors such as braking time and battery state-of-charge in electric vehicles [58]–[61]. Despite implementation similarities, high-fidelity torque control and strict passivity are not treated in any of those works. Conversely, current control performance for motor braking is a focus of other recent works, such as [62]–[65]; however the proposed control schemes do not guarantee strictly passive behavior. Equivalent switching schemes to the one presented here have additionally been proposed for control of knee prostheses (e.g., [19], [20], [23], [66]), and a similar idea of exploiting the constitutive behavior of a motor has also been applied to haptic interfaces [67], [68].

Unlike prior works, this work explicitly treats the problem of achieving high-fidelity torque control in a brushless DC motor, while physically guaranteeing strictly passive behavior. Unlike the empirical models employed in [19], [20], the authors derive a physics-based model of torque as a function of motor velocity and control duty cycle, which among other features, explicitly considers changes in the current continuity regime (unlike [57]–[60], in which either continuous or discontinuous modes are assumed). Based on the developed model, a design trade-off is found between controllability and energy generation (the authors are unaware of other work in which this trade-off is explored). Experiments are presented that validate the proposed model and analyses. The model is then employed in a feedforward plus feedback torque controller, which is shown to provide notably improved performance relative to either control component used alone. As such, the derived controller is an important component of achieving desirable torque-control performance.

3.2 Modeling Motor Behavior

This section characterizes the bounds of behavior achievable when using a motor as a brake in a strictly passive manner, as proposed here. The results of this section are summarized in Fig. 3.2. Note that this section assumes continuous behavior in order to frame the problem, while the following sections assume a switching control structure to more explicitly explore it.

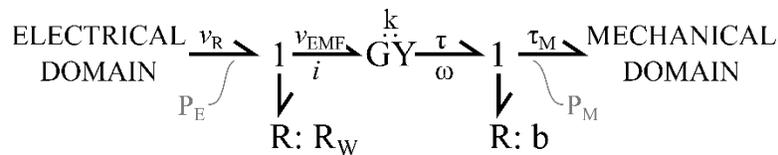


Fig. 3.1. Bond graph model of a motor.

3.2.1 Performance Bounds when using a Motor as a Brake

A simplified model of a DC motor is shown in bond graph form in Fig. 3.1. A gyrator element represents the electromotive phenomenon which connects the electrical and mechanical domains, a resistance with value R_W on the electrical side represents the motor windings (for now we ignore the electrical inductance) and a resistance with value b on the mechanical side represents mechanical friction. The electrical and mechanical power flowing in and out of the motor, P_E and P_M respectively, are given by:

$$P_E = v_R i = R_W i^2 + v_{EMF} i \quad (3.1)$$

$$P_M = \tau_M \omega = \tau \omega - b \omega^2 \quad (3.2)$$

where i , v_{EMF} , τ , and ω are the motor current, back-EMF voltage, torque, and angular velocity, respectively; v_R is the voltage across the motor leads; and τ_M is the total output torque (which includes friction). Using the gyrator relationships, torque and back-EMF are written as $\tau = ki$ and $v_{EMF} = k\omega$, where k is the motor constant. P_E and P_M can be rewritten in terms of the electromechanical damping (EMD) on the mechanical port of the gyrator, defined as the impedance term $z = \tau/\omega$.

$$P_E = \left(\frac{R_W}{k^2} z^2 + z \right) \omega^2 \quad (3.3)$$

$$P_M = (z - b) \omega^2 \quad (3.4)$$

Curves for P_E and P_M are schematically represented for an arbitrary velocity ($\omega \neq 0$) in Fig. 3.2 (b). Note that generator behavior is achieved when $P_M < 0$ and $P_E < 0$ (i.e., when power flows in the negative direction according to the arrow convention in Fig. 3.1). However, upon inspection of (3.3) and (3.4), one can observe that P_E becomes positive for $z < -k^2/R_W \triangleq -b_E$. This condition is referred to as plugging or reverse current braking and is generally undesirable because it requires an expenditure of electrical power ($P_E > 0$) in order to extract mechanical power ($P_M < 0$). Note that b_E is the reflected mechanical damping, which corresponds to the electrical resistance R_W reflected across the gyrator; as such, any desired damping value greater than this will result in plugging the motor.

Assuming $\omega \neq 0$, the conditions $z = -b_E$ and $z = 0$ are the only roots of P_E , and $P_E < 0$ between the two cases. To satisfy (3.1), the case $z = -b_E$ must correspond to $v_R = 0$. Therefore, $z = -b_E$ is the motor behavior when leads are shorted (i.e., when voltage across the leads is zero). The case $z = 0$ corresponds to the open-leads behavior, in which there is no motor torque because $\tau = z\omega = 0$, which implies $i = 0$. Therefore, the behavior of a motor under the proposed passive switching scheme must be limited to $z = -b_E$ for 100% duty cycle (short-circuit case) and $z = 0$ for zero duty cycle (open leads case). Since $P_E \leq 0$ within that range of z , expenditure of electrical power cannot occur. Thus, for the strictly passive control approach examined here, the motor must remain in the generator regime (i.e., motor cannot operate in the active or plugging regimes).

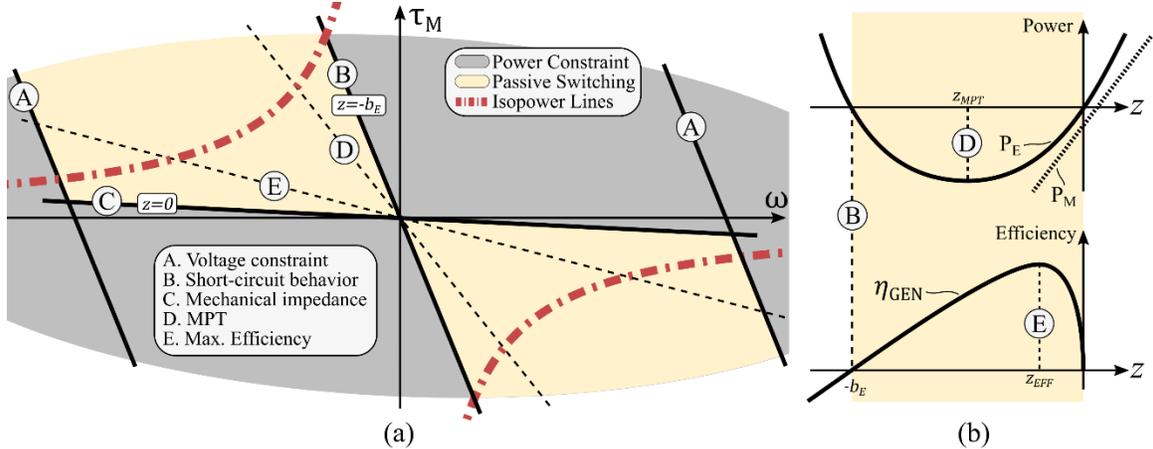


Fig. 3.2. Achievable behavior (continuous) for motor under passive switching. Areas indicated in beige represent achievable behavior with the proposed method of passive switching. The gray ellipsoid in (a) represents the heat dissipation constraint in motors in general, which exceeds the isopower lines at the same power level. Lines A, B, and C represent the voltage constraint, short-circuit behavior, and open leads behavior, respectively. Lines D and E represent the MPT and maximum efficiency conditions. Power and efficiency curves are represented schematically in (b) as a function of EMD.

3.2.2 Achievable Behavior of Motor relative to an EMB

Dissipation of heat in the armature of motors and brakes is a limiting factor for the achievable behavior of these devices. With the simplified model, heat losses can be modelled for both in terms of the output torque τ_M and velocity ω . In the case of the motor, power lost as heat corresponds to the power in the two resistance elements of the bond graph model in Fig. 3.1,

$$P_{HEAT(MOT)} = R_W i^2 + b\omega^2 \leq P_{HEAT(MAX)} \quad (3.5)$$

In the case of a traditional brake, all mechanical power is dissipated as heat,

$$P_{HEAT(BRAKE)} = \tau_M \omega \leq P_{HEAT(MAX)} \quad (3.6)$$

It is assumed that both motor and brake have a maximum amount of heat $P_{HEAT(MAX)}$ that can be dissipated continuously. The motor has two additional factors that limit achievable behavior. First, the voltage available in the electrical domain is limited to that of the battery, $|v_R| \leq v_E$. Second, specifically for motors under passive switching, behavior is further constrained to $-b_E \leq z \leq 0$, as previously demonstrated.

The combination of these constraints results in the regions of achievable continuous behavior in terms of mechanical power (P_M) for motors as schematically represented in Fig. 3.2 (a). The beige regions represent the achievable behavior of a motor under passive switching, and therefore the control method proposed in this paper is bounded by the extents of the beige regions. Note that the behavior of EMBs ($P_{HEAT(BRAKE)}$) is constrained by the isopower lines, which the motor far exceeds. Importantly, a holding torque (torque at zero velocity) cannot be passively achieved by a BLDC due to the intrinsic damping behavior.

3.3 Controlling a Motor as a Brake

Methods of torque control for BLDC motors are well known; in the case that control of impedance is desired, the reference torque can be computed from the combination of reference impedance and measured velocity. Methods of tracking a desired torque in strictly passive configurations have been studied in [57], [59] by leveraging the equivalence of the BLDC motor with a boost DC-DC converter and assuming either continuous or discontinuous current regimes to simplify the analysis. This section describes a model that explicitly considers operation in both the continuous and discontinuous regimes, and a controller based on it.

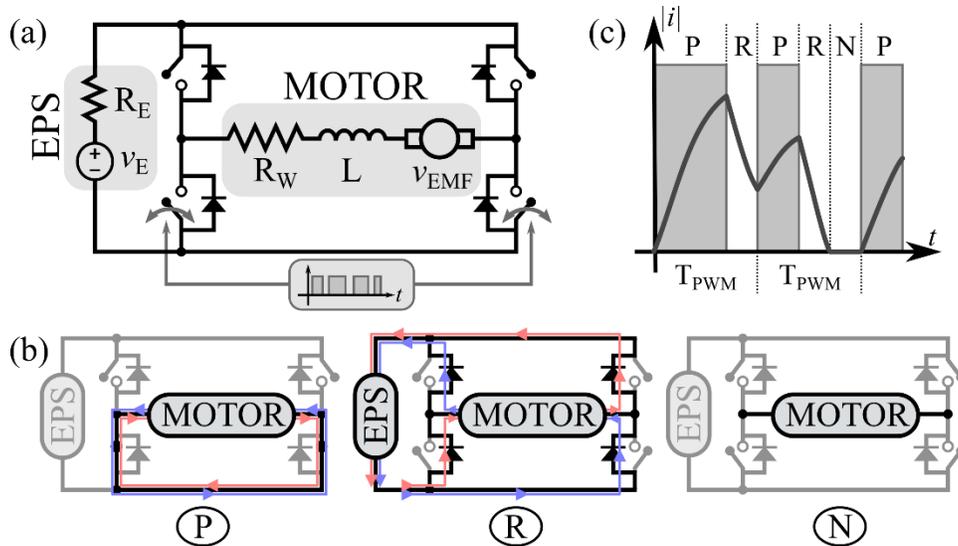


Fig. 3.3. (a) H-bridge circuit representation with motor and EPS model; (b) the three equivalent circuits (P, R, and N) under passive switching; and (c) an example of motor current behavior.

3.3.1 Control Structure and Modes

Fig. 3.3 (a) shows an H-bridge circuit that connects the motor to an electrical power system (EPS), which is modelled as a voltage source in series with a resistor. The H-bridge is a simplification of the three-phase bridge used to control BLDC motors in practice. The H-bridge representation facilitates analysis and yields equivalent results. A power supply is not required to implement a motor used in a strictly passive manner, as is proposed here; however, since a motor (unlike an EMB) is primarily a transducer, dissipated mechanical power needs to be channeled to some electrical outlet. In this work, we assume a battery will be used, and that the energy stored in the battery will be discharged for other purposes as appropriate. The MOSFETs are represented in the circuit as switches with parallel diodes.

The passive switching scheme leads to three equivalent behaviors as indicated in Fig. 3.3 (b): (P) short circuit behavior in the forced phase of PWM with the low-side switches active; (R) energy generation behavior in the unforced phase with all switches off and diodes conducting power from the motor to the EPS; and (N) open leads behavior in the unforced phase if the diodes are not conducting. Note that the high-side gates are not required for the control method presented here, so they remain open at all times. They are included in the bridge schematic to make clear that the approach is easily integrated into a standard motor interface, but diodes could be used instead. Note also that, unlike in active control schemes, the motor is never actively connected to the EPS via any of the MOSFET devices. Motor current increases in magnitude during phase P and declines during phase R. If current reaches zero during that phase, behavior N will be in effect until the next PWM cycle. The current behavior is exemplified schematically in Fig. 3.3 (c) in absolute terms (i.e., could be positive or negative depending on the direction of motion).

Table 3.1. Forced and unforced behaviors for the switching model.

Behavior	λ	f	Duration
P	$\lambda_P = \frac{L}{R_W + 2R_{ON}}$	$f_P = -\frac{v_{EMF}}{L}$	$t_P = uT_{PWM}$
R	$\lambda_R = \frac{L}{R_W + 2R_D + R_E}$	$f_R = -\left(\frac{v_E + 2v_D}{L}\right) \text{sgn}(i) - \frac{v_{EMF}}{L}$	$t_R = (1 - u)T_{PWM}$

The motor current behavior can be modeled in phases P and R as a first order differential equation:

$$\frac{di}{dt} + \frac{i}{\lambda} = f \quad (3.7)$$

where the forcing function f , time constant λ , and phase duration specific to each behavior (P, R) are given in Table 3.1. Rotor velocity is assumed constant within a PWM period, so f can be treated as a constant as well. MOSFET devices are modelled by their series resistance R_{ON} when active and diodes are modeled with a forward bias voltage term v_D and a series resistance R_D . The fixed EPS voltage is v_E , with series resistance R_E . Normalized control duty cycle is represented by u and the period of the PWM signal is T_{PWM} . Behavior N has no dynamics and can be modeled simply as $i_N = 0$. The differential equation in its general form has a closed form solution presented in (3.8).

$$i(t, i_0) = (i_0 - \lambda f) e^{-\frac{t}{\lambda}} + \lambda f \quad (3.8)$$

The closed form solution can be used to study the motor behavior as a function of velocity and control duty cycle. Note that the initial condition i_0 is explicitly expressed as an argument in the closed form solution. Average values for current in each phase can also be determined in closed form as in (3.9).

$$i^{AVG}(t, i_0) = \lambda f - \frac{\lambda(i_0 - \lambda f)}{t} \left(e^{-\frac{t}{\lambda}} - 1 \right) \quad (3.9)$$

To determine if behavior N will occur once the current reaches steady state, we determine whether the current magnitude can reach zero within t_R (assuming current starts from zero and reaches some magnitude in phase P). This is accomplished by solving for t_N in $i_R(t_N, i_P(t_P, 0)) = 0$, which yields:

$$t_N = \lambda_R \log 1 - \frac{\lambda_P f_P}{\lambda_R f_R} \left(e^{-\frac{t_P}{\lambda_P}} - 1 \right) \quad (3.10)$$

If t_N is less than the duration of the unforced phase, then current will reach zero before the end of the PWM period. That condition inhibits current build-up from one PWM cycle to the next: current magnitude raises from zero and returns to zero at each PWM period. Subscript D will be used to indicate that behavior (for current discontinuity [69]) and subscript C is used for the opposite case, in which there is current build-up. Average motor current over a PWM period in the discontinuity case is given by,

$$i_D^{AVG} = \frac{t_P i_P^{AVG}(t_P, 0) + t_N i_R^{AVG}(t_N, i_P(t_P, 0))}{T_{PWM}} \quad (3.11)$$

The average motor current in the continuous case is,

$$i_C^{AVG} = ui_P^{AVG}(t_P, i_R^*) + (1 - u)i_R^{AVG}(t_R, i_P^*) \quad (3.12)$$

where i_R^* is the current at the transition from behavior R to P and i_P^* from P to R. The two transition values can be determined by solving the simultaneous equations given by $i_P^* = i_P(t_P, i_R^*)$ and $i_R^* = i_R(t_R, i_P^*)$. The resulting linear system of equations is shown in (3.13), with the matrix-form solution in (3.14). With that arrangement, one can determine the average output of the PWM system in closed form (i.e., without requiring a simulation of the physical system).

$$\begin{cases} i_P^*(t_P, i_R^*) = (i_R^* - \lambda_P f_P) e^{-\frac{t_P}{\lambda}} + \lambda_P f_P \\ i_R^*(t_R, i_P^*) = (i_P^* - \lambda_R f_R) e^{-\frac{t_R}{\lambda}} + \lambda_R f_R \end{cases} \quad (3.13)$$

$$\begin{bmatrix} i_P^* \\ i_R^* \end{bmatrix} = \begin{bmatrix} 1 & -e^{-\frac{t_P}{\lambda}} \\ -e^{-\frac{t_R}{\lambda}} & 1 \end{bmatrix}^{-1} \begin{bmatrix} \lambda_P f_P \left(1 - e^{-\frac{t_P}{\lambda}}\right) \\ \lambda_R f_R \left(1 - e^{-\frac{t_R}{\lambda}}\right) \end{bmatrix} \quad (3.14)$$

Since the PWM frequency is typically well above the fundamental mechanical frequencies, we assume $\tau = ki_{AVG}$. Average current i_{AVG} is treated as a function of duty cycle u and velocity ω (recall that $v_{EMF} = k\omega$ is present in both f_P and f_R), and is either given by i_D^{AVG} or i_C^{AVG} depending on the current regime ($t_N \leq t_R$).

Finally, the motor EMD can be determined by,

$$z(u, \omega) = \frac{\tau}{\omega} = \frac{ki_{AVG}(u, \omega)}{\omega} \quad (3.15)$$

Fig. 3.4 (a) shows the nominal EMD behavior of a motor as a function of duty cycle and velocity. Motor current is simulated for four test points to demonstrate the current build-up effect in Fig. 3.4 (b). Note that EMD and motor current are expressed in absolute terms to facilitate visualization. In practice, achievable behaviors are all $z \leq 0$: current will always oppose velocity. Nominal motor parameters were obtained from the datasheet of a Maxon EC22 4-pole 90 W motor (part number 323217) which is the motor used in the experiments presented in this work. Nominal MOSFET parameters were selected from the datasheet of International Rectifier IRFU1018EPBF, which are the MOSFETs used in the 3-phase bridge circuit. All parameter values are listed in Table 3.2.

Table 3.2. Parameter values.

Parameter	Value
R_W	0.323 Ω
k	10.5 mNm/A
L	28.3 μH
v_E	24 V
R_E	0.1 Ω
R_{ON}	8.1 m Ω
v_D	0.65 V
R_D	18.2 m Ω
T_{PWM}	22.2 μs

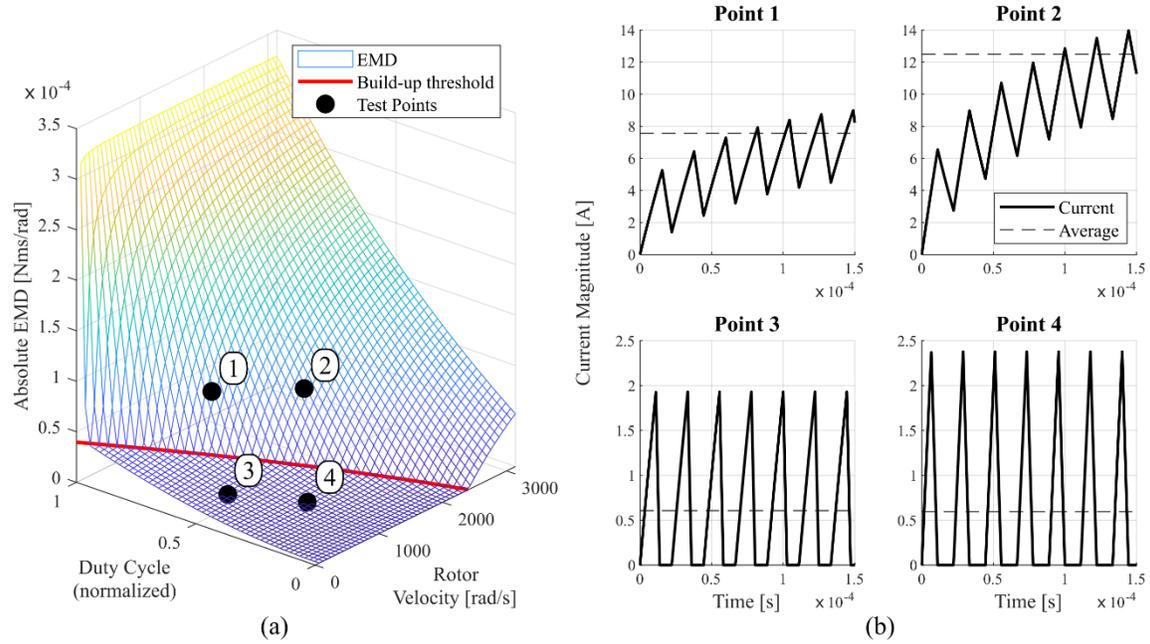


Fig. 3.4. Motor EMD behavior as a function of duty cycle and velocity in (a) and simulated motor current profiles in (b) for four test points. Test points 1 and 4 are at the same rotor velocity (1000 rad/s) with different duty cycles (70% and 30%, respectively); test points 3 and 2 are at the same duty cycle (50%) but different velocities (500 rad/s and 1700 rad/s, respectively).

There are a few important characteristics of the surface in Fig. 3.4 (a) that should be noted. First, note that the EMD behavior at 100% duty cycle is constant. It corresponds to the short-circuit behavior ($z = -b_E$) and is independent of rotor velocity (line B in Fig. 3.2). Second, behavior at 0% duty cycle corresponds to $z = 0$ (line C in Fig. 3.2) only up to a maximum velocity ω_{MAX} . At that point, v_{EMF} is large enough in magnitude to forward bias the diodes, which limits achievable behavior to the voltage constraint lines indicated as A in Fig. 3.2). ω_{MAX} is given by,

$$\omega_{MAX} = \pm \frac{v_E + 2v_D}{k} \quad (3.16)$$

Third, most achievable behavior occurs in the continuous case, which has implications for control sensitivity (i.e., linearity) at low velocities. Note that more than 90% of the EMD range is limited to around 10% of the duty cycle range in that region. The build-up threshold line indicated in red represents the transition between continuous and discontinuous current regimes which can be found by setting $t_N = 0$ in (3.10).

3.3.2 Impedance Control

Three approaches are taken to controlling motor impedance under passive switching. First a feedforward approach is presented which consists of numerically inverting the developed model. Errors in the nominal parameter values, unmodelled effects (e.g., winding temperature, MOSFET switching dynamics, capacitance in the EPS, etc.), and the high degree of control nonlinearity will lead to steady state errors in EMD tracking. The second method utilizes a closed-loop (feedback) PI controller (similar to [59]) to set the control duty cycle and track a reference motor current. The feedback method should be robust to model uncertainties and reduce steady-state tracking error. Unlike the feedforward approach, the feedback method has control lag. Finally, the third method consists of the combination of the two previous

methods: the inverted model is used as a feedforward controller, associated with the feedback PI controller. As shown subsequently in experimental implementation, the third method provides the best performance when tracking a desired impedance, combining the desirable characteristics of the other two methods.

3.3.2.1 Feedforward Method

To control the motor as a brake in a feedforward manner the model for $z(u, \omega)$ in (3.15) must be inverted so duty cycle u can be determined as a function of desired EMD z_{DES} and velocity ω . Unfortunately, the piecewise structure that arises from the current continuity issue and the switching between the forced and unforced phases of PWM make it impractical to invert the model analytically. Instead, the model is used to obtain a partial derivative of z with respect to u , and the Newton-Raphson method [70] is applied.

The Newton-Raphson method consists of iteratively determining approximations for the ideal u^* which corresponds to $z_{DES} = z(u^*, \omega)$. The method yields satisfactory results because the EMD surface is continuous, monotonically increasing, and has no inflection points. The partial derivative was constructed in such a way that values can be obtained even at the threshold between continuous and discontinuous behavior (where the behavior is not analytically differentiable). The method starts with an initial guess u_0 and uses the following equations to iterate starting at $j = 1$,

$$e_{j-1} = z(u_{j-1}, \omega) - z_{DES} \quad (3.17)$$

$$u_j = u_{j-1} + e_{j-1} \left(\frac{\partial z(u_{j-1}, \omega)}{\partial u} \right)^{-1} \quad (3.18)$$

$$e_{\%j} = \left| \frac{z(u_j, \omega) - z_{DES}}{z_{DES}} \right| \quad (3.19)$$

Two stopping criteria are employed: 1) relative error $e_{\%j}$ is sufficiently small compared to other sources of error and uncertainty in the system; or 2) the number of iterations j has reached a hard limit M . When a stopping criterion is reached, the duty cycle approximation of the last algorithmic iteration is assigned as the feedforward duty cycle u_{FF} for that control iteration, which is applied to the H-bridge. The limit on algorithmic iterations ensures that a duty cycle will be provided within the control loop period. Duty cycle is saturated to constrain it to $0 \leq u_j \leq 1$ at every iteration of the algorithm. If z_{DES} is unachievable, the method reaches the maximum number of iterations with u_{FF} saturated at either 0% or 100%. The initial guess for one control loop iteration is given by the resulting u_{FF} of the previous, which allows the method to converge across multiple control loop iterations if more algorithmic iterations are necessary. On startup, the initial guess is $u_0 = 0.5$. The relative error threshold was chosen as 0.1% and $M = 5$. With those settings, the method was observed to converge within the control loop period and provide adequate results.

3.3.2.2 Feedback Impedance Control

Measurement of average motor current i_{AVG} is used in the second control method. A proportional-integral (PI) controller with anti-windup was implemented in the discrete time domain. Since current direction will always oppose ω due to the motor's constitutive behavior, the controller only considers the magnitude of current and velocity. The setpoint for the controller is given by,

$$i_{DES} = \left| \frac{Z_{DES}\omega}{k} \right| \quad (3.20)$$

The PI control law is given by,

$$u_{FB} = k_p e + k_i e_{INT} \quad (3.21)$$

where $e = i_{DES} - |i_{AVG}|$ is the tracking error and e_{INT} is the anti-windup integral term. The duty cycle is constrained to $0 \leq u_{FB} \leq 1$ by saturation. The anti-windup method implemented here prevents accumulation of integral error by resetting the integral term via back-calculation if control is saturated [71]. This allows the control system to respond immediately if error changes sign.

3.3.2.3 Combined Impedance Control

The resulting duty cycle from the feedforward method u_{FF} and feedback method u_{FB} are added together in the third method. Duty cycle saturation and anti-windup must now use the total duty cycle: $u_{FB} + u_{FF}$.

3.4 Other Control Considerations

Several factors impact the motor's controllability under passive switching and its capability to generate power. As presented, controllability is affected by control nonlinearity at low velocities as well as limited range of EMD at high velocities (see Fig. 3.4 (a)). The nonlinearity issue is in part a result of the imbalance in magnitude between f_P and f_R at low velocities, because $|f_P| \ll |f_R|$ for small ω . The range limitation issue is intrinsic to electric motors due to the voltage constraint (i.e., it is not a characteristic of passive switching exclusively).

The only solution to the range limitation issue is to raise the EPS voltage. Unfortunately, that also worsens the imbalance between the two forcing functions at low velocities, affecting the linearity problem. Conversely, reducing the voltage is expected to improve linearity at low velocities at the expense of reducing ω_{MAX} .

An alternative solution that addresses the linearity problem with no effect on ω_{MAX} is to vary the PWM frequency. This approach leverages the imbalance between f_P and f_R and behavior N to raise the EMD surface in the discontinuous current region. Interestingly, lowering the PWM frequency helps. Recall that in the discontinuous region current in the unforced PWM phase reaches zero before the end of the PWM cycle. Therefore, stretching the duration of a PWM period has the effect of prolonging the P phase with little change to the unforced behavior (still limited by t_N). That is, current reaches a higher magnitude in phase P and has effectively the same behavior in phase R.

However, stretching the PWM period has a negative effect on energy generation, which only occurs during behavior R. As mentioned, behavior R maintains effectively the same duration due to the imbalance in forcing functions and current discontinuity while behavior P becomes longer in absolute terms. Therefore, the amount of time within a PWM period reserved for generation becomes proportionally smaller, resulting in less generation. Note also that $P_{HEAT(MOT)}$ in (3.5) depends on i^2 , therefore heat losses will be greater when current ripple has larger amplitude.

In the following sections, metrics for controllability and energy generation are developed so effects and tradeoffs can be visualized and interpreted quantitatively rather than qualitatively. The metrics are evaluated with varying EPS voltage and PWM frequency, since both parameters contribute to the presented issues and can be selected by the system designer.

3.4.1 Controllability and Power Generation Metrics

Metrics for behavior range, control linearity, and generation are defined here using the motor model developed in section 3.3, the ideal bond graph model in section 3.2, and a defined maximum application velocity ω_D . The latter represents the maximum velocity that should be used in practice; in this study, it was selected $\omega_D = 700 \text{ rad/s}$.

First, the range metric S_R is given by the relative difference between EMD at 100% duty cycle and at 0% duty cycle,

$$S_R = \frac{z(1, \omega_D) - z(0, \omega_D)}{z(1, \omega_D)} \quad (3.22)$$

S_R will evaluate to one when the full range of EMD behaviors is available at ω_D .

A second metric is used to indicate the linearity in behavior between z and u . Note that linearity here is a measure of how evenly EMD can be controlled throughout the achievable range. The ratio between the areas under $z(u, \omega)$ and an ideal linear case serves to indicate linearity,

$$S_L = \frac{1}{\omega_D} \int_0^{\omega_D} \frac{2 \int_0^1 z(u, \omega) - z(0, \omega) du}{z(1, \omega) - z(0, \omega)} d\omega \quad (3.23)$$

Similar to the range metric, the linearity metric evaluates to one in the best-case scenario, if z varies linearly with u . Note that $z(0, \omega)$ is subtracted to isolate the range limitation issue from the linearity issue (i.e., behavior may still be considered linear even if $z(0, \omega) \neq 0$).

The ability of the actuator to generate power may be of interest as a third metric. Two important operating regimes for a motor operating as a generator are maximum efficiency and maximum power transfer. Maximum efficiency $\eta_{GEN} = P_E/P_M$ occurs at $z_{EFF} = b - \sqrt{b^2 + bb_E}$ [3]. Similarly, maximum power transfer (MPT) into the electrical domain is achieved at an EMD value $z_{MPT} = -b_E/2$. Schematic plots of η_{GEN} , P_E , and P_M are shown in Fig. 3.2 (b). Unlike the maximum efficiency condition, the MPT condition does not depend on the mechanical impedance parameter b . Note that b is not as straightforward as parameter b_E and generally requires empirical models. Importantly, both MPT and maximum efficiency conditions are within the bounds of EMD for the proposed switching scheme. For purposes of a metric of interest, the MPT can be used to benchmark generation. Using P_E from (3.3) and z_{MPT} , the maximum power into the EPS in an ideal case is $P_{E(IDEAL)} = -b_E\omega^2/4$, while the model-based power is,

$$P_{E(MODEL)} = \frac{\min\{t_R, t_N\}}{T_{PWM}} v_E i_R^{AVG} \quad (3.24)$$

Again, ω_D is used here to compare the two, resulting in a normalized power generation score S_P ,

$$S_P = \frac{\max P_{E(MODEL)}}{P_{E(IDEAL)}} \Bigg|_{\omega=\omega_D} \quad (3.25)$$

A numeric solver is used to find the duty cycle that yields $\max P_{E(MODEL)}$ for ω_D .

3.4.2 Effect of Battery Voltage and PWM Frequency

The three metrics (control linearity, range, and generated power) are first presented in the top plot in Fig. 3.5 as functions of EPS voltage. Voltage is varied from a lower bound of 1 V to a maximum of 48 V while

PWM frequency is fixed at the nominal 45 kHz . The tradeoff between control linearity and range of achievable impedance is clear from the results. It can be observed that as EPS voltage rises, the range metric quickly reaches its maximum value of one while linearity monotonically decays. There is also a noticeable variation in the energy generation metric. The effect can be attributed in part to the proportionally higher diode losses (due to the constant forward bias voltage relative to the EPS voltage). At the nominal EPS voltage of 24 V , the range limitation issue is not present; the system is close to its maximum power transfer capability; and control linearity is poor relative to lower voltages.

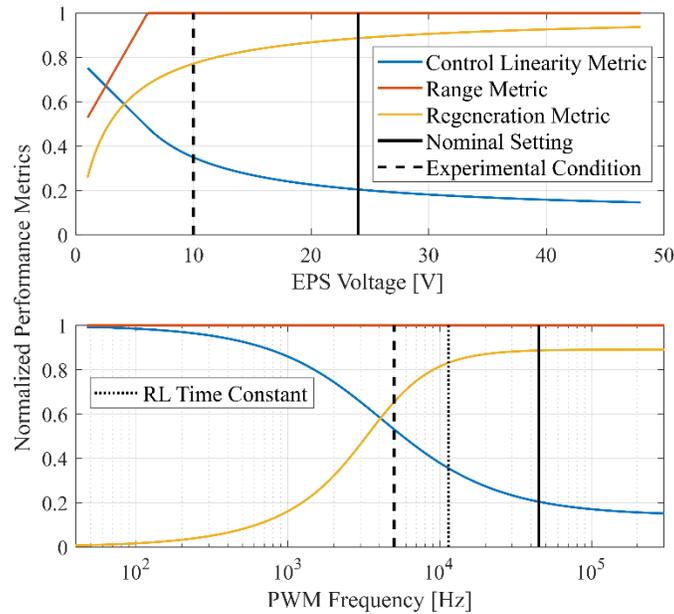


Fig. 3.5. Controllability and energy generation metrics for varying EPS voltage (top) and PWM frequency (bottom).

In the bottom plot in Fig. 3.5, the performance metrics are presented as a function of PWM frequency, which varies from a lower bound of 50 Hz to an upper bound of 200 kHz . EPS voltage is fixed at the nominal 24 V . The range limitation issue does not occur when PWM frequency changes, so that metric remains at the maximum value of one. The tradeoff between linearity and generation occurs between 1 kHz and 50 kHz : lower values of PWM frequency will favor controllability and higher values will favor generation. The S-shape of the two curves indicates that, for purposes of using a motor as a brake, there is little advantage in going below or above those values (for the parameters of the motor used in the study). At the nominal PWM frequency of 45 kHz , the tradeoff is largely skewed towards higher power generation at the expense of nonlinear behavior. The electric circuit time constant (λ_p) is indicated for reference. Note that the generation metric does not consider switching losses, which are larger at higher PWM frequencies. PWM frequencies lower than 20 kHz may cause audible noise. If frequency is too low, switching may be perceptible in the mechanical domain as vibration.

In summary, controllability and generation metrics are inversely related with respect to the choices of EPS voltage and PWM frequency. Somewhat counterintuitively, a lower PWM frequency provides increased resolution in controllability. A control designer may choose both parameters based on the performance trade-offs, as outlined in this section, and their application objectives. Favoring controllability over generation may be particularly important in the case of low application velocities, in which control nonlinearity and sensitivity are challenging, as previously discussed.

3.5 Experimental Implementation

An experimental test setup was built in order to validate the motor model, as well as evaluate the relative performance of the strictly passive impedance control approaches. The setup, shown in a photograph in Fig. 3.6, includes two motors: one BLDC motor acts as a brake under the proposed switching scheme, while the other (brushed) motor simulates an active load, driving the first. The two motor shafts are directly connected: there is no transmission and minimal system inertia.

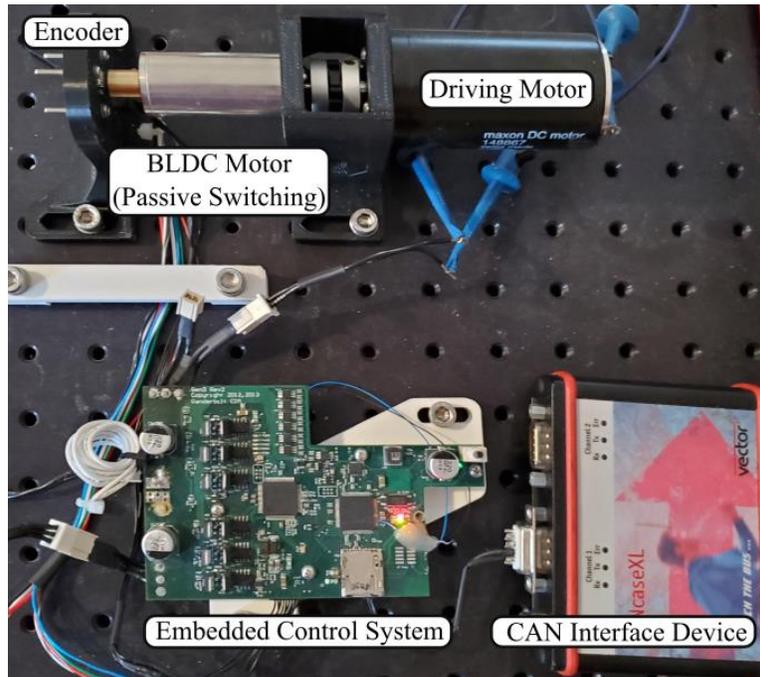


Fig. 3.6. Testing setup with two motors, encoder, control electronics, and CAN interface device.

The BLDC motor is controlled via three MOSFET half-bridges, while the brushed motor is controlled via a MOSFET H-bridge. Electrical current is measured via series shunt resistors, which are connected to current sense amplifiers to enable measurement of current in each phase. Rail voltage is also measured. A 1024 count per revolution magnetic incremental encoder is installed in line with the motor shafts. A real time differentiating filter in the control electronics provides velocity estimates from the encoder. Hall effect sensors in the BLDC motor indicate commutation sector, which is used in conjunction with phase current measurements to approximate motor torque.

The control electronics are powered by a linear power supply, which can source and sink electrical power at a constant voltage. The control electronics is connected to a computer through a CAN bus. Mathworks MATLAB and Simulink with the Desktop Real-Time kernel are used to interface with the control electronics at 500 Hz. The computer and power supply are not shown in Fig. 3.6.

A schematic representation of the control system is presented in Fig. 3.7. The feedforward controller was implemented in Simulink, providing duty cycle u_{FF} to the control electronics at 500 Hz. The feedback controller that defines duty cycle u_{FB} was implemented directly in the control electronics, running at 5.8 kHz. The two controllers can be separately enabled and disabled through the computer interface (“mode” selection). To control the driving motor, a PI velocity controller implemented in the computer

provides a current reference i_{VEL} to another current controller in the control electronics, which regulates the driving motor current (i_{DRV}) by varying the duty cycle u_{DRV} . Gains of the velocity controller are scheduled with the reference motor EMD, which is necessary due to the wide variations in system impedance caused by the motor operating as a brake.

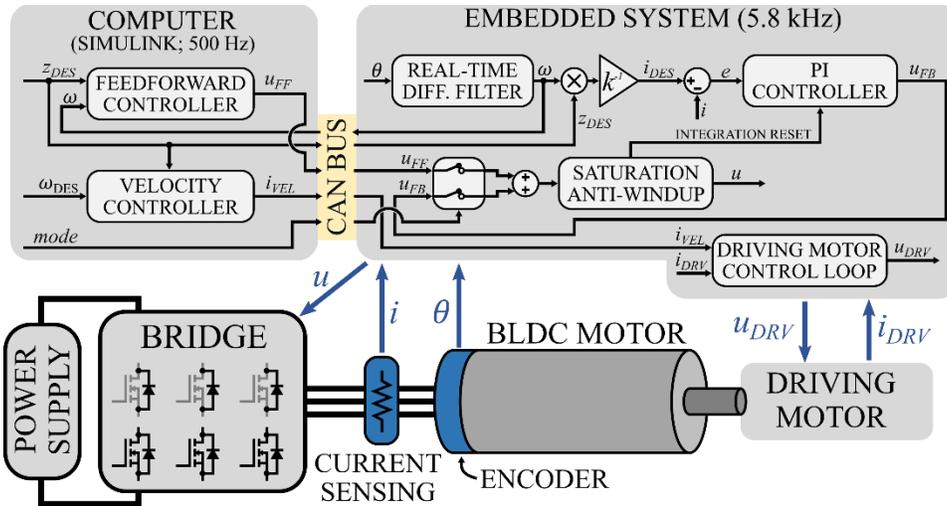


Fig. 3.7. Schematic representation of the control system with computer and embedded system sharing the control tasks.

When carrying out model validation experiments, the second motor was disconnected from the control electronics output and powered by a secondary power supply. This electrically isolated the BLDC motor under study, which facilitates characterization of generated power. Note that the main power supply still provides energy to power other components of the control electronics, but that power consumption is mostly invariant, and therefore is subtracted out when characterizing generation.

3.5.1 Model Validation

A first set of experiments were performed to validate the developed model and evaluate the tradeoffs associated with controllability and energy generation. The driving motor was powered by the secondary power supply. Two voltage and two switching frequency configurations were tested. First, the nominal configuration of 45 kHz and 24 V; then a lower voltage setting with the same PWM frequency: 45 kHz and 10 V; and finally, a lower switching frequency setting with the original voltage: 5 kHz and 24 V. The experimental procedure consisted of setting the velocity reference to one of {25, 50, 75, 100%} of $\omega_D = 700 \text{ rad/s}$ and the passive PWM duty cycle to one of {2, 4, 6, ... 100%}. In total, six hundred settings of velocity, duty cycle, and hardware configuration were tried, each for four seconds. Motor current, EPS current, and velocity were recorded at 500 Hz. The experimental routine included pauses between trials to prevent motor overheating. Trials were interrupted if motor current reached 12 A. The period corresponding to the transient of the velocity controller was removed from the analysis. Results are presented in Fig. 3.8, where experimental data is compared to the behavior predicted by the model, according to (3.11), (3.12), and (3.24). Power generated is calculated by taking the difference in average EPS current between idle and operating modes and multiplying that by the nominal EPS voltage.

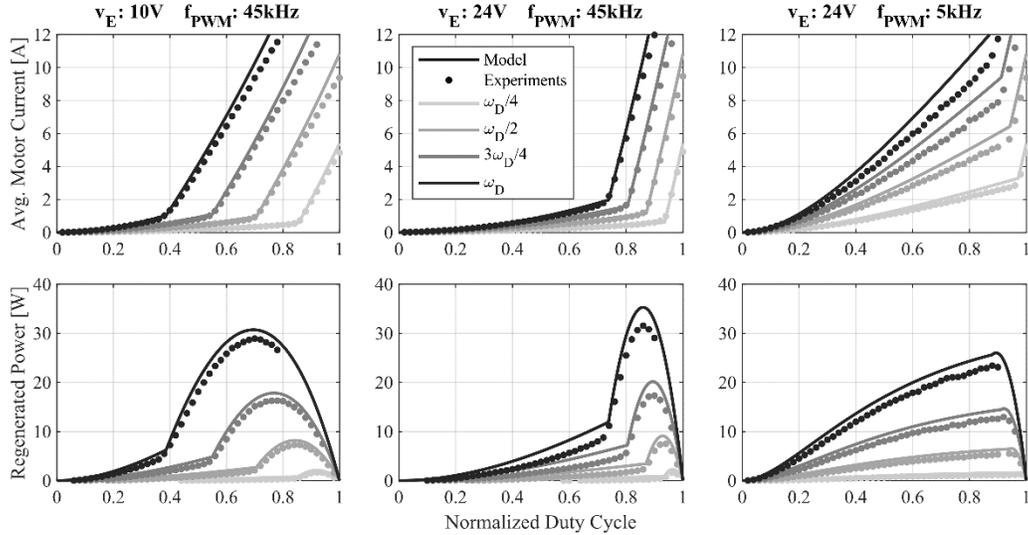


Fig. 3.8. Model validation results.

Experimental results closely match the developed model in all three configurations, demonstrating model validity. Note particularly that the change in behavior between the continuous and discontinuous current regime is accurately captured (and has a substantial nonlinear effect on the PWM control of torque). The tradeoff between controllability and generation is apparent: the nominal case in the center plots has the worst level of controllability but achieves the highest levels of power into the EPS. Lowering the EPS voltage (plots on the left) results in improved controllability with slightly lower power. Likewise, lowering the switching frequency greatly improved controllability but again power is reduced. The error between experiments and model in almost every trial indicates that EMD and generated power are slightly overestimated by the model. Note that experimental data is not available for motor currents above 12 A, for that reason some power curves are incomplete. Relative errors in motor current behavior are between 8 and 15% for duty cycles above 20%. Errors are consistent across duty cycle and velocity. Relative errors in energy generation are higher, between 10 and 30% (for power above 5 W).

3.5.2 Controlling Impedance

A second set of experiments were performed to demonstrate and compare the efficacy of the three strategies for impedance control (i.e., feedforward controller, feedback controller, and combined). These experiments consisted of three tracking tasks (the first and second described here, the third subsequently). The first task consists of tracking a 0.5 Hz sinusoidal EMD reference under constant velocity. The second, a complementary task, consists of tracking a constant EMD reference under sinusoidally varying velocity with the same frequency. In the first case, the EMD sinusoidal reference swings between 5% and 85% of maximum EMD and velocity is set to $\omega_D/2 = 350 \text{ rad/s}$. In the second task, velocity varies between $3\omega_D/8$ and $7\omega_D/8$ and EMD is set to 50% of the maximum range. The maximum setting $z = -b_E$ is purposefully avoided as it was observed in the model validation experiments that the model overestimates impedance. Results are presented in Fig. 3.9: velocity tracking, EMD tracking, and passive switching control duty cycle.

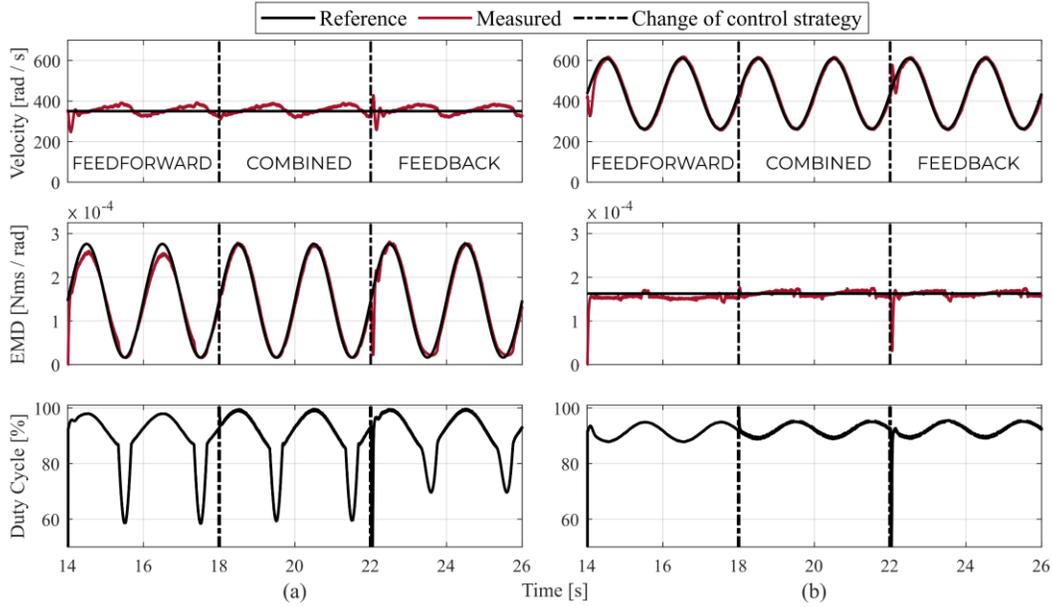


Fig. 3.9. Two EMD control experiments. In (a) the control task consists of tracking a sinusoidal EMD reference under constant velocity. The complementary task is shown in (b).

It can be observed that all three methods perform similarly well in tracking EMD. The feedforward method exhibits tracking error in both tasks, but the error vanishes when the PI controller is enabled. Performance degrades in the first task, Fig. 3.9 (a), when only the feedback controller is used. The PI controller alone is too slow to reach the low duty cycle setting due to the control linearity issue. That effect will be explored in more depth with the third task. All three methods are robust to velocity variations. A transient can be seen in the transition between combined and feedback-only control modes; the PI controller takes the system back to the desired behavior when the feedforward term is removed.

The third task consisted of tracking a square wave EMD reference with a fixed velocity setting of $3\omega_D/8$. The square wave switches between 1 and 75% of maximum EMD. The rise and fall dynamics are shown in detail in Fig. 3.10 for the three control strategies along with the control variable. The feedforward strategy presents desirable dynamic behavior, responding quickly and with no overshoot, however with noticeable steady-state error. The feedback strategy achieves desirable tracking but demonstrates a slow response in the discontinuous region followed by overshoot in the continuous region. This behavior is a result of control nonlinearity as presented in the previous section. The combined strategy accomplishes both desirable tracking and a quick response, preserving the best characteristics of the other two.

3.6 Discussion

The control approaches proposed here allow for either torque or EMD tracking with little modification to the controllers. We opted to control motor impedance (EMD) instead of torque directly in this work because short-circuiting the motor leads results in a fixed impedance as the limiting behavior (i.e., $z = -b_E$). That is in contrast with typical commutation schemes in which duty cycle can be interpreted as the application of a proportion of the rail voltage to the motor.

Only one BLDC motor was experimentally tested in this work. It is important to note that the observed nonlinearities and the tradeoffs related to achievable behavior, energy generation, and controllability are

generalizable to any DC motor. The choice of motor must be made based on power capability, application velocity, and desired impedance. The nonlinearity resulting from discontinuous current behavior will be present as long as the application velocity is below the system's ω_{MAX} . Similar experimental observations are presented in literature: 1) in [19], [66] the experimentally-obtained bounds of achievable passive behavior are equivalent to the model-based results presented here; 2) an experiment presented in [20] demonstrates the relationship between controllability and PWM frequency; and 3) the energy regeneration experiments in [57] demonstrate equivalent behavior to what was found in the model validation experiments here. As such, the work presented here provides a model-based explanation of published, experimentally observed, highly nonlinear effects in the torque-control of motors using identical or equivalent switching schemes.

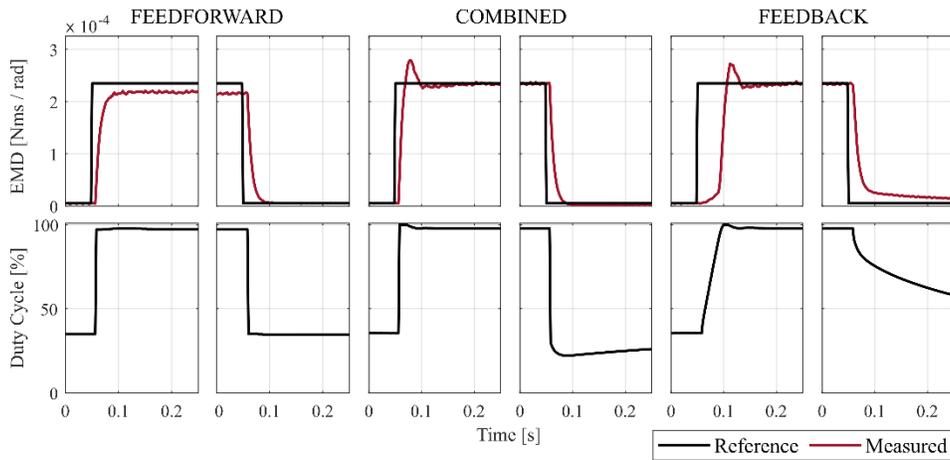


Fig. 3.10. The third EMD control experiment consisted of tracking pulses of EMD reference behavior using the three controller configurations (feedforward alone, feedback alone, and feedforward and feedback combined).

Linear approximations for PWM behavior could have been used to simplify the model. However, we wished to investigate the effect of PWM periods in the same order of magnitude of and even lower than the electrical system time constant, which required the complete exponential model.

Motor current was measured by the control electronics by detecting each commutation sector. At every sector, motor current is measured by averaging the currents in the two phases that align most closely with the magnetic field and ignoring the third phase. We note, however, that the proposed commutation scheme does not rely on the commutation sector. This current measurement method may contribute to some of the model error due to the ignored motor phase.

3.7 Conclusion

BLDC motors are generally more power dense than electrically-modulated brakes because, unlike brakes, they can transduce mechanical power into an electric energy storage device. This paper considers the use of BLDC motors as brakes, and makes several contributions which have not previously been presented, including: 1) a model of a BLDC motor as a brake which includes the highly nonlinear contributions of switching effects; 2) a model-based control approach that provides high-fidelity strictly passive control of impedance; and 3) a discussion of control trade-offs, some counter-intuitive, associated with the relationship between storage voltage and PWM frequency, as they affect range of controllable impedance, power generation, and control linearity.

4 On the Benefits and Limitations of Modulated Damping with Passive Motor Control

This chapter has been submitted in the Fall of 2022 for review as a journal paper in ASME’s Journal of Dynamic Systems, Measurement, and Control and has been accepted for publication in the Spring of 2023. The paper is currently in production and should be published in the coming months. We opted to use the same outlet as the work presented in the previous chapter because this work complements the previous. Specifically, while the previous work demonstrated control of emulated damping behaviors in a strictly passive manner, this work compares the performance of such methodology relative to a conventional, active motor control approach. Results are presented based on a theoretical analysis, simulations, and experiments.

4.1 Introduction

In emulated damping control, the torque imposed by a current-controlled motor is computed as a function (typically linear) of the measured velocity. This control structure is commonly employed in mechatronic control systems, either as an emulation of physical damping, or as a component of a servo-control system. Although this control component is theoretically strictly energetically passive, it is not guaranteed to be so when implemented with conventional motor control strategies. As has been well established, a motor emulating braking behavior in this manner can in fact add energy to the system through so-called energy leaks (e.g., filter dynamics in velocity measurement, sampling discretization, measurement uncertainty, etc.) which renders the control system prone to instability and behavioral inaccuracy [52]–[54], [72]–[74]. Those considerations are particularly relevant to applications that require high performance impedance control, high level of behavior fidelity, or in which control instability would be problematic.

As an alternative to the conventional damping control feedback structure described above, the authors have previously described a strictly passive control approach for achieving a desired torque from a brushless DC (BLDC) motor [5]. The passive approach employs a combination of feedback and feedforward controllers and leverages what was referred to as “passive switching” to operate the motor as a brake, in which damping torques are a result of the motor’s constitutive behavior. The control method presented in [5] is shown to achieve high-bandwidth and high-accuracy tracking of damping, which are not treated in other works in which similar or equivalent switching techniques are employed (e.g., to control knee prostheses [19], [20], [55], [66], haptic interfaces [67], [68], [75], and electric vehicles [57]–[59]). However, the prior work lacks a comparison of the proposed passive motor control approach with a conventional approach to emulated damping. That is, even though the results of [5] were positive and encouraging, the advantages of stability and behavioral fidelity attributed to the proposed technique were not as clear.

In the present work, the authors present theoretical, simulated, and experimental results that compare the relative characteristics of damping emulation via a conventional control method, relative to the proposed passive control method. In the conventional control implementation, the BLDC motor current is controlled with a feedback proportional-integral (PI) controller and a complementary switching bridge with block commutation. Experiments are performed with both systems to control damping behavior with different magnitudes and asymmetries (i.e., different settings of damping for forward and reverse directions of rotation). Obtaining high-bandwidth asymmetric damping passively from a motor may have particular utility for haptics applications (e.g., collisions in a virtual environment [76]), control of lower-leg prostheses (in which asymmetric behavior is typically implemented with mechanical components [9], [10], [77]–[79]), as well as active vehicle suspension systems [80], [81].

The passive approach is shown to provide advantages of control robustness and absence of energy leaks, as demonstrated by a stability analysis of conventional closed-loop damping control, followed by simulations and experiments. Experimental results are analyzed and discussed. The results indicate that passive switching is an attractive method for motor control (providing higher behavior fidelity, smooth operation, and control robustness) if the control task requires mechanical impedances within the constitutive limitations. Finally, this work also includes an improved formulation of the feedforward controller, which is more computationally tractable and efficient relative to the one described in [5].

4.2 Passive and Conventional BLDC Motor Control

In this section we review motor current control techniques with the conventional complementary switching approach as well as with the proposed passive switching scheme. Models are developed for both cases for subsequent stability analyses and simulation. A comment on BLDC electronic commutation (which passive switching does not require) is presented for completeness.

4.2.1 Torque Control and Switching Techniques

A simplified electric model that describes the behavior of a BLDC motor shown in the circuit schematic in Fig. 4.1 (a) and given by the differential equation in (4.1).

$$v_M = L \frac{di_M}{dt} + Ri_M + v_{EMF} \quad (4.1)$$

where v_M is the voltage at the motor leads for a motor with inductance L and resistance R , v_{EMF} is the back electromotive force (back-EMF) induced in the motor armature due to the rotation of the permanent magnet relative to the windings, and i_M is the resulting motor current. Electronic sector commutation is disregarded at this point. Motor torque and back-EMF are given by $\tau_M = ki_M$ and $v_{EMF} = k\dot{q}$, respectively, where k is the torque constant and τ_M and \dot{q} are respectively the motor torque and rotational velocity. Given the relationship between torque and current, torque control is achieved by controlling motor current.

In traditional motor control schemes, control of current i_M is achieved by PWM modulation of motor phase voltage v_M with an H-bridge structure and current sensing. Note that v_{EMF} is an input disturbance to the feedback controller. With the passive switching approach, current control is achieved by effectively modulating the application of v_{EMF} . The key difference between the two methods is that in passive switching energy can only be provided to the system by v_{EMF} as a result of motion, while in traditional methods both v_M and v_{EMF} can provide energy to the electric system. In the following sections, the hardware implementation of complementary switching (representing conventional control methods) and passive switching will be discussed in greater detail.

4.2.1.1 Complementary Switching

As is shown in Fig. 4.1, an H-bridge structure with PWM can be used to apply a voltage of $v_M = \pm v_E$ to the motor leads where v_E is the battery voltage. Two diagonally opposing MOSFET switches are engaged during the forced portion of PWM, where positive voltages correspond to the main diagonal and negative voltages correspond to the secondary diagonal (respectively shown as cases (iii) and (iv) in Fig. 4.1 (b)). Both low-side switches are active during the unforced portion of PWM, which corresponds to $v_M = 0$ at the motor leads which is shown in Fig. 4.1 (b) case (i). By averaging between forced and unforced PWM phases, an arbitrary average voltage (within the limits of the battery voltage) can be obtained at a given PWM command.

The duty cycle command is determined by a feedback controller and is combined with electronic commutation to apply the command to the MOSFET gates in the appropriate H-bridge diagonal. For a duty cycle $u \in [-1,1]$, the effective control action can be modelled as a continuous system with an equivalent applied voltage $v_M = uv_E$. A PI control architecture is commonly applied to this system in order to achieve good tracking of reference current i_{REF} and disturbance rejection; the output of the controller is given by (4.2)-(4.4) below, where e^{CS} is the tracking error for the complementary switching (“CS”) case, e_I^{CS} is the error integral, and k_p and k_I are the proportional and integral gains, respectively.

$$e^{CS} = i_{REF} - i_M \quad (4.2)$$

$$\dot{e}_I^{CS} = e^{CS} \quad (4.3)$$

$$u^{CS} = k_p e^{CS} + k_I e_I^{CS} \quad (4.4)$$

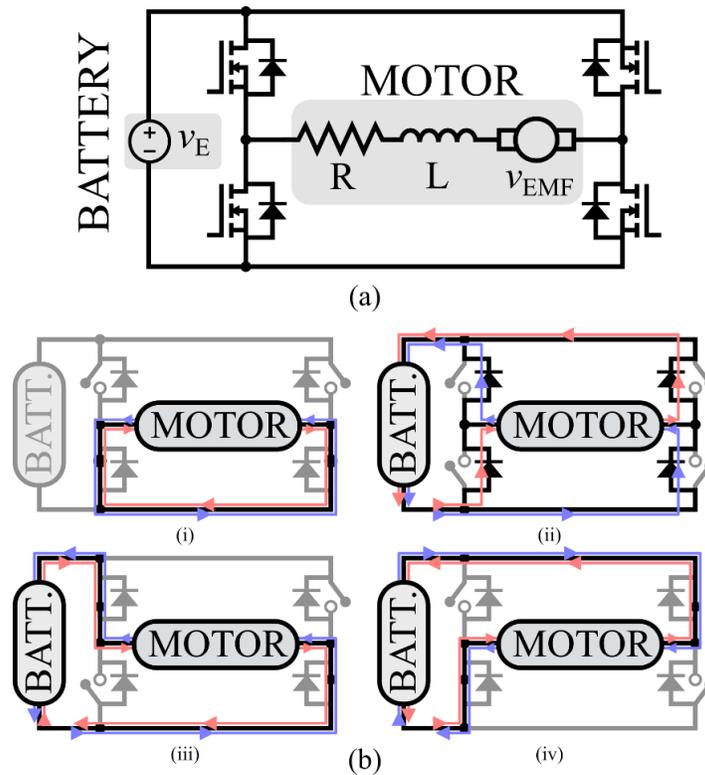


Fig. 4.1. Simplified motor model within an H-bridge circuit in (a) and four switching configurations with highlighted positive and negative current paths in (b): (i) unforced PWM phase for complementary switching and forced PWM phase for passive switching, (ii) unforced PWM phase for passive switching, (iii) and (iv) forced PWM phase for complementary switching with positive and negative voltage commands, respectively. Note that the battery never delivers power to the system with passive switching (i.e., current flow through the battery always opposes the battery voltage with passive switching).

4.2.1.2 Passive Switching

Passive switching is achieved with a standard H-bridge by quickly alternating between shorting and opening the motor leads. Shorting the leads is the forced PWM phase in passive switching but corresponds to the same behavior as the unforced phase of PWM in complementary switching (low-side MOSFET switches are active, Fig. 4.1 (b) case (i)). The open leads state occurs in the unforced PWM phase of passive switching and is implemented by simply keeping all MOSFET switches open. During the unforced phase, energy accumulated in the inductor is delivered to the battery through diodes, as shown in Fig. 4.1 (b)

case (ii). The diodes prevent energy flow from the battery back to the motor. As such, the function of the battery in passive switching is analogous to an energy dissipator in rheostatic braking; with the benefit of energy regeneration into the battery, which can be useful for other control tasks.

Since motor current is strictly a result of v_{EMF} , the duty cycle u can be interpreted as an application of between 0 to 100% ($u \in [0,1]$) of v_{EMF} on the RL circuit, as in (4.5). The resulting i_M dynamics (and therefore the resulting torque) reacts to the multiplicative combination of the control signal u and back-EMF voltage.

$$0 = L \frac{di_M}{dt} + Ri_M + uv_{EMF} \quad (4.5)$$

Note that battery voltage does not appear in this simplified model but has important implications to controllability. While back-emf voltage is the only driver of motor current during the forced portion of PWM, the opposing battery voltage dominates the dynamics during the unforced portion (driving $|i_M| \rightarrow 0$ but never changing the sign of i_M due to series diodes). The choice of both PWM frequency and battery voltage has critical implications to control sensitivity and efficiency of energy regeneration. Likewise, the simplified model ignores nonlinear effects such as diode clipping (i.e., discontinuous current regime) and motor velocity limitations when $|v_{EMF}| > v_E$ (diodes become forward-biased). Despite these shortcomings, the model is useful to understand the resulting motor behavior with passive switching. A more complete model of the forced and unforced PWM phases of passive switching as well as a discussion on aspects of controllability and regeneration efficiency are presented in [5].

A slightly modified PI controller structure is used for passive switching since only magnitude of current can be affected (i.e., the scheme does not allow control of directionality, motor current always opposes velocity). The tracking error is defined as the difference between reference and actual current magnitudes as in (4.6) with the superscript “PS” indicating passive switching. The error integral and control output are analogous to the complementary switching case in (4.3)-(4.4).

$$e^{PS} = |i_{REF}| - |i_M| \quad (4.6)$$

A feedforward controller is used in combination with the feedback controller to address the challenging controllability. The PI controller will be considered in isolation in this theoretical analysis, while the feedforward controller will be presented subsequently in section 4.4.2.

4.2.2 Sector Commutation

Unlike what is represented in the simplified circuits of Fig. 4.1, BLDC motors have three electrical phases and typically rely on electronic commutation to energize the phases in a manner that aligns the magnetic field generated in the windings with the magnetic field of the permanent magnet in the rotor. Commutation can be implemented continuously using field-oriented control (FOC) techniques or, more commonly, can be done discretely with low-resolution absolute position information from Hall Effect sensors (referred to as block commutation). The rough discretization of rotor position into sectors causes undesirable torque ripple but is significantly simpler to implement than the smoother FOC alternatives (from the perspective of sensing, computational resources, and algorithmic complexity).

For the purposes of this paper, the conventional control method with complementary switching was implemented with standard block commutation. Passive switching, on the other hand, does not require any form of commutation. Since motor current is induced via back-EMF, the motor phases that align with

the permanent magnet are subject to the largest portion of induced currents. In that way, continuous commutation occurs by construction with passive switching. This characteristic results in noticeable absence of vibration which is confirmed in the experimental section of this work.

4.3 Application of passive switching to damping emulation

In this section we introduce a closed-loop damping control structure (a subset of the more general impedance control approach [50], [51]) and show the advantageous robustness that passive control provides by analyzing the stability of the passive approach with linear and nonlinear analyses, and contrasting to instability for conventional implementations resulting from so-called energy leaks.

4.3.1 Motor-based Damping Emulation

Many applications employ variable (linear or nonlinear) damping behaviors, including microprocessor-controlled knee prostheses, haptic displays, adaptive automobile suspensions, and vibration control in manufacturing processes. The variable damping behaviors can be implemented using different approaches, among them motor-controlled emulated damping. Consider a rotational mechanical system represented in (4.7) with intrinsic inertia I_0 and damping B_0 , subject to a motor torque τ_M and external torque τ_E (which represents the interaction force between the mechatronic system and the environment). The resulting angular motion is denoted by q with derivatives for angular velocity and acceleration.

$$I_0\ddot{q} + B_0\dot{q} = \tau_E + \tau_M \quad (4.7)$$

To achieve motor-controlled emulation of damping, velocity is measured, and motor torque is controlled via current-control feedback. In the particular case of viscous damping, a motor torque reference is computed according to (4.8), where τ_{REF} is the motor torque reference, \hat{q} is the measured version of the angular velocity \dot{q} , and B is the desired damping coefficient prescribed by the control task. The role of a motor controller is therefore to make motor current (and therefore motor torque) track the reference τ_{REF} .

$$\tau_{REF} = -B\hat{q} \quad (4.8)$$

In a more general case, the reference torque can be any function of measured velocity such that the product of reference torque and velocity is negative. The control structure of emulated damping is shown in Fig. 4.2. In the ideal case of $\tau_M = \tau_{REF}$ and $\hat{q} = \dot{q}$, the emulated dynamics become as in(4.9), where the emulated damping behavior effectively adds to the intrinsic damping of the system.

$$I_0\ddot{q} + (B_0 + B)\dot{q} = \tau_E \quad (4.9)$$

To summarize this section, this paper compares the characteristics and performance of motor-controlled emulated damping, as shown in Fig. 4.2 (a), with two different motor control approaches: 1) a conventional control approach, which employs a complementary switching scheme, which results in a “motor control” block in Fig. 4.2 (a) with behavior schematically represented in Fig. 4.2 (c); and 2) a passive control approach, which uses the same general control block diagram shown in Fig. 4.2 (a), but with a passive switching approach, which results in a “motor control” block with behavior in Fig. 4.2 (d). The structure of both the complementary and passive switching schemes is described in sections 4.2.1.1 and 4.2.1.2, respectively.

4.3.2 Robustness of Damping Behavior

Due to several implementation factors, however, including control dynamics in the current tracking loop and filtering and computational dynamics in the velocity measurement, damping emulation with the conventional control method can cause instabilities to occur in the controlled system, as discussed by other researchers (e.g., [52]–[54], [72], [73]), and also briefly presented here.

Assume velocity is estimated by real-time differentiation of position, such that a dynamic exists between \dot{q} and \hat{q} , which is modeled here as a first order system with time constant T_F . Using the electrical and mechanical models respectively in (4.1) and (4.7), assuming the PI current control structure with conventional motor control as presented in (4.2)–(4.4), and damping emulation with $i_{REF} = -Bk^{-1}\hat{q}$, the complete system dynamics can be represented by the state-space model in (4.10).

$$\frac{d}{dt} \begin{bmatrix} \dot{q} \\ \hat{q} \\ e_I^{CS} \\ i_M \end{bmatrix} = \begin{bmatrix} -B_0/I_0 & 0 & 0 & k/I_0 \\ 1/T_F & -1/T_F & 0 & 0 \\ 0 & -B/k & 0 & -1 \\ -k/L & -k_P v_E B/(kL) & k_I v_E/L & -(R + k_P v_E)/L \end{bmatrix} \begin{bmatrix} \dot{q} \\ \hat{q} \\ e_I^{CS} \\ i_M \end{bmatrix} + \begin{bmatrix} 1/I_0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \tau_E \quad (4.10)$$

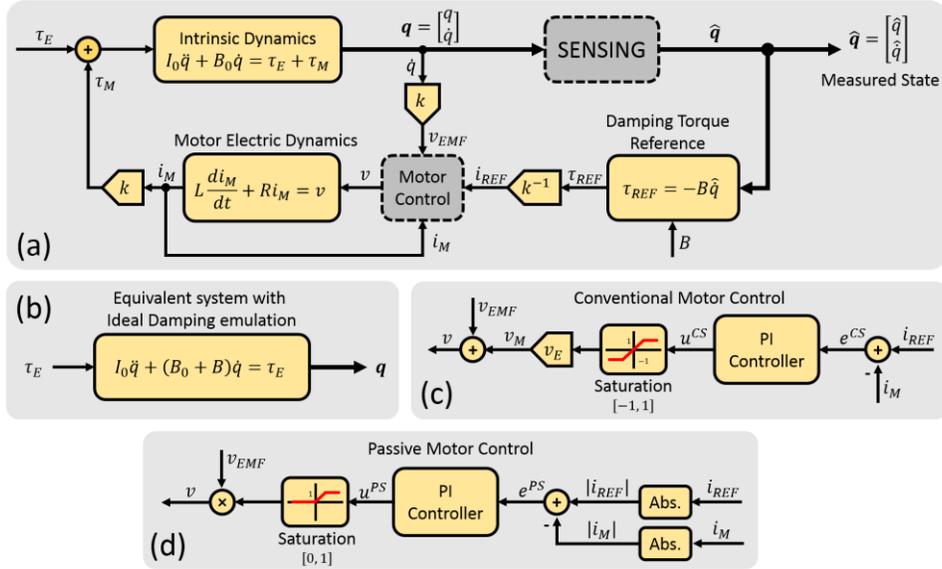


Fig. 4.2. Damping control schematic in (a) and equivalent closed-loop system in (b), assuming ideal sensing (no errors, no added dynamics) and perfect tracking of damping torque reference. In this paper, the term “conventional control” refers to the control structure shown in (c) which replaces the block entitled “Motor Control” in the overall damping control structure in (a), whereas the term “passive control” refers to the control structure in (d).

Assuming an input torque τ_E and output system velocity \dot{q} , the root locus as a function of B is shown in Fig. 4.3, which demonstrates the limits of stability for a given limiting value of B . The exact value of B in the limit of stability depends on the control gains, the real-time differentiating filter time constant, the period of the digital control loops, and other implementation parameters. The limit of stability for the experimental system used in this work is found experimentally in section 4.5.1.

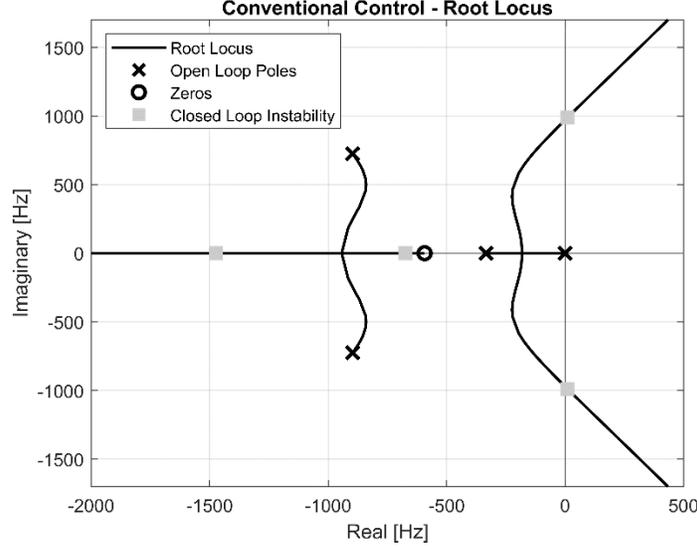


Fig. 4.3. Root locus plot demonstrating the instability of feedback damping control with conventional motor control techniques with respect to damping parameter B .

A similar analysis can be conducted for passive motor control using the relationship in (4.5) and the control error as defined in (4.6). Note that the mechanical system dynamics and behaviors of \dot{q} and \hat{q} remain unchanged. Motor current for damping control with passive switching is given by the nonlinear differential equation in (4.11), where $\text{sat}(u^{PS}, 0, 1)$ is a saturation function with low-level limit 0 and high-level limit 1. The behavior of the PI controller with respect to the system state is represented in (4.12).

$$\frac{di_M}{dt} = -\frac{R}{L}i_M - \frac{k}{L}\text{sat}(u^{PS}, 0, 1)\dot{q} \quad (4.11)$$

$$\dot{u}^{PS} = k_P \left[-\frac{B}{kT_F}(\dot{q} - \hat{q}) \text{sgn} \hat{q} + \left(\frac{R}{L}i_M + \frac{k}{L}\text{sat}(u^{PS}, 0, 1)\dot{q} \right) \text{sgn} i_M \right] + k_I \left[\frac{B}{k}|\hat{q}| - |i_M| \right] \quad (4.12)$$

There is some additional complexity in analysis due to the nonlinearity of saturation, product terms, and absolute values. The homogeneous version of the system ($\tau_E = 0$) contains equilibria at $\dot{q}^* = \hat{q}^* = i_M^* = 0$ and arbitrary u^{PS} . That is to be expected since motor current requires velocity to be imposed; the control signal can assume any value at the equilibrium because $\dot{q} = 0$ implies $i_M = 0$. Stability of the system (about the origin as an operating point) can be proven with the Lyapunov-Malkin theorem [82], in which the pseudo-state u^{PS} and nonlinear terms are separated from the portion of the system which is strictly linear and stable, as presented in (4.13).

$$\frac{d}{dt} \begin{bmatrix} \dot{q} \\ \hat{q} \\ i_M \end{bmatrix} = \overbrace{\begin{bmatrix} -B_0/I_0 & 0 & k/I_0 \\ 1/T_F & -1/T_F & 0 \\ 0 & 0 & -R/L \end{bmatrix}}^A \begin{bmatrix} \dot{q} \\ \hat{q} \\ i_M \end{bmatrix} + \overbrace{\begin{bmatrix} 0 \\ 0 \\ -\frac{k}{L}\text{sat}(u^{PS}, 0, 1)\dot{q} \end{bmatrix}}^g \quad (4.13)$$

The Lyapunov-Malkin theorem requires that all eigenvalues of A be stable and that both $\mathbf{g}(\dot{q}, u^{PS})$ and \dot{u}^{PS} vanish at $\dot{q} = \hat{q} = i_M = 0$. Since both conditions hold, the theorem states that the origin is an asymptotically stable equilibrium point for the reduced system in (4.13) and that u^{PS} is likewise stable (as

would be expected of the integral action of a PI controller, it does not necessarily converge to zero when the error signal is zero). Note that the stability of the origin does not depend on B .

4.3.3 Simulation Results

The representative emulated damping system illustrated in Fig. 4.2 was simulated with both the conventional and passive motor controllers for two simulation cases. Both systems were subject to the same sinusoidal torque input τ_E in each simulated case, and both systems employed the same motor parameters, specifically set to the nominal parameters of the drive used for the (subsequently described) experimental validation.

In the first simulated case, gains for the PI controllers were adjusted independently for each system to achieve adequate motor torque tracking with similar magnitude of errors. The input torque signal is a sinewave with frequency of 5 Hz and desired damping ratio of $B = 1.4\text{ Nm} \cdot \text{s}/\text{rad}$. Motor torque tracking results are shown in Fig. 4.4. The top plots show the comparison of steady-state motor torque, relative to the torque which would be obtained by an ideal damper ($-B\dot{q}$). The larger bottom plot shows motor torque against velocity. The ideal behavior is a line passing through the origin which both systems approximate. In the case of traditional motor control, the combination of measurement and control lag (as can be seen in the comparison with an ideal damper) traces an ellipsoid which crosses into the active quadrants of the power plane. In the haptics literature, that phenomenon is termed an “energy leak” because the motor adds energy to the mechanical system, which is a contributing factor to instability of the larger system. Passive switching guarantees that the motor behavior is strictly dissipative (the red curve converges to the origin when velocity changes sign) but does not guarantee behavioral accuracy, which depends on well-adjusted controller gains and accuracy of velocity estimation.

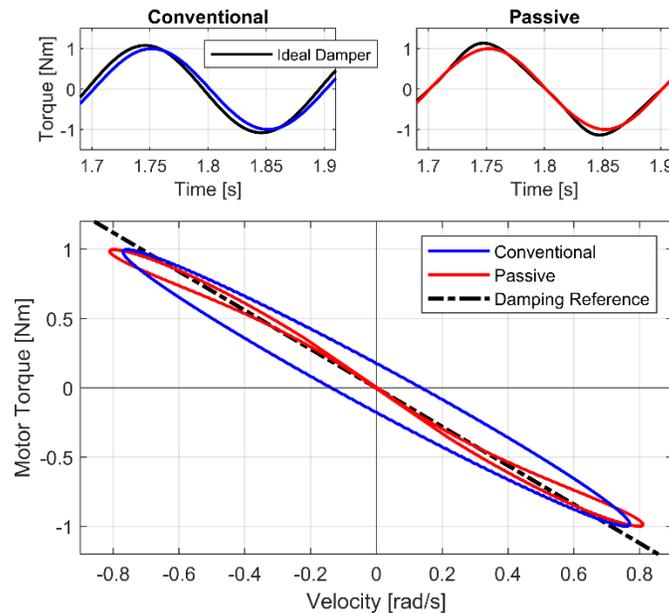


Fig. 4.4. Simulation results demonstrating the difference in achieved damping behavior between conventional and passive motor control compared to an ideal damper. The behavior obtained with conventional control crosses into the active quadrants of the power plane.

In the second simulated case, the PI controller gains were raised and purposefully kept identical for both systems to evaluate system behavior in the case of inadequate control gains. Input torque is a sinewave at 5 Hz. Desired damping starts at a low value of $B = 0.1 \text{ Nm} \cdot \text{s}/\text{rad s}$ and then at $t = 1 \text{ s}$ the command is increased by an order of magnitude to $B = 1.5 \text{ Nm} \cdot \text{s}/\text{rad}$. The results in Fig. 4.5 demonstrate the instability of the conventional strategy to the higher setting of commanded damping. The higher control gains adversely affect the performance of the passive system, but do not cause instability. Motor impedance (ratio of motor torque to motor velocity) in the top plot of Fig. 4.5 is calculated by dividing motor torque by motor velocity (the calculation diverges when velocity changes sign).

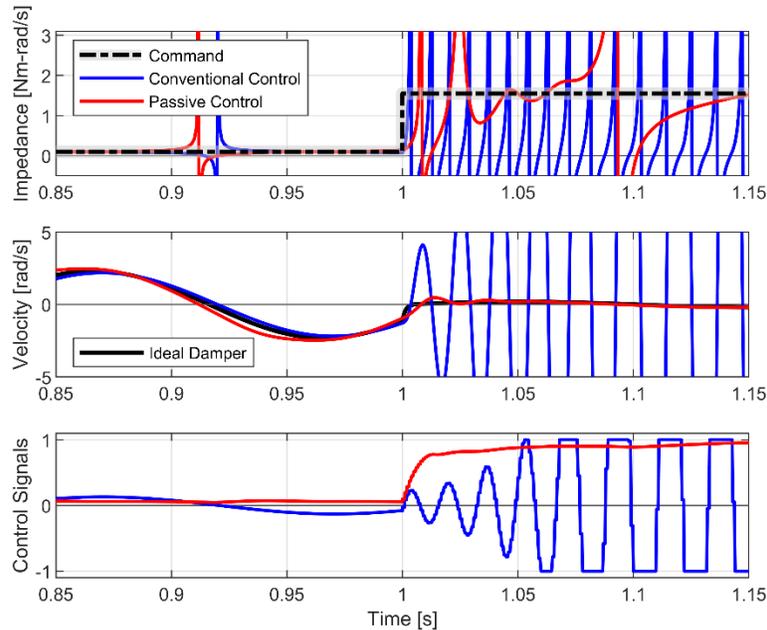


Fig. 4.5. Conventional and passive control systems, both (intentionally) configured with gains that are too high to provide impedance tracking. The top plot shows the desired and actual motor impedances; the middle plot shows velocity; and the bottom plot shows the control signal to the motor. Note that the actual impedance must go to infinity during velocity zero-crossings, by definition. In the case of conventional control, a step change in damping command (shown at $t = 1 \text{ s}$) results in instability. Impedance tracking performance is also adversely affected for the control passive case, but no instability occurs (as indicated in the velocity plot).

4.4 Implementation

4.4.1 Experimental Apparatus

An experimental setup was constructed to examine the performance differences between conventional and passive approaches to damping emulation. The apparatus was constructed to perform experiments with a Maxon EC-4pole 22 90W BLDC motor (model no. 32318). A planetary gearhead (Maxon GP 22 HP) providing 72:1 reduction was installed at the motor output, and a 10 cm lever arm affixed to the gearhead shaft to allow an operator to interface with the system and backdrive the motor. The setup includes a custom embedded system with BLDC current control hardware which is powered by a bipolar power supply providing 16 V to the system. Passive and complementary switching motor control schemes were implemented in the embedded system so control methods could be compared. A photograph of the apparatus is presented in Fig. 4.6. Various parameter values are presented in Table 4.1.

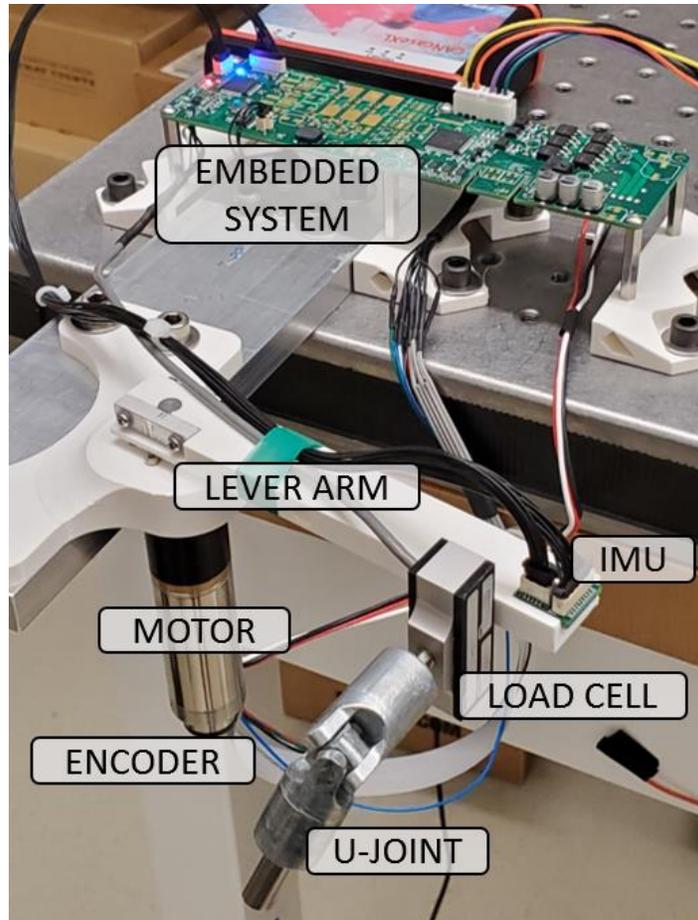


Fig. 4.6. Experimental apparatus. Photo shows the actuator assembly with motor, gearhead, lever arm, load cell, U-joint, and IMU sensor as well as the custom embedded system and CAN interface device in the background.

Table 4.1. Parameters of the experimental apparatus. Motor parameters (winding resistance, inductance, and torque constant) are presented as shown in the device datasheet; the power supply voltage is adjusted with a multimeter before the experimental procedures; the PWM frequency of the drive was selected in the firmware of the low-level microcontroller; the transmission ratio value was obtained from the spec sheet of the part provided by the manufacturer; the lever arm length is the nominal value used in the design of the rapidly-prototyped part; and moment of inertia, intrinsic damping, and coulomb friction terms were determined with a system identification procedure described subsequently.

Parameter	Symbol	Value
Motor winding resistance	R	0.527Ω
Motor winding inductance	L	0.0503 mH
Motor torque constant	k	0.014 Nm/A
Battery (power supply) voltage	v_E	16 V
PWM Frequency	f_{PWM}	45 kHz
Transmission ratio	n	$72.1:1$
Lever arm length	l	0.1 m
Moment of inertia	I_0	$0.002 \text{ kg} \cdot \text{m}^2$
Intrinsic damping	B_0	$0.013 \text{ Nm} \cdot \text{s/rad}$
Coulomb friction	τ_0	0.0173 Nm

The motor has a built-in 128 pulse encoder on its shaft and hall sensors to indicate commutation sector. The embedded system contains signal conditioning and acquisition circuitry for high-bandwidth measurement of motor current. A load cell rated for 100 N was separately calibrated and installed on the lever arm to record the interaction forces between operator and system. A U-Joint was installed between load cell and operator interface to isolate moments and forces which would interfere with the load cell measurements of tangential forces. An IMU sensor was also installed at the end of the lever arm to provide an external source (i.e., not used for control) of kinematic measurements.

Within the custom electronics are two microprocessors which implement the hierarchy of the damping emulation control structure: a general-purpose microcontroller (PIC32MZ EF) implements the outer damping emulation control loop which defines i_{REF} at 1 kHz , and a microcontroller optimized for signal processing (dsPIC33F) implements the inner current control loops (PI current controllers) at 5.8 kHz . The dsPIC33 is responsible for measurements of motor current and encoder while the PIC32 interfaces with the IMU and takes measurements from the load cell. The PIC32 also exchanges parameters and sensor measurements via a CAN interface with a computer running MathWorks Simulink Desktop Real-Time. Note that the computer is only used to configure the embedded system and record trial data, no real-time control tasks are implemented in the computer.

Parameters I_0, B_0, τ_0 presented in Table 4.1 are defined at the gearhead output. These values were found using the load cell and motor encoder to estimate angular velocity and acceleration; the motor was electrically disconnected but kept mechanically connected to the system. Parameter values were estimated with a least-squares regression procedure. Regression data was obtained by moving the lever back and forth with varying velocities for about 30 seconds (collected at 1 kHz). In the following experimental results, motor torque is presented in terms of the measured current ($\tau_M \approx kni_M$) and in a couple cases in terms of friction and inertia-compensated load cell force (F) as in (4.14).

$$\tau_M \approx lF - B_0\hat{q} - I_0\hat{\dot{q}} - \tau_0 \quad (4.14)$$

4.4.2 Passive Motor Control Implementation

The theoretical comparisons (between conventional and passive motor control) presented in Section 3 employed only PI-type feedback control for both approaches, in order to provide uniform comparisons; the experimental implementation of the passive controller, however, additionally employs a feedforward controller. This is because, unlike conventional motor behavior, passive motor behavior entails a substantial static nonlinearity, which is readily addressed using a feedforward controller as described in this section. Note that the passive motor control method described in [5] also employs a feedforward component, although the implementation was simplified in this paper. Specifically, the feedforward controller developed in the previous work consists of numerically inverting the nonlinear switching model, which is more computationally complex than is necessary for a typical microcontroller implementation.

The nature of the static nonlinearity in the passive motor behavior is indicated in the top row in Fig. 4.7, which shows steady-state motor current as a function of passive switching duty cycle and velocity. The plot in (a) shows for reference the model presented in (4.5) which was used for all analytical discussions in sections 4.2 and 4.3; in (b) is the switching model as developed (and experimentally validated) in the prior work [5] where motor current is represented by the average PWM current; and in (c) is the simplified model employed herein to reduce computational requirements. The bottom row of Fig. 4.7 contains three examples of current behavior (zoomed in to show individual PWM periods) which illustrate the two

distinct regimes of operation, namely the “continuous” and “discontinuous” current regimes. In the continuous current regime, current builds up from phase to phase, eventually reaching a steady-state PWM behavior. In the discontinuous regime, current returns to zero before the end of every period, diode clipping preventing it from changing sign, so it remains at zero until the following PWM period.

As presented in Fig. 4.7 (c), the continuous current regime can be adequately approximated by fitting a plane to data points obtained by the switching model strictly in that regime, while the discontinuous current regime is captured with a lookup table. The table contains duty cycle settings and is indexed by measured velocity (\hat{q}) and desired current (i_{REF}). Plane fitting of the continuous regime was performed by least squares using the relationship $i_M \approx \alpha_1 u + \alpha_2 \hat{q} + \alpha_3$, where the resulting parameters $\alpha_{\{1,2,3\}}$ are used to find the feedforward control duty cycle u_{FF} as described in (4.15).

$$u_{FF} = \alpha_1^{-1}(i_{REF} - \alpha_2 \hat{q} - \alpha_3) \quad (4.15)$$

Note that the discontinuous regime has limited current amplitude (around 2 A for the experimental system) and limited velocity ($|\dot{q}| < k^{-1}v_E$, after which the H-bridge diodes become forward biased by the elevated back-EMF), therefore the discontinuous regime is finite and can be recorded in the lookup table with as fine resolution as will fit in the memory of the microprocessor in the embedded system. The processor used for the experimental implementation presented here permitted current resolution of 5 mA between 0 and 2 A and rotor velocity resolution of about 9 rad/s between 0 and 1200 rad/s; the table occupies 56 kB of memory (2 bytes per entry). The table also has the function of indicating which current regime is expected for each setting of velocity and desired current. The special value of -1 indicates to the embedded application to calculate the duty cycle from the continuous regime plane fit. Similarly, if current reference and/or velocity are out of the bounds of the lookup table, the system assumes the continuous current regime is in effect and uses (4.15). This feedforward controller is implemented within the general-purpose processor to take advantage of the floating-point hardware unit and abundant free memory in the device.

Importantly, the feedforward controller is defined using only nominal parameter values, so the method relies on the parallel PI controller (as in section 4.2.1.2) to achieve zero-error tracking. Refer to the prior work for a complete discussion, with experimental results, on the role of each controller component.

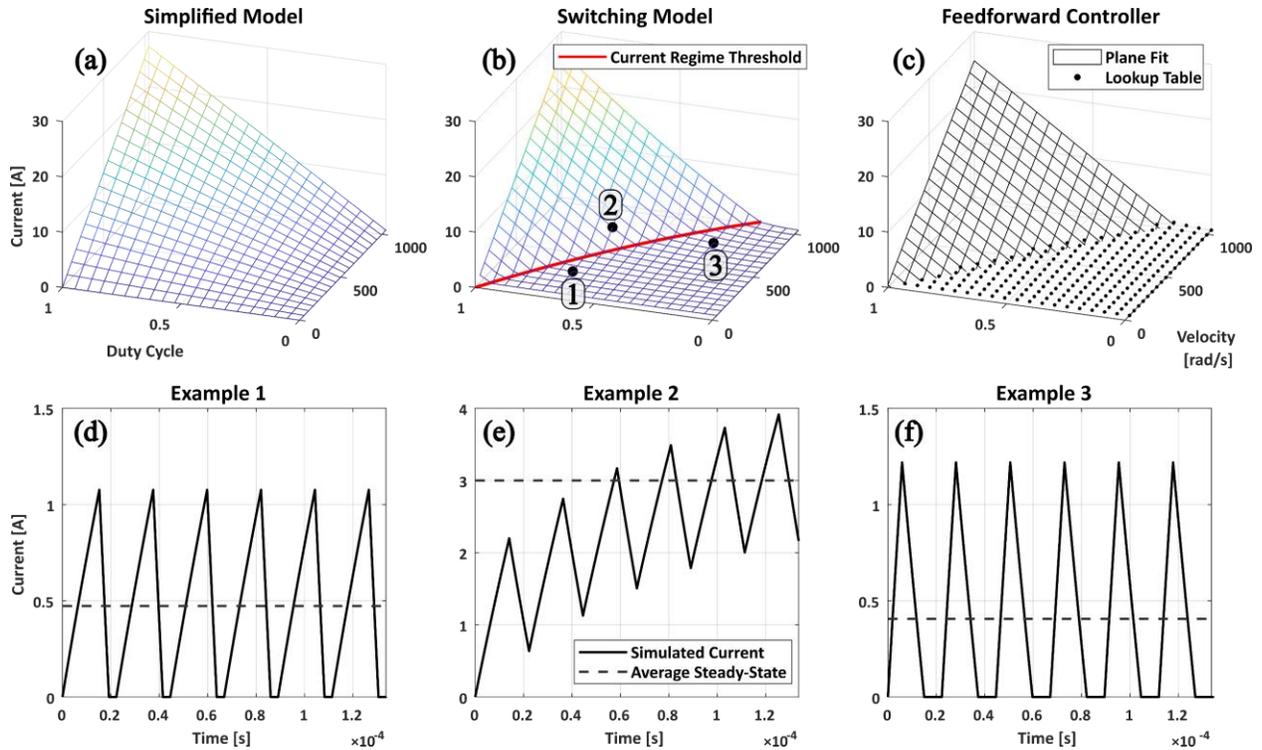


Fig. 4.7. In (a) is the steady-state magnitude of current for the simplified passive switching model as a function of duty cycle and velocity; in (b) is the same but for the complete switching model as developed in [5], highlighting the highly nonlinear effect of PWM with current discontinuity; in (c) is the proposed formulation of the feedforward controller with a plane approximating the continuous current regime region and a lookup table for the discontinuous region; in (d) and (f) are examples of simulated current behavior, showing individual PWM periods, in the discontinuous regime; and finally in (e) is an example of current behavior in the continuous regime. The duty cycle and velocity settings for each example are indicated by the numbered markers in (b).

4.5 Experimental Results and Discussion

4.5.1 Limit of Stability and Vibration

Emulation of increasing levels of symmetric damping was used to compare the performance of the two control schemes. In these experiments, a human operator moved the lever arm back and forth in a sinusoidal-like pattern at around 1 Hz , reaching velocities of up to 3.5 rad/s at the gearhead output (which equates to about 250 rad/s at the motor shaft); the damping reference B started at $0.115\text{ Nm} \cdot \text{s/rad}$ and was increased by the same amount until $1.146\text{ Nm} \cdot \text{s/rad}$. In these trials, the conventional damping emulation with complementary switching was no longer stable beyond $B \approx 0.69\text{ Nm} \cdot \text{rad/s}$, while the passive motor control remained stable for the entire range. For completeness, passive switching was also tested close to the short-circuit damping behavior and remained stable throughout.

The linear acceleration signal from the accelerometer axis aligned with the tangential direction of motion was used to indicate the level of vibration in each implementation (in units of 9.81 m/s^2 , or “ g ”). Noise energy as presented in Fig. 4.8 is the RMS value of the FFT of the signal for frequencies above 60 Hz (up to 500 Hz since data is acquired at 1 kHz). Note that motion happens at much lower frequencies, between 0.5 and 2 Hz . The conventional and passive motor controllers started with similar levels of vibration for low values of reference damping but quickly diverged: the noise level with conventional control rose as a function of emulated damping ratio, while the noise level as a function of emulated

damping with passive motor control fell. Note that increasing damping in a mechanical system should help attenuate vibration, which is what occurred in the passive control case. Note finally that passive control is able to achieve a much higher level of emulated damping relative to conventional control due to limits of stability. Joint velocities and raw accelerometer readings in time are also shown in Fig. 4.8 for trials with damping reference of about 0.57 Nms/rad , serving as examples of the analyzed data. It can be clearly seen that the conventional control strategy has higher magnitude of noise in the accelerometer signal, particularly when velocity changes sign.

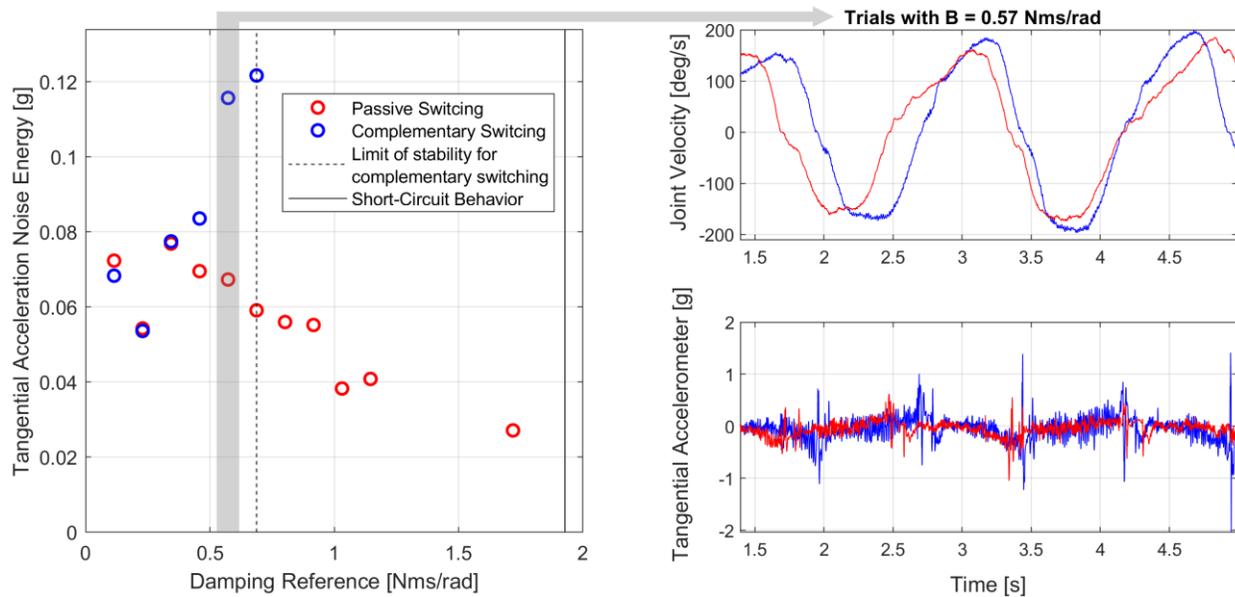


Fig. 4.8. Vibration metric as measured by FFT analysis of the tangential accelerometer sensor mounted in the lever arm, for passive and conventional control methods. Each data point in the left plot represents an individual trial with the fixed damping reference indicated in the x-axis. The vertical lines show the experimentally determined limit of stability for the conventional control method and the constitutive limitation of the passive control method. On the right, two plots show a 3.5 s time window of measured velocity and raw accelerometer readings for the trials with $B = 0.57 \text{ Nms / rad}$, demonstrating a higher magnitude of vibration for the conventional method under an equivalent velocity profile.

4.5.2 Asymmetric Damping Behavior

Another set of experiments was performed to examine the relative behavior of both control schemes for the emulation of asymmetric and unilateral damping references (i.e., directionally-dependent emulated damping). The operator moved the lever arm at about 1 Hz , performing rapid changes in direction of rotation. In Fig. 4.9, the first set of results are presented in terms of velocity (encoder and gyroscope) and torque (current and load cell) presented in a 3 second window; motor current and encoder data are also presented in a power plane. The positive velocity has a zero-torque reference ($B_{pos} = 0$) while the negative direction follows $B_{neg} = 0.573 \text{ Nm} \cdot \text{s/rad}$. Both motor torque as well as the load cell readings indicate that the conventional motor control method leads to larger magnitudes of behavior inaccuracy and energy leaks, since torque is positive for a positive velocity. The planetary gear and output shaft coupling contain some backlash, as can be observed by the velocity mismatch between encoder and gyroscope velocity readings.

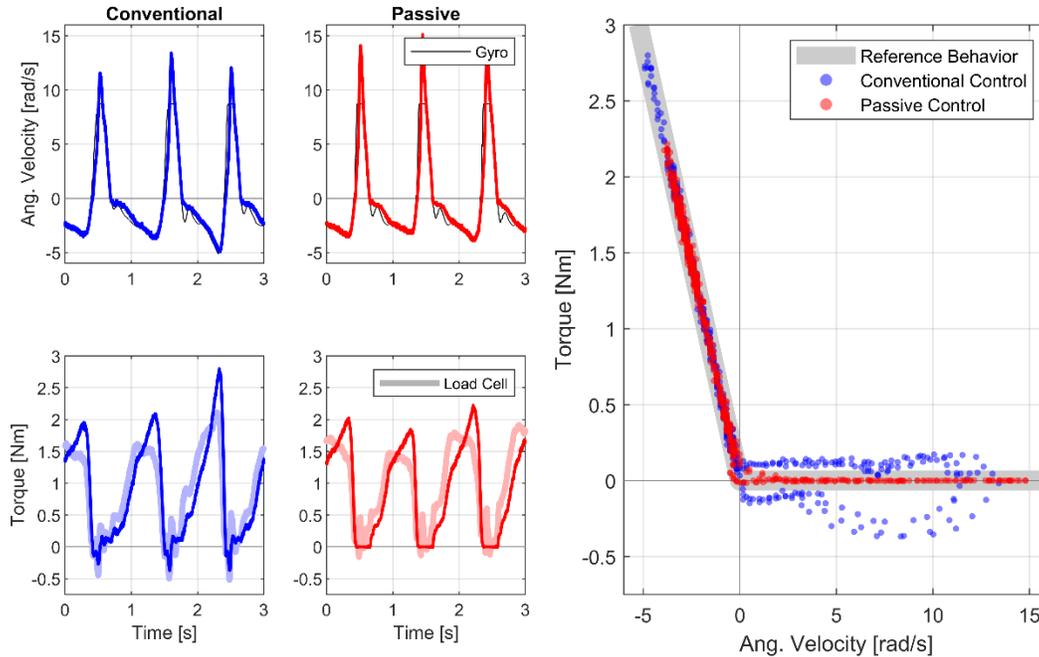


Fig. 4.9. Results of first asymmetric damping experiment. Data is shown for internal sensors (motor encoder and current) which are used in by the control system as well as external sensors (gyroscope and load cell) which are only used for validation. Only internal sensor data are shown in the power plane on the left.

Note that the damping reference for negative velocities is close to the experimentally determined limit of stability. To evaluate even greater asymmetries, gains of the conventional motor controller were lowered (in order to avoid instability in the conventional control case). In Fig. 4.10, a new set of results is presented with $B_{pos} = 0.005 \text{ Nm} \cdot \text{s}/\text{rad}$ and $B_{neg} = 1.146 \text{ Nm} \cdot \text{s}/\text{rad}$ (double of the previous case). While the lowered controller gains allow the conventionally controlled system to achieve higher magnitudes of emulated damping, performance is worse than in the previous case (with higher gain and lowered limit of stability). Again, that effect is clearly observed with measurement of motor current and can clearly be seen in load cell data. As such, the system with conventional control must balance a tradeoff between stability margins and accuracy of behavior. That tradeoff does not exist with the passive switching system, which achieves higher behavioral fidelity without risk of instability.

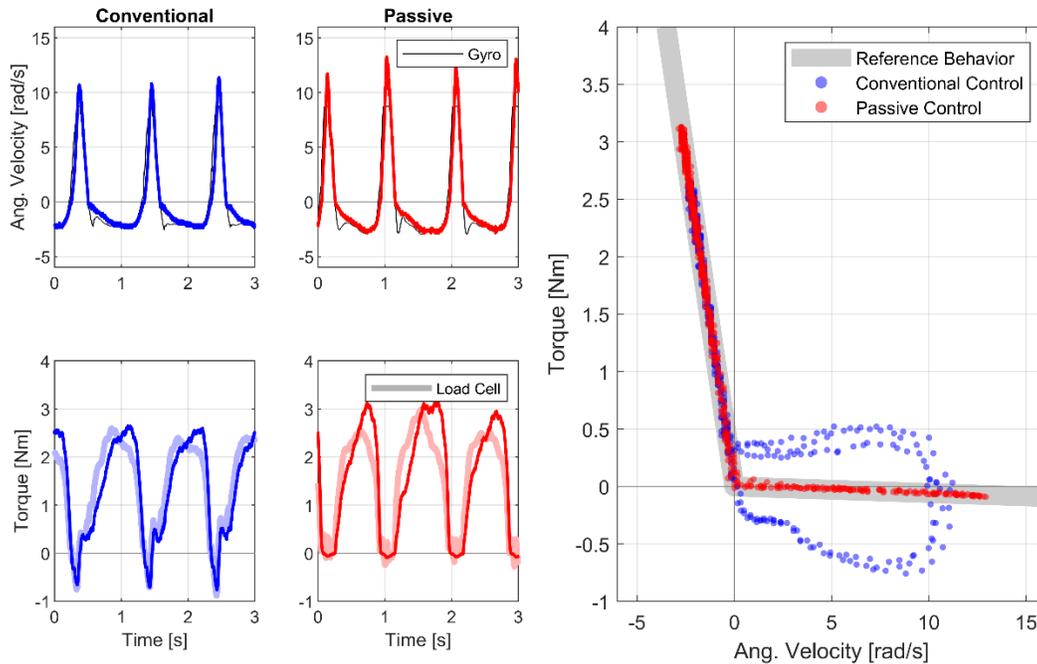


Fig. 4.10. Results of the second asymmetric damping experiment.

The passive control method consists of an additive combination of feedforward and feedback controllers. Note, however, that either control element can be used independently, which is highlighted in this final set of results. Current tracking performance for same experiment as above (Fig. 4.10) are presented in Fig. 4.11. Results are shown for the conventional method, the passive method, and the passive method with the feedback controller disabled (that is, with only the feedforward controller). As discussed, the conventional method suffers from poor tracking when the reference behavior changes abruptly, demonstrating the energy leak phenomenon. The importance of the feedback controller to the passive method is clear from the results: while the feedforward-only case demonstrates fair results, it cannot track the reference with zero error. The duty cycle plots show the immediate change in behavior that the (passive) feedforward controller achieves when the reference abruptly changes, which does not happen in the conventional case.

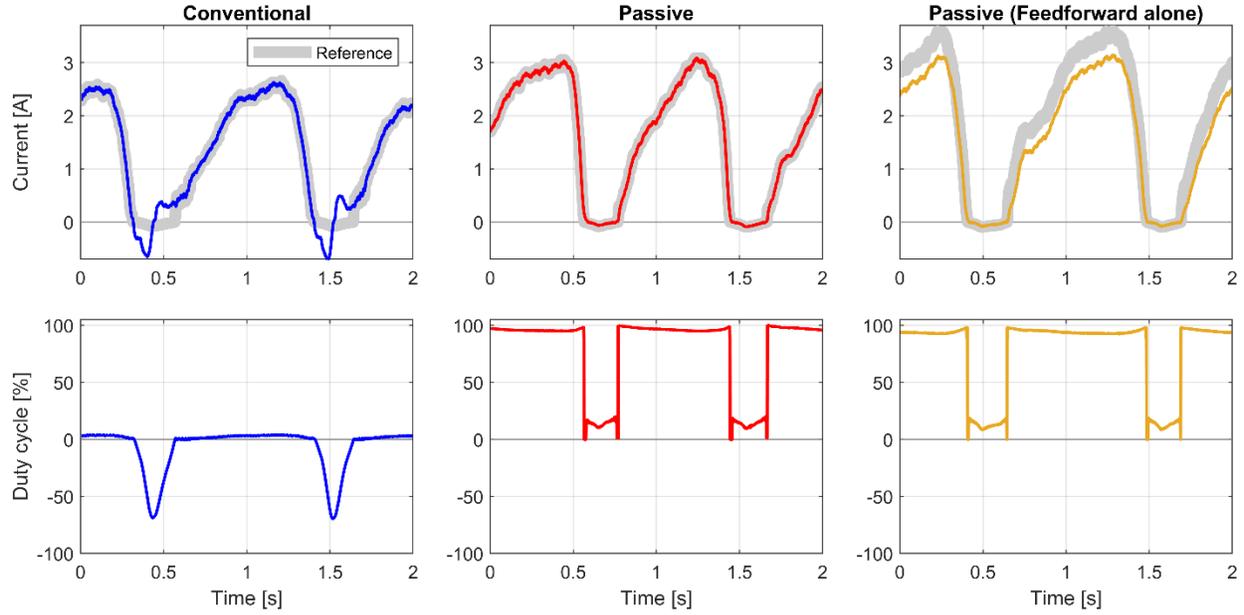


Fig. 4.11. Current tracking and controller output for the same asymmetric damping emulation task presented in Fig. 4.10. The three columns show results with the conventional motor control approach on the left, passive in the middle, and passive limited to using only the feedforward component on the right.

4.6 Discussion

Relative to a conventional damping control approach, the passive control method demonstrated lower levels of vibration, better stability characteristics, and improved behavioral accuracy. Specifically, energy leaks did not occur with the passive method and the system does not exhibit instability, both of which are consistent with the previously presented theoretical and simulation analyses, and both of which are problematic for the conventional implementation. The absence of energy leaks is demonstrated with measurements of motor current and confirmed by measurement of user interaction force (load cell installed at the user interface with the experimental apparatus). The level of vibration for both passive and conventional implementation was measured by an accelerometer; in the case of passive control, results show that the level of vibration falls as emulated damping grows; the opposite occurs with conventional control (vibration increases until the system reaches instability). Moreover, a tradeoff was observed with the conventional control method between behavioral accuracy and the stability margin for damping emulation: when the system allows for higher levels of damping to be achieved, behavioral accuracy worsens. That tradeoff does not exist in the passive control case, since stability is guaranteed.

Despite these observed advantages, the passive control method has important behavior limitations relative to conventional control methods, which should be emphasized here to balance the comparison. As has been shown, and as the name suggests, the passive approach constrains the motor to the passive quadrants by construction (a desirable feature when providing braking behavior). However, the behavior is further constrained due to the short-circuit limitation within the passive quadrant. From (4.5), assuming steady-state operation, the motor current becomes as indicated in (4.16).

$$i_M = -\frac{k}{R}u\dot{q} \quad (4.16)$$

If control is saturated at 100% (motor leads kept shorted all the time), then the resulting motor behavior is a fixed level of damping, referred to here as “short-circuit damping”. The range of achievable behaviors with the passive approach is shown schematically as the green region in Fig. 4.12. That limitation has important implications to braking functions. For example, the simple task of providing a constant opposing torque is unachievable in low velocities depending on the required magnitude of impedance. For the system in the experimental apparatus, $B_{MAX} = (kn)^2/R \approx 1.93 \text{ Nm} \cdot \text{s/rad}$.

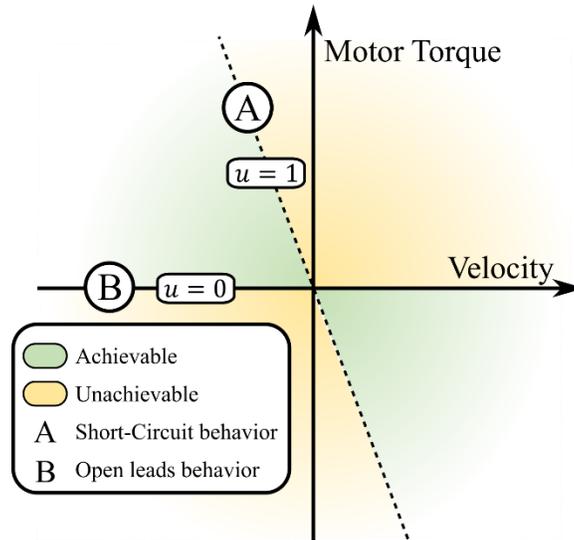


Fig. 4.12. Achievable and unachievable behaviors with the passive control approach as a function of motor current (torque) and velocity. The corresponding duty cycle command for passive switching is shown for reference.

We note, however, that achieving behaviors that surpass short-circuit damping with any motor control technique requires the use of battery power (i.e., requires expenditure of electrical power in order to remove mechanical power from the system). Despite being a nonintuitive result, this characteristic has been well-established by other authors and is explored in the prior work as well [3]–[5]. In other words, energy regeneration is only achievable up to the short-circuit damping behavior. As such, passive motor control constrains the motor to a region of generative behavior, which is a subset of the range of conventionally controlled regions of operation. If motor and drivetrain parameters can be selected such that the region of generative behaviors fulfills the needs of a given application, passive motor control will provide enhanced damping emulation, relative to conventional control. It should also be noted that, since passive control does not require modifications to the standard BLDC control bridge, a control system may be able to change control strategies in real-time if the required behavior surpasses the constitutive limitation, preserving the versatility of the BLDC motor.

The passive motor control method also has some additional implementation complexity relative to the conventional method, since it relies on the combination of feedback and feedforward controllers, while the conventional method requires only feedback control. Note, however, that the passive approach does not require electronic commutation, which simplifies the implementation of the BLDC motor control hardware. Additionally, the feedforward element has been simplified in this work (relative to [5]), now consisting of a plane fitting procedure and a lookup table, both of which are suitable for implementation in an embedded context.

4.7 Conclusion

The authors present herein a control method of the strictly energetically passive implementation of damping emulation with a BLDC motor. The proposed motor control method is compared directly with a conventional approach, both analytically and experimentally. The comparisons highlight improvements in stability, smoothness, and behavioral accuracy with the passive motor control scheme. The passive approach further does not require commutation, which simplifies implementation in a BLDC motor. Relative to the conventional approach, however, our approach is limited to a narrower range of the torque-speed plane, and is best implemented with a feedforward component, which adds (or at least trades) complexity. For applications requiring damping emulation, the authors recommend employing the passive control method when possible (i.e., selecting motor and transmission parameters such that passive motor control will satisfy the range of required emulated damping behaviors). Alternatively, if energetically dissipative behaviors are required outside of the passive motor control regime, or if a given application requires periods of damping emulation combined with separate periods of non-energetically-passive control, it may be advantageous to switch between control methods, particularly since both control approaches employ the same BLDC motor driver MOSFET bridges. That is, the passive control method can be employed for improved performance and stability when behavior is within the passive control regime, and the conventional control approach used otherwise.

5 A method for optimized energy regeneration in a knee prosthesis

This chapter has been adapted from a draft journal publication which will be submitted for review in the coming months.

5.1 Introduction

Harnessing electric power from human locomotion has been a research goal in lower-limb prostheses for several decades, particularly in knee joints [6]. During walking, the knee demonstrates a net negative energetic profile [83], dissipating a significant amount of mechanical energy per stride (approximately 10 J per stride for an average-sized individual walking at a comfortable speed). A microprocessor-controlled knee prosthesis (MPK) capable of harnessing part of this dissipated energy to power its internal systems, sensors, and actuators can potentially employ lighter, smaller batteries or trade that performance advantage for longer battery life. In certain instances, it may even be possible to achieve sufficient regeneration such that a prosthesis that can operate indefinitely without requiring a recharge. An important aspect of such regeneration, however, is that it should not affect the functional behavior of the prosthesis. As such, it should not measurably alter the interaction between the user and prosthesis, relative to a desired non-regenerative behavior; that is, the presence of regeneration should not increase user energy expenditure, not alter the kinematic or kinetic patterns of the user or prosthesis during walking.

Net dissipation of mechanical energy in walking gait occurs primarily (and arguably entirely) between late stance and terminal swing, when the knee joint is either driven by the hip to initiate swing or is modulating the swing trajectory to provide ground clearance and a comfortable level of deceleration to prepare for stance. The behavior of the knee during earlier phases of stance requires higher levels of torque to support the weight of the user, occurring with approximately net zero mechanical energy. The challenge to energy harnessing from the knee joint arises from the order-of-magnitude difference in required impedance (ratio of torque to velocity) between stance and swing phases; although the regeneration potential occurs in swing with low impedance (low torques and high velocities), the system must still be able to implement the high-impedance stance behaviors (relatively high torques and low velocities).

As has been well-established, direct current (DC) motors are only able to regenerate energy within a limited band of impedances in the passive quadrants [3], [12], [24]. Motors also demonstrate diminishing efficiency in the active quadrants as the desired torque-to-velocity ratio rises, reaching peak active efficiency around the same magnitude of impedances that allow for regeneration in the passive quadrants. Therefore, in order for a DC motor to operate efficiently and regenerate energy, the drive (including motor and transmission) must be properly sized for the application.

A motorized prosthesis that can regenerate energy in swing phase may require too much power to implement the elevated levels of impedance desired for stance due to diminishing efficiency (examples in [19], [20], [66], [84]). Other authors have shown promising emerging harnessing results with such devices, such as in [66], but the lack of stance support ultimately limits their applicability and safety. Conversely, if a drive is optimized for high-impedance behaviors (e.g., to assist with stair ambulation or sit-to-stand tasks) it will likely not be able to regenerate energy in swing (without supplemental exertion by the user), as the band of regenerating impedances is shifted upwards due to the higher intrinsic friction associated with such drives (example in [13]).

The prototype prosthesis used in this work employs an electronically controllable transmission (ECT) mechanism with solenoid-operated clutches that can quickly vary the transmission ratio between motor and knee joint by an order of magnitude [28]. The actuator of the ECT prosthesis is put in the low-gear state for late-stance and swing, which enables regeneration in those phases of walking, and in the high-gear for the earlier phases of stance in order to support the weight of the user. Other research prototypes that employ actuators with variable transmission ratio have recently demonstrated encouraging results in energy harnessing. For instance, the device in [26] employs a passive spring-based mechanism that increases the transmission ratio of the system in tandem with the output torque. The authors report significant energy regeneration in walking when using an alternative, “passive” controller. The earlier device in [25] developed by the same research group uses an actively variable transmission and also reports significant energy regeneration in walking.

Controllers for the ECT prosthesis have been developed for the tasks of level walking, slope walking, stair ambulation, and sit-to-stand and stand-to-sit transitions, as presented in [27]. The walking controller has two modes of operation, depending on the estimated walking speed of the user. Active behaviors are employed with low walking speeds (below around 0.7 m/s) to assist the user in reaching adequate levels of knee flexion. For medium and fast walking speeds, the controller behaves entirely passive, with the motor behavior constrained to the regeneration band by virtue of the passive motor switching scheme employed [12]. Although energy regeneration is inevitable given the motor control approach, the original controller was designed solely with the objective of achieving adequate swing-phase motion.

In this work, a new walking controller for medium to fast walking speeds is developed and tested on an amputee participant. The new control laws are defined based on an optimal control problem constructed with a cost function that penalizes deviations from the kinematic goals of walking while simultaneously favoring regenerating as much electric energy as possible. The optimization routine, which runs offline, takes as input various parameters of the ECT device, walking kinematics from datasets of healthy and amputee subjects, and explicit kinematic goals such as peak flexion and terminal swing velocity. The Pontryagin maximum principle formulation is used to construct the optimal control problem with calculus of variations, which is then solved numerically with the shooting method. The authors believe this to be a novel approach to design control laws for a knee prosthesis [85], so the energy optimization approach is described in detail in Appendix I for a generic mechatronic system which has terminal goals for its state but that does not specify trajectories. The application of the method to the ECT prosthesis within walking tasks is detailed in Appendix II.

The optimization results inform the “shape” of real-time control laws for swing flexion and extension, which are parametrized and implemented in the prosthesis. Just like with heuristics-based control laws, parameters can be adjusted when fitting the prosthesis to the user [85]. Note that the optimization approach taken here requires knowledge of quantities that are not directly measurable from embedded sensors in the prosthesis and is also non-causal, requiring knowledge of past and future state values. Those characteristics make online human-in-the-loop optimization impractical with the current optimization scheme, although optimization is still possible if a different optimization strategy were adopted [86]. Moreover, having fixed, non-trajectory-based behaviors affords greater predictability to the device, which has advantages to user confidence, learning, and motor adaptation [87].

In addition to the optimization of swing behaviors, an exploratory control law for late stance is proposed. Note that late stance is more difficult to model since the foot is still in contact with the ground and the

lower legs effectively form a closed chain system. The new controller fits well within the framework of controllers developed for the device (for instance, the same state transition conditions are used) so new and original behaviors can be selected via a computer interface. In that way, the original controller for medium to fast walking can be “switched on” to serve as a baseline for performance comparisons.

Experiments ultimately show that the optimized control laws for swing significantly increase energy regeneration in swing. A difference of about 1 J per stride is recorded both in the motor and by measurements of battery current, raising the overall efficiency of the system by between 5 to 10%, depending on the walking speed. The exploratory control law for late stance was also shown to improve regeneration in that phase of gait, however the gains in regenerated energy were limited.

5.2 Prosthesis Hardware

The design of the ECT prosthesis used in this work will be presented here in an abbreviated format, see [27] for details of the design and discussions of the approach. The device was conceptualized as a two-gear mechatronic prosthesis capable of switching between configurations of transmission ratio within a few milliseconds, enabling the use of a single motor to provide adequate behavior in all phases of gait. Specifically, a high-gear setting is used to produce the elevated knee joint torques required in stance, and a low-gear setting is used to minimize drive impedance and therefore maximize back-drivability in swing. Photographs of the device are presented in Fig. 5.1 (a).

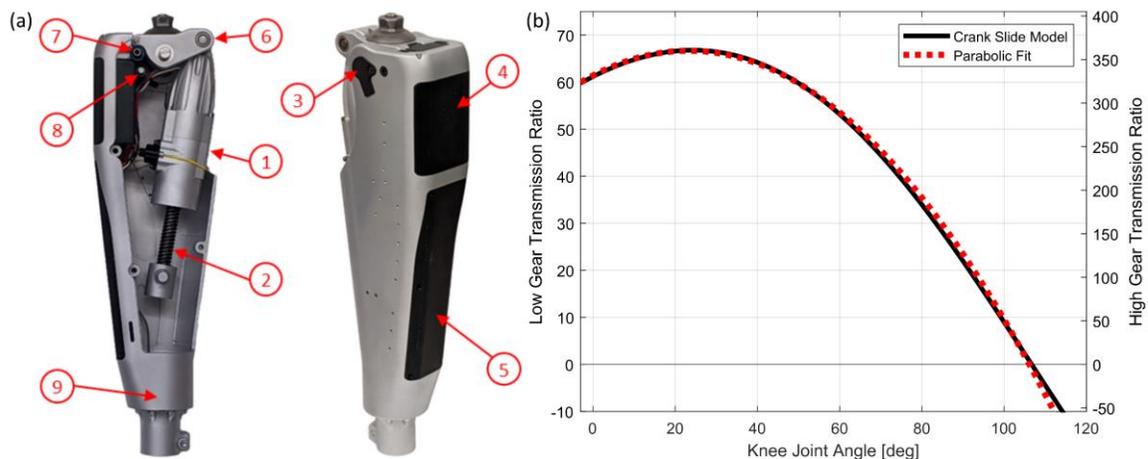


Fig. 5.1. In (a) are two photographs of the ECT knee prosthesis prototype obtained from [28] with one half of the housing removed on the left and fully assembled on the right; the numbered arrows point to the actuator (1), lead screw (2), absolute encoder (3), battery pack (4), housing for the custom embedded system (5), crank (6), extension and flexion hard stops (7 and 8, respectively), and load cell (9); in (b) is a plot of the motor-to-joint transmission ratio with respect to knee angle for high and low-gear configurations based on the crank slide model as well as a parabolic fit, which is computationally simpler.

The ECT actuator consists of a brushless direct-current motor (EC 4-pole 22 90W; part number 323217; Maxon; Sachseln, Switzerland) with a modified two-stage planetary gearhead and lead-screw output. Modifications to the gearhead consisted of adding two solenoid-operated clutches to alternate between grounding the ring gear of the transmission to the actuator body (which is the typical planetary gear application and results in the high-gear setting) or grounding the first stage planetary gear (which results in the low-gear setting). The high- and low-gear configurations cause a five-fold difference in transmission ratio magnitude: 29.2:1 and 5.4:1, respectively, between motor and lead screw. High-gear is the default configuration when both solenoids are de-energized, which is a safe configuration in case of a power-off

fault. The ring clutch was built with a unidirectional element, so the ring gear can freely rotate in one direction regardless of the state of the associated solenoid. The unidirectionality of the ring gear provides two characteristics which are desirable for a knee prosthesis: 1) when the device is in high-gear, the transmission is underconstrained in the extension direction, so the user can freely extend the leg (this behavior is implemented with a check valve in hydraulic devices [9], [10], [78], [79]); and 2) when the carrier clutch is engaged and the ring clutch is disengaged, the transmission is overconstrained in the flexion direction and provides low-gear behavior in the extension direction, which is useful to prevent knee buckling in case of disturbances such as stumble and foot scuff during swing extension. The combination of the two solenoid-operated clutches and unidirectionality of the ring clutch create three behaviors of interest which align well and produce desirable conditions for walking stance, flexion, and extension. The specific mapping between hardware configuration and phases of gait presented in Table 5.1.

Table 5.1. Relationship between ECT hardware configuration, resulting actuator behavior in either direction of knee joint rotation, and associated phase of walking gait for which each configuration is useful.

Clutch Solenoid State		Knee Joint Direction of Motion	Actuator Behavior	Gait Phase
Carrier	Ring			
De-energized	De-energized	Flexion	High-Gear	Stance
		Extension	Free	
Energized	Energized	Flexion	Low-Gear	Flexion
		Extension		
Energized	De-energized	Flexion	Locked	Swing Extension
		Extension	Low-Gear	
De-energized	Energized	Any	Free	Not used

The actuator is mounted on a crank-slide mechanism in the leg housing which transforms rotation of the actuator’s lead screw into rotation of the prosthesis at the knee joint. The transmission ratio between actuator output and knee joint is variable with respect to knee angle; the system geometry was designed such that the crank slide transmission ratio is higher for lower levels of knee flexion (associated with higher torque phases of gait such as stance and terminal swing) and lower for higher levels of flexion (associated with swing). The effective transmission ratio between motor and knee joint is shown for the low- and high-gear settings with respect to knee angle in Fig. 5.1 (b). Flexion and extension hard stops are integrated in the housing.

Besides coupling the actuator to the knee joint, the ECT housing also contains a custom strain-gage load-cell to measure axial lower leg loading, an absolute encoder at the knee joint, a custom embedded system with motor and solenoid control hardware, and a nominally 16 V battery pack which consists of four 18650 Li-Ion cells connected in series. Additional sensing in the system includes a motor shaft encoder; a 6-axis IMU; electrical current sensing for the battery, motor, and solenoids; and battery voltage readings. The embedded system is responsible for reading all system sensors (signal conditioning and acquisition), controlling motor current, operating solenoids, and exchanging real-time control data with a computer which implements high-level control tasks (running Mathworks Simulink Desktop Real-Time).

Motor current can be controlled actively using the standard technique of complementary switching with block commutation or passively by alternating between short-circuiting and opening the motor leads in a PWM manner, as presented in [12]. As implemented, passive motor control is preferred because it guarantees energy regeneration into the battery by precluding the motoring and reverse current braking

regimes. Note that providing torque at zero velocity is not possible with the passive scheme because the behavior is limited to reflected electrical impedance of the motor, corresponding to a maximum amount of damping. As such, when passive control is employed, the prosthesis behaves as an energy-regenerating mechanical damper which can be modulated in real-time between 0% (no motor torque) and 100% (motor short-circuit).

Controllers which leverage passive and active behaviors have been developed for level walking, slopes, and stair ambulation tasks [27] and tested with an amputee participant. Results presented so far demonstrate that the ECT can match or exceed the performance of MPK devices by adding power if and when needed (e.g., swing flexion assistance in slow walking gait, powered stance extension for upstairs walking and sit-to-stand transitions). Importantly to this study, the intrinsic impedance of the ECT drive is sufficiently low that it can operate entirely passively for level walking at comfortable walking speeds.

5.2.1 Model for the ECT drive

A model for the ECT drive that relates motor current with knee joint torque in the low gear is presented here, including drive friction and the nonlinear effect of the crank slide mechanism. Models for knee joint power and motor power follow from the derivation of kinetics. The resulting relationships shall subsequently be applied within the optimization framework developed to find energy-optimized control laws for the swing phase.

5.2.1.1 Kinetics

Knee joint torque τ_K expressed in (5.1) includes commanded motor torque (through motor current i) and drive friction. Motor current is multiplied by the motor torque constant k and by the angle-dependent transmission ratio $n(\theta_K)$ between motor and knee joint in the low-gear setting. As shown in Fig. 5.1(b), the nonlinear effect of the crank-slide mechanism can be adequately represented with a parabolic fit as in (5.2), which significantly simplifies analysis. Friction is modeled based on a characterization of the ECT in low gear as presented in [28] including Coulomb and viscous friction terms (τ_C and b , respectively) as well as a motor-current dependent term β , which effectively adds to the baseline Coulomb friction. The latter is an effect of the sliding interaction in the lead screw (i.e., interaction force amplifies friction). Drive inertia is negligible relative to the lower leg inertia and therefore is not considered. Values for the parameters used in (5.1) and (5.2) are in Table 5.2.

$$\tau_K = n(\theta_K)ki - b\dot{\theta}_K - (\tau_C + \beta|i|) \operatorname{sgn} \dot{\theta}_K \quad (5.1)$$

$$n(\theta_K) = \alpha_2\theta_K^2 + \alpha_1\theta_K + \alpha_0 \quad (5.2)$$

5.2.1.2 Power

Knee joint instantaneous power P_K is given by the product of knee torque in (5.1) and knee angular velocity $\dot{\theta}_K$: $P_K = \tau_K\dot{\theta}_K$. Motor electrical power P_E (the integral of which this work seeks to optimize) is given by the product of motor current i and back electromotive voltage (back-EMF) plus resistance losses (Ri^2), where back-EMF is given by motor velocity $\dot{\theta}_M$ multiplied by the motor torque constant k . The transmission ratio model in (5.2) is used find electrical power in terms of knee angular kinematics, as shown in (5.3). The value of winding resistance R is in Table 5.2.

$$P_E = ki\dot{\theta}_M + Ri^2 = n(\theta_K)ki\dot{\theta}_K + Ri^2 \quad (5.3)$$

Table 5.2. Parameter values for ECT friction and transmission ratio models, and nominal parameter values for the motor used in the prosthesis

Parameter	Symbol	Value
Coulomb friction torque	τ_c	0.033 Nm
Viscous friction coefficient	b	0.0022 Nms/rad
Friction sensitivity to motor current	β	0.055 Nm/A
Parabolic model for transmission ratio	α_2	-31.91 rad^{-2}
	α_1	25.96 rad^{-1}
	α_0	61.39
Torque constant	k	0.0105 Nm/A
Winding resistance	R	0.323 Ω

5.3 Optimizing Energy Regeneration

In this section we discuss the energetics of level-walking amputee gait within the context of the three ECT clutch states as presented in Table 5.1 (which provide a convenient underlying structure to separate controlled behaviors). Then, we introduce the optimization framework used to inform new controlled behaviors that improve energy regeneration in walking gait while achieving kinematic goals. The existing walking controller for the ECT is presented alongside the proposed behaviors since it serves as a baseline for energy metrics.

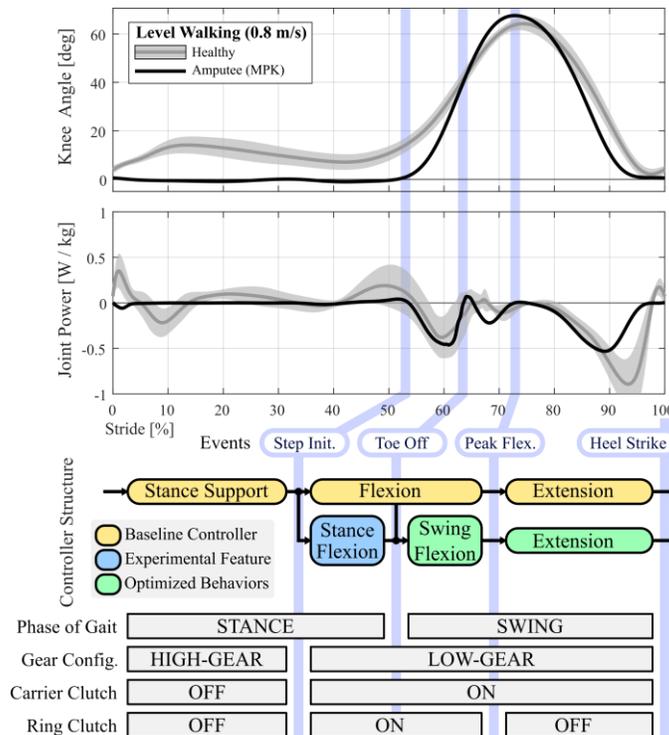


Fig. 5.2. Summary of walking gait energetics and ECT control structure.

5.3.1 Energy Regeneration Potential in Walking Gait

A typical level walking stride for a transfemoral amputee can be divided into four major sections: stance loading and support, stance flexion (or pre-swing), swing flexion, and swing extension. Representative knee joint kinematics and joint power are shown in Fig. 5.2 for healthy and amputee (wearing an MPK)

walking on a treadmill at 0.8 m/s. The four sections of gait are separated by the vertical lines representing gait events that cross the plots and are extended down to highlight the transitions in the ECT controller structure and hardware configuration.

Note that the stance phase differs significantly between able-bodied and amputee cases because amputees seek to prevent knee buckling during weight acceptance and single support phases of stance, generally maintaining the knee joint hyperextended from heel strike to pre-swing [88]. Therefore, stance support does not have any potential for energy regeneration, since the knee velocity is approximately zero. During that phase of walking, the ECT is configured in high-gear state in order to provide significant resistance to flexion.

The controller switches the prosthesis to the low-gear state for “step initiation” by energizing both clutches when the user starts to unload the leg and forward motion of the shank is detected with the IMU sensor. The prosthesis remains in the low-gear state throughout flexion and extension phases, until the following heel strike event is detected (at which point the prosthesis switches back to high-gear and the cycle repeats).

Late-stance flexion happens during double stance (following the contralateral heel strike, while the prosthesis side is unloading) and serves to configure the leg for swing. Flexion of the prosthetic joint is driven by the user’s hip during stance flexion because the foot still makes contact with the floor; the low-gear state of the ECT facilitates knee flexion. As shown in Fig. 5.2, this phase of gait has a substantial amount of dissipative knee joint power in both healthy and amputee cases, so there is potential for energy harnessing. Since motion is directly driven by the user, an alternate stance flexion control law shall be defined to evaluate the extraction additional energy at the expense of increased user effort. It is not clear if removing more energy than prescribed will adversely affect the user’s gait, so this controller state is strictly experimental (will be manually enabled or disabled for each experiment).

The low impedance of the ECT prosthesis in low-gear allows for swing flexion and extension to follow a biomimetic ballistic trajectory which is mostly governed by gravity and inertial forces [15], [89]. As such, the motor behavior is strictly passive in moderate to fast walking speeds, which permits energy regeneration. If walking speed is low, then energy must instead be added to the mechanical system to achieve a proper amount of flexion for toe clearance. The baseline controller contains a powered swing flexion state in which an assistive torque pulse is provided [27] (not presented here). It was found that addition of power is helpful for walking speeds below 0.7 m/s with the ECT [27], so energy regeneration in swing flexion is only possible with walking speeds above that threshold. In healthy gait, that crossover occurs between 0.35 to 0.5 m/s (based on an analysis of the dataset in [90]).

The ECT de-energizes the ring clutch at peak flexion to achieve unidirectional low-gear behavior in swing extension, with benefits to gait robustness and reduction of battery power demand. The prosthesis provides resistance to slow down the lower leg in preparation for stance, providing a comfortable level of terminal impact at the extension hard stop (users find that some impact is useful as a cue to indicate the leg is completely extended). The behavior in swing extension remains dissipative regardless of walking speed, with considerable potential for energy regeneration (more than the previous phases).

Energy is therefore available for harnessing by motor regeneration during stance flexion, swing flexion, and swing extension at comfortable walking speeds, with extension providing the largest amount of negative joint energy during terminal swing.

5.3.2 Framework for Optimizing Energy Regeneration in Swing

Motion of the lower leg during swing phase follows the behavior of a pendulum with movable pivot point, where the pivot point is the knee joint that moves with the user's thigh. As such, the behavior can be described with the dynamical model in (5.4), where shank angular acceleration $\ddot{\theta}_S$ is given by knee joint torque τ_K from (5.1) as well as terms corresponding to the effect of gravity and acceleration of the knee joint in the sagittal plane (i.e., movement of the knee joint itself, not to be confused with knee angular acceleration). The diagram in Fig. 5.3 shows all variables of interest for the lower leg model. Parameters m , I , and l represent the user-dependent lower leg mass, inertia by the knee joint, and distance between knee joint and lower leg center-of-mass (CoM), respectively. Knee joint acceleration produces an inertial force F_I at the CoM, which is considered an exogenous input to the system (i.e., the distal end of the thigh is considered entirely driven by the user). F_I is obtained by rotating the sagittal acceleration vector (\ddot{x}_K, \ddot{y}_K) by θ_S and taking the component perpendicular to the lower leg; the axial component is ignored since it does not cause a torque about the knee.

$$I\ddot{\theta}_S = -\tau_K - l \overbrace{[m\ddot{x}_K \cos \theta_S + m\ddot{y}_K \sin \theta_S]}^{\text{Inertial force } F_I} - \overbrace{lmg \sin \theta_S}^{\text{Gravity}} \quad (5.4)$$

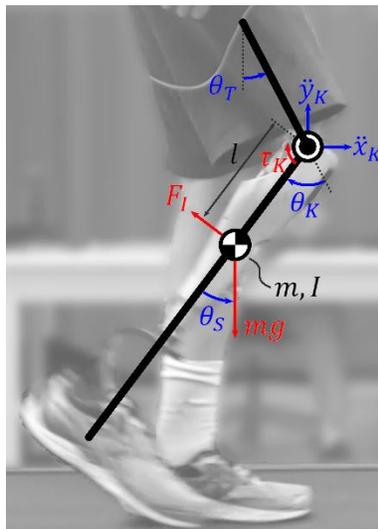


Fig. 5.3. Schematic diagram of the leg with variables and parameters of interest for the dynamical model in (5.4). The diagram is overlaid on a photo of the amputee participant wearing the ECT prosthesis in mid-swing during treadmill walking.

Although knee angle and angular velocity appear in the system equation (as well as in τ_K), those variables are not part of the state of the system. Instead, knee configuration is determined by the difference between thigh and shank angles: $\theta_K = \theta_T - \theta_S$, where shank angle is the system state and thigh angle is provided as an exogenous input (along the acceleration of the knee joint in the sagittal plane). In order to simulate the system, one must therefore provide user-specific parameters m , I , and l ; user- and task-specific data for the exogenous signals θ_T , $\dot{\theta}_T$, \ddot{x}_K , and \ddot{y}_K ; initial conditions for the system state: $\theta_S(0)$ and $\dot{\theta}_S(0)$; and simulation time T equal to the duration of the specific activity. All of which can be obtained from datasets of gait kinematics. Lastly, one must also provide a control law for motor current, which is the topic of the subsequent sections.

5.3.3 Optimal Swing Behaviors

The energy optimization approach described in Appendix I is applied to the lower leg mechanical system with a cost function that combines motor energy in the electrical domain (the time integral of (5.3)) and kinematic goals for terminal leg configuration. The outcome of the optimization is the unbounded motor current signal in time which minimizes the cost function. Refer to Appendix II for the full setup of the optimization problem (with definitions for cost function, Hamiltonian, and co-state variables) and the numerical algorithm used to solve the resulting boundary value problem. In order to impose kinematic goals for each phase of gait, the optimization is executed separately for swing flexion and swing extension (which conveniently eases numerical issues with the Coulomb friction term in (5.1)). The kinematic goals for swing flexion are specified as achieving biomimetic peak knee angle (from biomechanics data) and zero angular velocity; the kinematic goals for swing extension are zero terminal knee angle and low terminal knee angular velocity which has been found to be conducive to a comfortable level of terminal impact. The method also attributes weights to each kinematic goal in the cost function, as discussed subsequently.

Two datasets with walking gait are used to provide parameters and exogenous signals to the dynamical system. The first dataset, which was obtained from [91], contains level and ramp treadmill walking trials of ten able-bodied participants at different speeds, ranging from 0.8 m/s to 1.6 m/s. User-specific parameters of lower leg inertia, mass, and CoM location were preserved. Note that lower leg inertia varies during swing because the foot rotates about the ankle joint, so an average value was used as a proxy since that variation does not occur in the prosthesis case. The average-gait movement of each participant's thigh segment provided the necessary set of signals for thigh angle, thigh angular velocity, and acceleration of the knee joint in the sagittal plane. The second dataset contains walking data from one amputee participant walking with the ECT baseline controller in a treadmill at speeds ranging from 0.4 to 1.2 m/s. Lower leg inertia, mass, and CoM location were directly measured by disconnecting the ECT from the user's socket. Similar to the first dataset, averaged thigh motion provides the necessary exogenous signals for the model.

Results for both datasets are presented in Fig. 5.4 for swing flexion and Fig. 5.5 for swing extension. Treadmill walking speeds of 0.8, 1.0, and 1.2 m/s are shown to demonstrate the variation in optimized behavior that is associated with the different gait patterns of each walking speed. Plots show the simulated knee angle and knee angular velocity trajectories, the variation between the simulated and actual (nominal) knee angle, and the optimized motor current and motor impedance z . Motor impedance can be interpreted as a damping relationship, defined as the ratio of motor torque to motor velocity: $z = -ki/\dot{\theta}_M$ (negative so that a positive value indicates resistance to motion). The range of motor impedance that is achievable with passive motor control (and therefore the range in which energy regeneration is possible) is highlighted in green.

The cost function weights associated with each kinematic goals were manually adjusted for dataset 2 and kept unchanged for dataset 1. The gains were chosen to balance the accomplishment of kinematic goals and energy regeneration; values are presented in Appendix II.

In swing flexion, the goal of biomimetic peak knee angle is specified as the nominal value for each simulated scenario and is achieved within 1 degree in most cases. The goal of zero terminal velocity is also adequately achieved. The optimized behaviors consist of mostly flat motor impedance commands which present a high degree of user specificity. The impedance curves lay within the regenerating region (except

for discontinuities towards the end when velocities cross zero) and trend upwards as walking speed is increased.

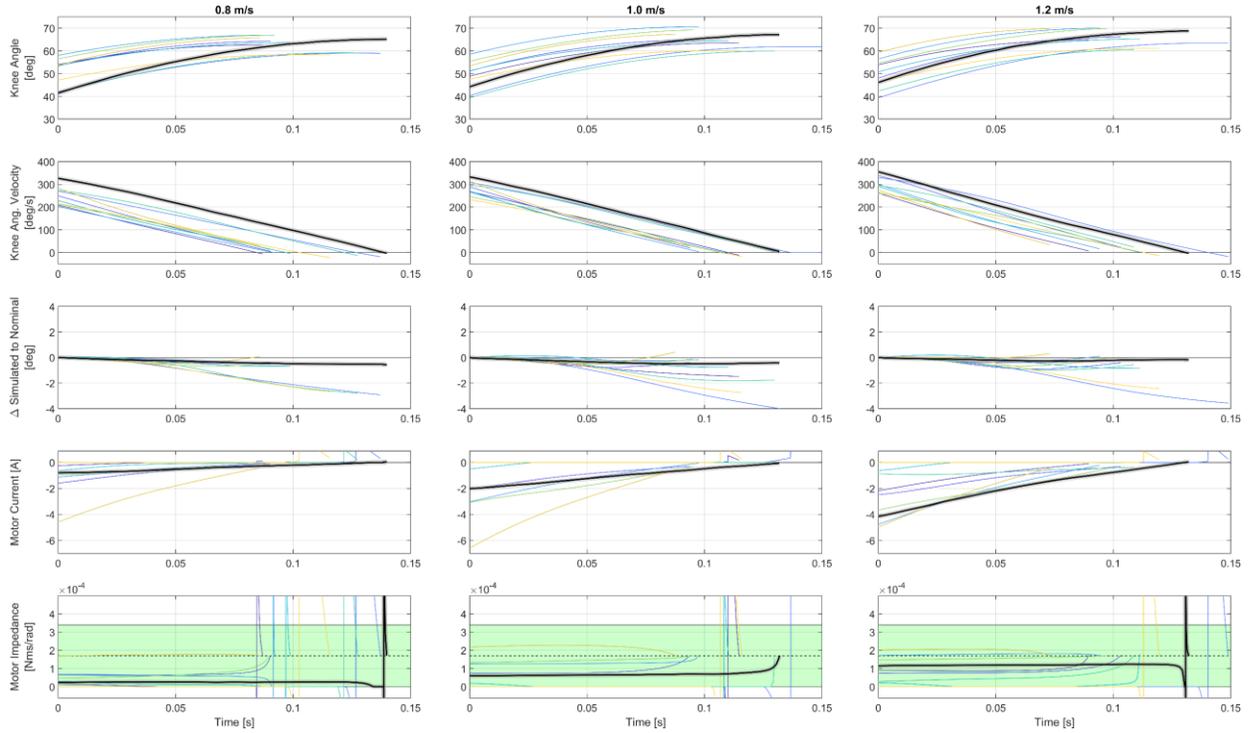


Fig. 5.4. Optimization results for swing flexion.

In extension, the kinematic goal of full leg extension at the end of the swing phase is also adequately achieved (within a few degrees). The specified terminal impact velocity is set at 80 deg/s and is achieved with a considerable amount of error, in the order of tens of degrees per second in some cases in the first dataset. The second dataset demonstrates lower error ($\dot{\theta}_K(T) = \{-83.5, -114.1, -117.7\} \text{ deg/s}$ for the respective walking speed settings) because the cost function gain associated with the terminal velocity goal was defined using that dataset. Two factors contribute to the higher degree of error in the velocity goal on the first dataset: 1) the prosthetic lower leg in the second dataset has lower inertia (normalized to body weight) than the lower leg of the able-bodied participants in the first dataset and the contribution of the exogenous torque inputs (thigh motion) has likewise lower magnitude, so the motor can more easily affect knee velocity; and 2) since the cost function incorporates both kinematic goals and energy in the electrical domain, the optimizer trades the velocity goal in order to avoid crossing into the motor plugging regime (the area above the green shaded region in the motor impedance plots), which has significant energy cost.

The optimal behaviors in extension consist of initially maintaining motor current at zero and then ramping up motor impedance until the end of swing. Motor current and motor impedance have higher magnitudes in the first dataset (for the same reasons enumerated above). Note that motor impedance converges to the maximum power transfer condition for the prosthetic case in the second dataset, which is one of the main observations of the general energy optimization problem explored in Appendix I.

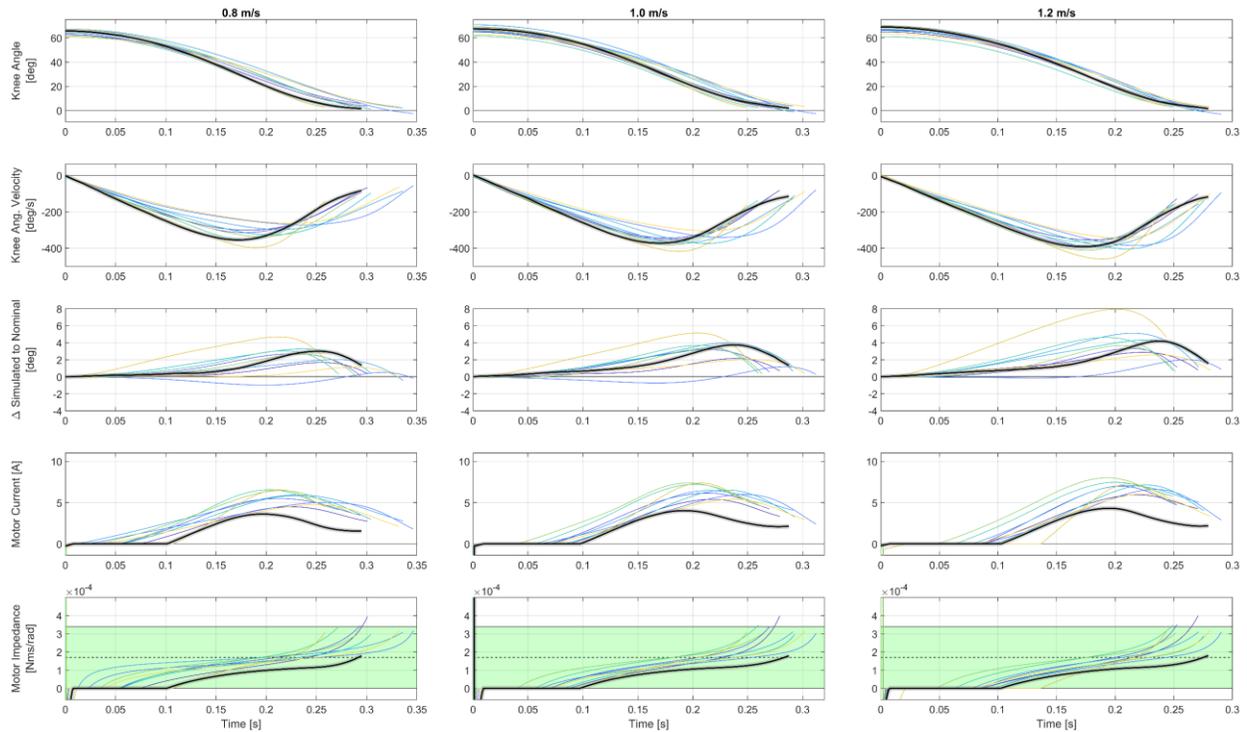


Fig. 5.5. Optimization results for swing extension.

The first dataset produces results that are somewhat artificial due to differences between amputee and able-bodied gait, both in terms of hip motion and physical parameters of the lower leg. However, the optimized trajectories still allow us to visualize the effects of variations in user-specific data and are therefore useful to motivate the selection of generalized control laws in the following section. Results from the second dataset have a higher degree of legitimacy since the input signals and parameters are based on a real amputee individual wearing the same prosthesis as in the model, the only caveat being that the model cannot infer how the user will react to the new controller. In other words, the assumption that thigh motion remains unaltered is merely an approximation in both able-bodied and amputee cases.

5.3.4 Optimized and Baseline Behaviors

The optimization results presented in Fig. 5.4 and Fig. 5.5 as well as intuition gained from the general energy optimization problem in Appendix I are used to inform real-time control laws for stance flexion, swing flexion, and extension (controller states highlighted in green in Fig. 5.2). The baseline behaviors for flexion and extension (highlighted in yellow in Fig. 5.2) are presented alongside. More information on the theory, development, and validation of the baseline behaviors is presented in [27]. All behaviors (both baseline and optimized) consist of real-time damping commands to the prosthesis which are scheduled by readings of the various prosthesis sensors (i.e., are time-independent).

A schematic representation of the resulting behaviors is presented in Fig. 5.6. Knee angle and ground reaction force were obtained from amputee walking and are presented in the top plot for context since the various control laws (as will be shown) depend on the real-time values of those signals. The plots in the second and third rows in Fig. 5.6 show the commanded motor damping for the baseline and optimized cases at the knee and at the motor, respectively. As such, the only difference between the second and third rows is the angle-dependent transmission ratio.

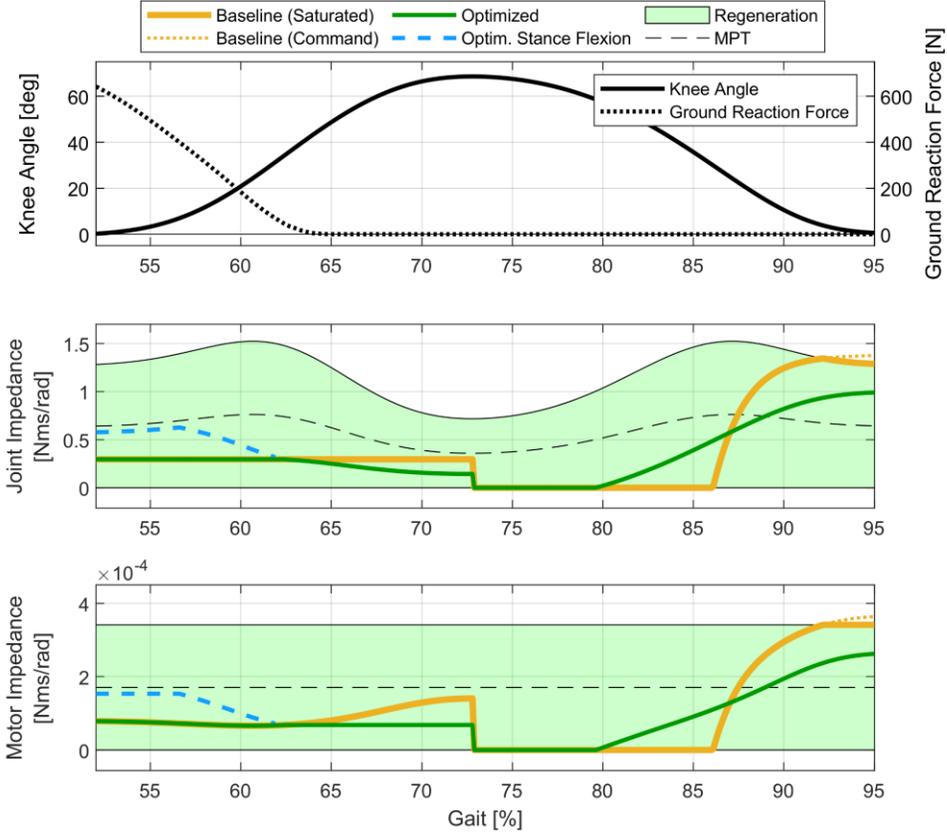


Fig. 5.6. Exemplification of baseline and optimized behaviors.

5.3.4.1 Stance Flexion

The proposed stance flexion control law is based solely on intuition of optimizing energy regeneration, since the developed dynamical model does not capture the behavior of the lower leg in stance flexion. It has been previously noted that the prosthesis user controls this phase of gait directly (unlike the swing phase, which is controlled indirectly through the inertial linkage of thigh and shank). Since the hip can produce about an order of magnitude more torque than the ECT in the low gear, the prosthesis motor may oppose motion in a greedy manner to maximize energy into the battery, and the user shall compensate accordingly, adapting their gait as necessary.

As is discussed in Appendix I, the “greedy” control law for energy regeneration consists of applying a maximum power transfer (MPT) damping command to the motor. In (5.5) is the proposed optimized stance flexion (*opt. st. flx*) control law, which consists of a motor damping command modulated by a normalized axial force signal \hat{F} which is presented in (5.6) as well as additional parameters β_0 and β_1 and walking speed estimate w (in m/s). The equations were constructed in such a way that a fixed motor damping (nominally MPT) is commanded while axial load $F > F_1$ and the prescribed swing flexion damping is commanded immediately before toe-off, when $F < F_0$. The command is linearly interpolated in-between (when $F_1 > F > F_0$) using F itself. Parameter F_0 has the same value as the axial load threshold used to transition the controller from stance flexion to swing flexion. The $sat(\cdot, u, v)$ function represents saturation with lower limit u and upper limit v . Blending of the damping command provides a smooth transition between stance and swing and helps match the level of resistance with the decreasing level of hip torque.

$$b_{opt.st.flex}^{mot} = \hat{F} b_{opt.sw.flex}^{mot} + (1 - \hat{F}) sat(\beta_0 + \beta_1 w, 0, 1) b_{MPT}^{mot} \quad (5.5)$$

$$\hat{F} = sat\left(\frac{F_1 - F}{F_1 - F_0}, 0, 1\right) \quad (5.6)$$

Parameters β_0 and β_1 and the walking speed estimate allow for tuning the magnitude of damping in stance flexion, giving some flexibility to this experimental control law if MPT resistance is inadequate for stance flexion. That is, the extraction of additional energy from this phase of gait need not occur with the MPT behavior (although that would be the ideal case), it suffices to increase the damping command towards MPT.

5.3.4.2 Swing Flexion

Motivated by the optimal behaviors found in the previous section, the swing flexion control law consists of a constant damping command at the motor. As can be seen in Fig. 5.6, the baseline behavior is in fact very similar, consisting of a fixed damping command applied at the knee joint and that serves for both stance and swing flexion. The baseline controller uses a relationship like in (5.7) to define that fixed damping command (“*bas. flex*”), consisting of the walking speed estimation w with scaling (γ_1) and offset (γ_0) parameters. In the optimized case, that value is sampled at toe off ($F < F_0$) and kept constant at the motor until peak flexion, which is shown in (5.8)

$$b_{bas.flex}^{knee} = \gamma_0 + \gamma_1 w \quad (5.7)$$

$$b_{opt.sw.flex}^{mot} = \frac{b_{bas.flex}^{knee}}{n(\theta_K)^2} \Big|_{toeoff} \quad (5.8)$$

Note that the baseline behavior will be used for stance flexion when the experimental stance flexion control law is not enabled.

5.3.4.3 Swing Extension

Similar to the baseline swing extension control law, the proposed behavior consists maintaining the motor initially deactivated and then ramping up damping until the end of swing (see Fig. 5.6). Despite structural similarities, however, the optimized behavior initiates the ramp-up significantly sooner and reaches a lower level of terminal damping compared to the baseline controller. Moreover, the shape of the damping curve is a fast first-order exponential in the baseline case, but more akin to a linear ramp in the optimized case (as prescribed by the optimization results).

The proposed behavior is presented in equation form in (5.9). The command is modulated by knee angle and scaled by both walking speed and an adjustable parameter δ_M . The product $\delta_M w$ ultimately determines the magnitude of damping achieved at the end of swing, when $\theta_K \approx 0$. Parameters δ_0 and δ_1 are respectively scale and offset the walking speed estimate w , which is then used alongside the peak flexion knee angle of the current stride ($\theta_{K(peak)}$) to modulate the point in the knee angle trajectory in which application of damping starts (i.e., when $\theta_K = \theta_{K(peak)} - \delta_1 w - \delta_0$).

$$b_{opt.sw.ext}^{knee} = \delta_M w \cdot sat\left(\frac{\theta_{K(peak)} - \delta_1 w - \delta_0 - \theta_K}{\theta_{K(peak)} - \delta_1 w - \delta_0}, 0, \infty\right) \quad (5.9)$$

Note that swing extension damping is commanded at the knee joint (the embedded system transforms the command to motor damping in real-time). Damping is commanded at the joint instead of the motor because the proposed behavior matches the optimization results more closely that way.

5.4 Experimental Evaluation

One amputee participant was recruited to test the proposed control laws with the ECT prosthesis. The participant is a male, 85 kg, 65 years of age, with transfemoral amputation on the right side. The amputation occurred in his youth for reason of trauma. The same participant validated the baseline controller and the powered features of the device, so he is well acquainted with the ECT and the existing walking controller. Testing consisted of two sessions of treadmill walking, the first one for training and finding adequate parameters for the control laws presented, and the second one for data collection. An annotated photograph of one of the trials is presented in Fig. 5.7(a). Data was collected using the prosthesis embedded system, consisting of knee angle and angular velocity, shank angle and angular velocity, shank axial ground reaction force, and motor and battery electric current.

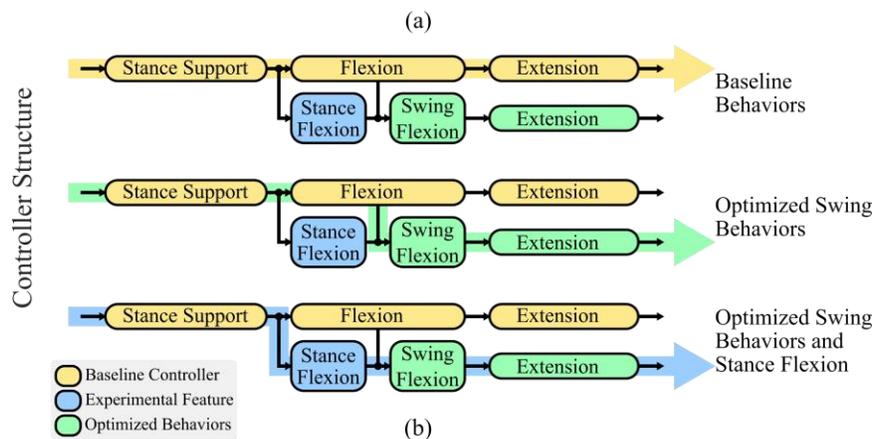
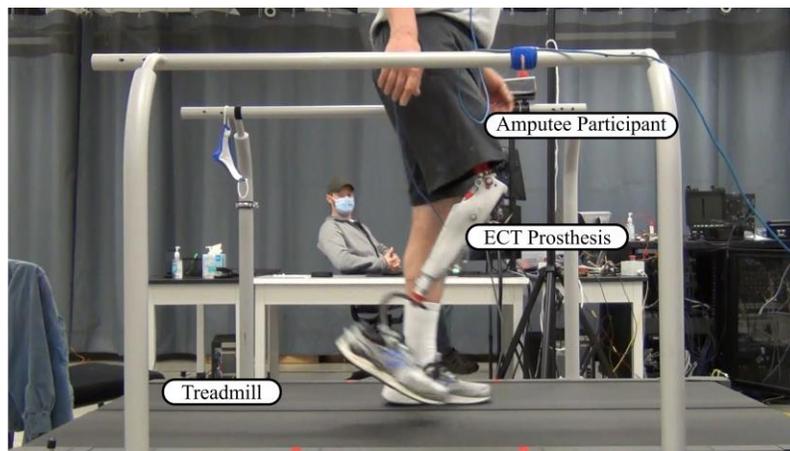


Fig. 5.7. A photograph of the amputee participant walking on the treadmill utilized in this study in (a) and a schematic representation of the three controller structures examined during the data collection.

All parameters were adjusted in the first session in order to achieve adequate walking behavior on speeds ranging from 0.8 to 1.2 m/s. The well-established baseline controller was used as a foundation to test and adjust optimized behaviors independently. That is, one by one, each baseline behavior was replaced by

each optimized behavior. Adjustments were made until the participant could successfully walk on any combination of behaviors. By the end, behavior combinations could be alternated in real-time, without warning the participant that a change was about to occur. In several cases, the participant could not notice the change or discern baseline from optimized. Changes in walking speed were likewise smooth, with the behaviors adjusting accordingly. Parameter values determined during testing are presented in Table 5.3. Even though swing flexion behaviors are different between baseline and optimized, the choice of values for parameters γ_0 and γ_1 served adequately for both.

Table 5.3. Parameter values for the optimized control laws found experimentally with the amputee participant.

Parameter	Value	Units
F_0	20	N
F_1	400	N
β_0	0	–
β_1	1	$(m/s)^{-1}$
γ_0	–0.6	Nms/rad
γ_1	0.9	$(Nms/rad)(m/s)^{-1}$
δ_M	1	$(Nms/rad)(m/s)^{-1}(deg)^{-1}$
δ_0	30	deg
δ_1	–20	$deg (m/s)^{-1}$

The second session occurred in the following week. First, the participant was given time to acclimate to the prosthesis with both controllers. The test protocol consisted of two sets of experiments: 1) automatically changing between the three combinations of optimized and baseline behaviors while the user walked on fixed treadmill speed settings and 2) fixed combinations of optimized and baseline behaviors with varying treadmill walking speed. As indicated in Fig. 5.7(b), three controller configurations were tested in each set of experiments: 1) all baseline behaviors, henceforth referred to as the “baseline” configuration; 2) baseline stance support and stance flexion with optimized swing flexion and extension, designated “optimized swing” configuration; and 3) baseline stance support with all remaining phases given by the proposed control laws developed in this work (stance flexion, swing flexion, and extension), designated “optimized stance and swing”.

5.4.1 Constant velocity experiments

For the first set of experiments, the real-time control software in the computer was used to automatically change the configuration of behaviors every 15 strides in the order of baseline, optimized swing, and optimized stance and swing. That protocol was followed for walking speeds of 0.8, 1.0, and 1.2 m/s. For each walking speed, the participant was given about 20 strides to reach steady state walking with the baseline controller before the automatic behavior transitions were activated. The first three strides after each transition are ignored in the analysis.

Average-stride results for this set of experiments are shown in Fig. 5.8. The angular trajectory of the knee matches closely between the three configurations shown, and all of them match able-bodied levels of peak knee flexion (data from [91]). Toe off happens at approximately the same point in the participant’s gait. Commanded damping at the knee is shown in the second row of plots, where the difference between the three configurations can be clearly seen, and compared to what is shown in Fig. 5.6. Electrical and mechanical power are shown in the bottom row of plots. Electrical power is estimated with the measured motor current and velocity (as shown in (5.3)) and nominal parameters for torque constant and phase resistance. Mechanical power includes the estimated motor and friction torques, as in (5.1). Note that the

power convention is positive for power delivery from the battery to the joint, so both power signals being negative indicates energy regeneration. The baseline and optimized swing controllers have similar power profiles in flexion, but notably different profiles in extension (since damping is applied later and with higher magnitude in the baseline controller case). Enabling the experimental stance flexion law only causes a material difference for the lower walking speed and does not appear to compromise the participant's gait, even though additional mechanical power is demanded.

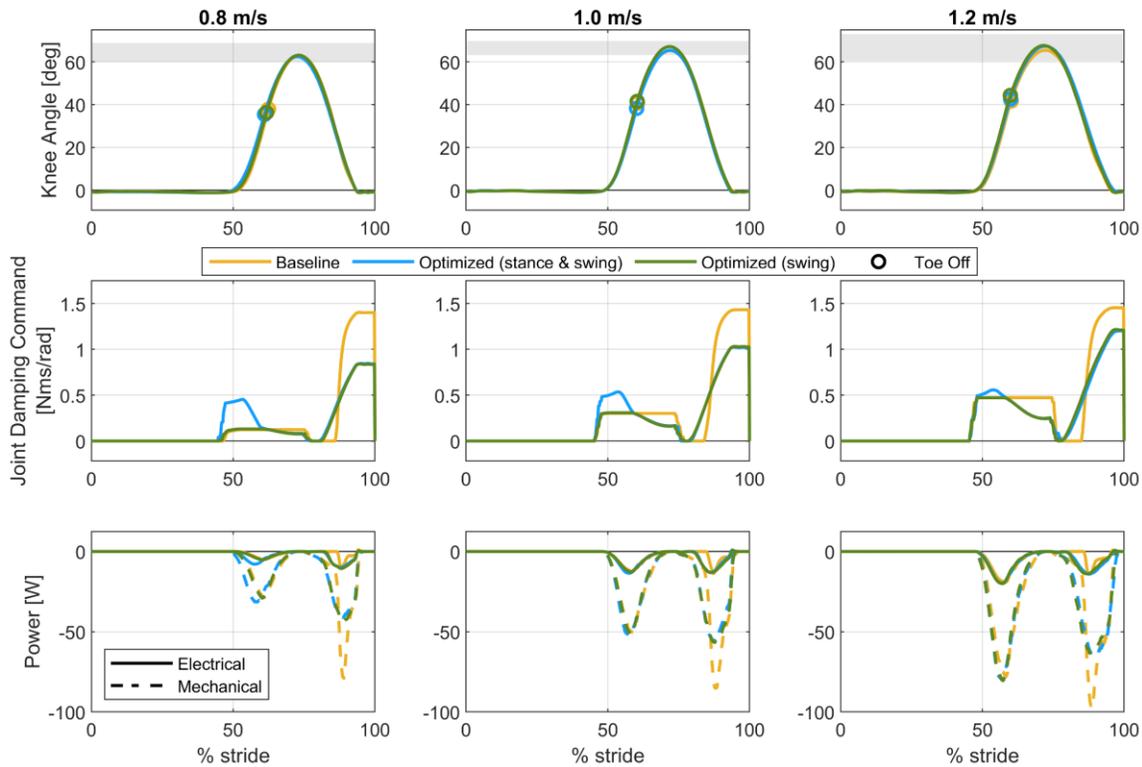


Fig. 5.8. Average-stride results for three fixed treadmill walking speeds (0.8, 1.0, and 1.2 m/s) and three controller configurations: baseline in yellow, optimized swing in green, and all optimized behaviors, including stance flexion, in blue.

5.4.2 Variable treadmill speed experiments

In the second set of experiments, the treadmill speed was varied every five strides, following a staircase pattern: 0.8, 1.0, 1.2, 1.1, and 0.9 m/s. The same three controller configurations were tested as in the first set of experiments. Results for this set of experiments are shown per stride, where the x-axis in Fig. 5.9 is the estimated walking speed as determined by the controller. The plots show three biomechanical metrics: peak knee flexion, stride duration, and swing phase duration. The only perceivable difference between the baseline and optimized controllers is a moderate increase in the peak flexion for faster walking speeds (about 4 degrees at 1.2 m/s). However, the variation stays within the able-bodied reference range.

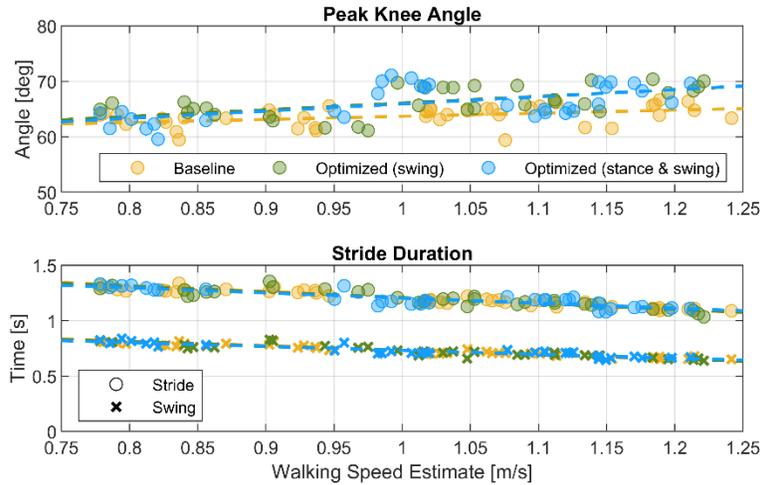


Fig. 5.9. Results of the set of experiments with varying walking speed. Biomechanical characteristics (peak flexion angle and stride timing) are shown per stride.

5.4.3 Energy Regeneration

A breakdown of mechanical and electrical energy per stride is shown in Fig. 5.10. Mechanical energy estimates are shown at the joint (including friction and motor torque) and at the motor (only motor torque). The trendlines show some difference in joint energy per stride between controller configurations, in the order of 0.5 J for slow walking speeds and up to 1 J for faster walking speeds, with the baseline controller always demanding less power. Note, however, that the observed difference of about 5% in relative terms is considerably less than the variance in the data. The three controller configurations demonstrate almost indistinguishable mechanical energy aggregates per stride at the motor. The ratio of mechanical energy at the motor and joint indicates the reverse drive efficiency of the system, between 30 and 40%. That is, about two thirds of the mechanical energy in the system is dissipated by drive friction.

Energy regeneration is presented next and can be used to infer the motor efficiency by taking the ratio of electrical to mechanical motor energy, calculated as between 50 and 60% for the optimized control laws and between 30 to 40% for the baseline controller. That is, about half the energy that is not dissipated as heat due to friction is dissipated as heat in the motor windings. Note that regeneration is considered negative power with the sign convention adopted in this work, so greater negative magnitudes indicate more regeneration. Regeneration is also further broken down among the three phases of gait that were optimized: stance flexion, swing flexion, and swing extension. As expected, the optimized control law with the experimental stance flexion behavior leads to more regeneration during that phase of gait, particularly in slow walking speeds (additional 0.35 J relative to optimization of swing alone at 0.8 m/s), but the effect is diminished for faster walking speeds. Regeneration in swing flexion is unaffected by the optimized control laws. Regeneration in swing extension, however, is significantly improved. The optimized behavior doubles energy regeneration in that phase of gait (between 0.4 and 0.7 J for the baseline controller and between 1 and 1.7 J for the two optimized cases).

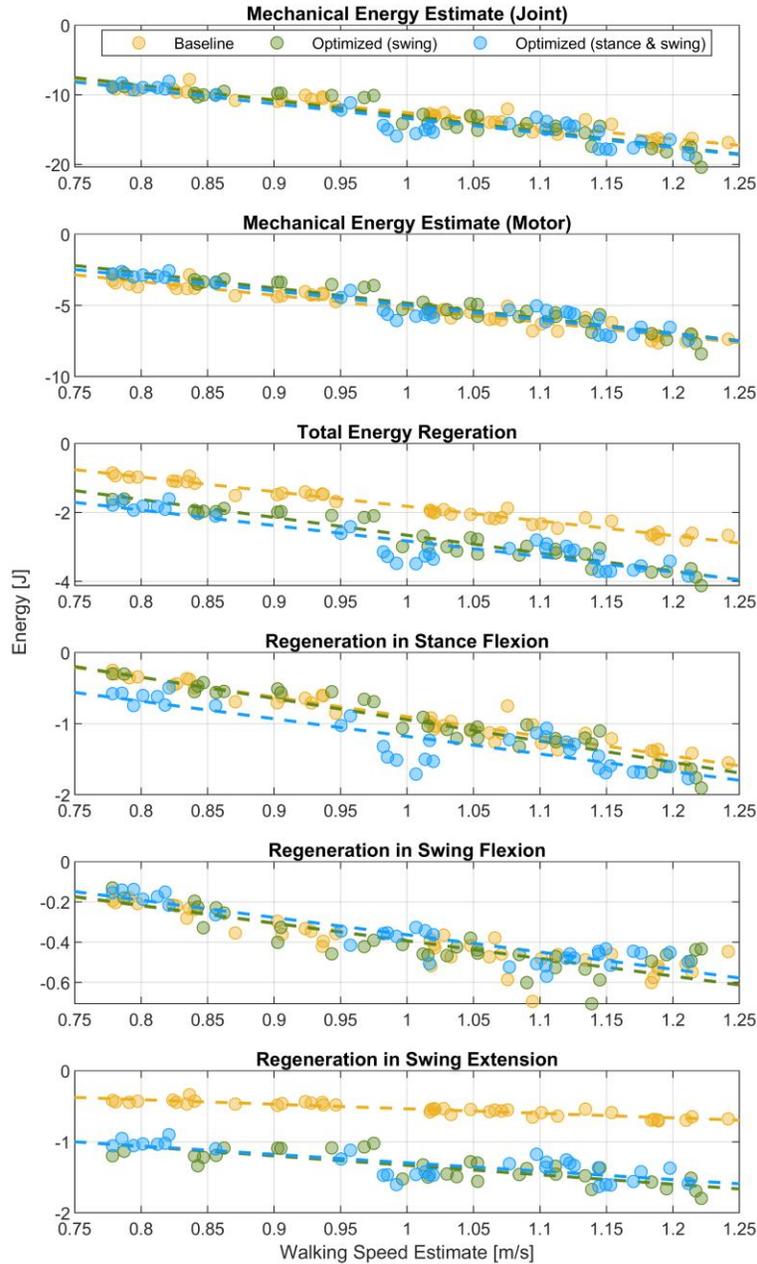


Fig. 5.10. Joint and motor mechanical energy and total motor electrical energy (regeneration) per stride in the first two rows of plots and a breakdown of energy regeneration per phase of gait in the bottom three rows.

The aggregate efficiency of the system, from mechanical energy at the joint to electrical energy at the motor leads is presented per stride in Fig. 5.11. The optimized control laws raise the efficiency metric from about 10% for slow walking speeds to 20% overall. Therefore, with the proposed controllers, one can expect to regenerate into the battery of the prosthesis up to 20% of the mechanical energy input at the knee joint (recall that two thirds of the mechanical energy is dissipated as heat due to friction, regardless of the control law, so efficiency has a ceiling at about 30%).

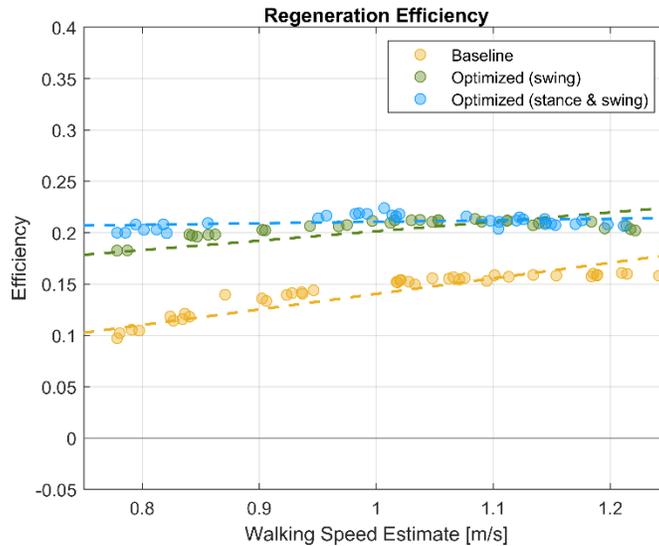


Fig. 5.11. Per-stride overall regeneration efficiency of the system with the three controller configurations (from mechanical energy at the joint to electrical energy output of the motor).

5.4.4 Battery power

Regeneration results were presented based on the estimated electrical power delivered by the motor to the system (see definition of P_E in (5.3)). Since the embedded system senses battery current in real-time, results can also be presented in terms of measured battery energy. The estimated values were used up to this point because the battery also powers other system-level components, hindering interpretation of the effect of each control law. Specifically, battery current includes current to power the embedded system and sensors (around 1.2 W or 75 mA continuously) and current to power the clutches. Crucially, as a provisional solution to clutch slippage that had been identified as an issue early into this project, the carrier clutch duty cycle command is scheduled with motor current. In that way, clutch force increases to accommodate motor torque and maintain engagement, but the solution leads to a significant (and variable) amount of energy expenditure.

Average-stride plots of motor, clutch solenoids, and battery current are presented in Fig. 5.12 for the three walking speeds in the first set of experiments. Clutch operation consists of a brief spike period (few milliseconds) which is achieved by commanding 100% duty cycle to the power semiconductor components used to control the solenoids, followed by a continuous hold period with lower duty cycle. As discussed in sections 2 and 3.1, the ring clutch is activated at step initiation and deactivated at the transition of flexion to extension. The carrier clutch is activated at the same time but remains active until the end of swing. The absolute value of motor current reference modulates the hold duty cycle of the carrier solenoid. Hold current in the ring clutch is about a quarter of an Ampere. Hold current in the carrier clutch peaks at about 1.5 A and averages close to 1 A.

It is clear from the plots of battery current that the combination of static demand from the embedded system and the variable demand from solenoids completely thwarts regeneration at the battery, even though the motor is always regenerating (as per the discussion of passive motor control in section 2 and results in Fig. 5.8). Battery power demand is lowered with increasing walking speeds, as regeneration at the motor regeneration is increased. Even though motor current (and therefore the magnitude of torque

which opposes motion) has higher magnitude with the baseline controller, that does not contribute to lowering the battery current demand because the behavior is achieved at low efficiency.

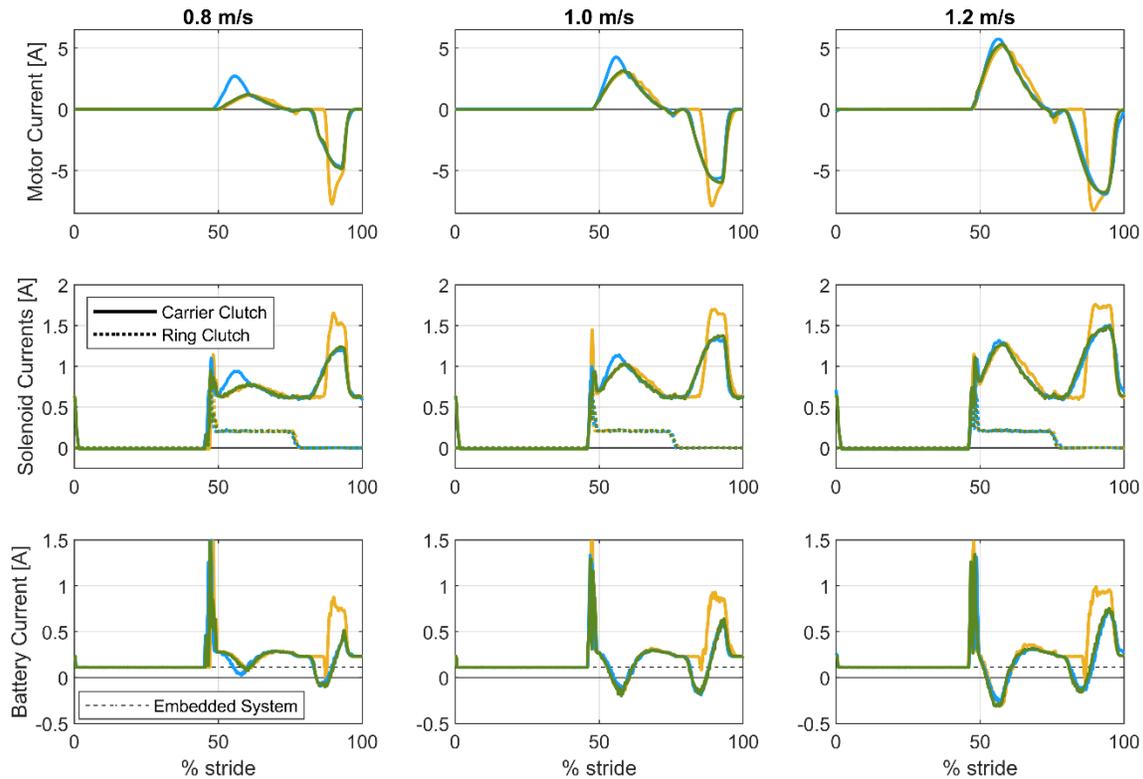


Fig. 5.12. Average-stride electrical current signals for the motor, solenoids, and battery.

Finally, battery energy per stride from the second set of experiments is presented in Fig. 5.13 along with a breakdown of the three system components that contribute to regeneration or cause battery energy expenditure: motor, embedded system along with its sensors, and clutches. The difference between baseline and optimized control laws is similar at the motor and at the battery: about 1 J per stride lower energy demand with the optimized control laws. Note that while the stance optimized configuration causes more regeneration, as previously discussed, that additional power is not reflected in the battery because the higher motor current in stance flexion causes an increased demand of power in the carrier clutch.

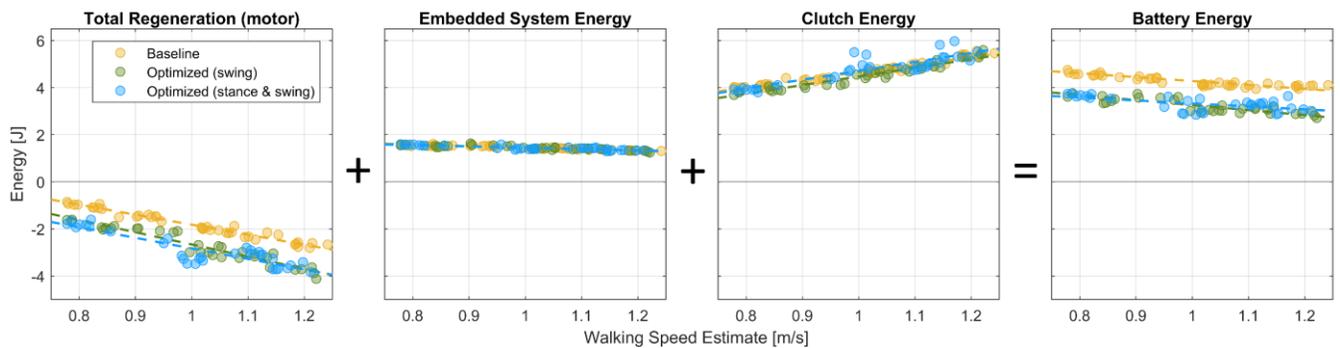


Fig. 5.13. Breakdown of battery energy per stride with varying treadmill walking speeds. The leftmost plot of total regeneration at the motor contains the same information as presented in Fig. 5.10.

5.5 Discussion of Results

The ability to regenerate energy in level walking at comfortable walking speeds gives battery-powered prosthesis a second dimension to the tradeoff between battery life and weight and size of the battery pack. When regeneration is optimized, such device may be designed with less battery capacity, resulting in reduced device size and weight, while still providing the same battery life to the user (given some assumed level of activity). Alternatively, if the battery pack remains the same, the device would gain in battery life, affording a greater level of independence to the user. Regardless of which approach interests the designer, the energy regenerating device has a clear advantage over non-regenerating MPKs.

This work presented a method to design modulated-passive damping control laws informed by optimal control techniques that demonstrate improved energy regeneration. The “modulated damping” control laws make the device operate as a viscous damper which is modulated by knee angle and estimated walking speed. This control approach supports behavior predictability and puts the user in control of the device. The device does not have any pre-programmed trajectories for the knee joint and does not need to estimate the user’s intentions, it simply responds with torque to the user-imposed movement. Despite not enforcing a trajectory, the optimization was constructed with goals to synergistically help the user achieve healthy gait kinematics; namely, sufficient peak knee flexion for ground clearance, and comfortable deceleration in terminal swing.

In the presented experimental evaluation, the overall regeneration efficiency was increased by about 10% over baseline behaviors (going from 10 to 20% efficiency in low speeds and from 15 to 20% efficiency in high speeds) with no appreciable change to walking kinematics and almost identical mechanical power requirements (serving as a strong indication that the user was not exerting additional effort to walk with the optimized control laws). While the motor can regenerate between 1.5 and 4 J per stride with the optimized control laws (was between 1 and 3 J with the original controller), the elevated energy requirement of the clutches (between 4 and 6 J) and the modest energy requirement of the embedded system (around 1.2 W continuous) make the net energy requirement of the overall system remain positive even with the optimized behaviors (between 3 and 4 J per stride). The values here provided are based on the experiments with a male participant weighing 85 kg, there would likely be less regeneration for an individual that weighs less, but power consumption by the embedded system and clutches would remain the same. While regeneration by the motor is not yet enough to achieve net-zero energy walking, it effectively cuts battery power consumption in half.

Experiments have shown that, for walking speeds below 1.0 m/s, the system is able to extract additional energy from the user with the exploratory stance flexion control law which raises the damping command, bringing it closer to MPT during the double support phase. That seems to have only a minimal impact on the mechanical energy required at the knee joint and effects a minor increase in regeneration efficiency of 2 to 3% for those walking speeds. Swing initiation in healthy gait has contributions from both hip and ankle (i.e., push-off) power [92] but the ankle contribution does not occur to the same extent in amputee gait (which has well-known negative effects [7]). Even though knee joint power is negative during this phase of gait (mostly due to the negative power contribution of the rectus femoris in able-bodied individuals), the net contribution of the entire leg musculature to movement of the individual’s center of mass is positive [93]. Therefore, it is unclear if late stance should be used to harness more power than prescribed and doing so may lead to aggravation of the user’s pathological gait, additional compensatory actions, and increased metabolic cost of walking. Within this context, the authors believe it is hard to

justify the additional complexity and potential for additional user effort given such limited gains. However, the increased energy harnessing could be desirable when the state-of-charge of the battery crosses some low threshold. Note that the additional torque required at the motor corresponds to additional carrier clutch activation energy due to the clutch slippage mitigation scheme, so this control law has no net effect on the battery power requirement in the current implementation of the system.

The walking speed estimate employed in the control laws (and which was used to plot the data in Figs. Fig. 5.9, Fig. 5.10, Fig. 5.11, and Fig. 5.13) is calculated from the average value of shank angular velocity during stance support, scaled by a user-specific parameter. That estimation approach leverages the inverted pendulum characteristic of the prosthetic stance support phase and has been shown in other works to represent walking speed with high fidelity [27]. The user-specific scaling parameter was found experimentally by fitting the data from the first set of experiments, which occurred with three fixed treadmill speed settings.

During experimentation, the duty cycle of the carrier clutch had to be significantly increased to address clutch slippage. A control law was defined to modulate the carrier clutch duty cycle based on commanded motor current, such that the clutch holding force is adjusted according to motor torque. We estimate that the activation of the carrier clutch required at least four times more power per stride than the activation of the ring clutch. If the carrier clutch is redesigned such that it can be activated with a similar amount of power as the ring clutch, the energy requirement of the clutches could be reduced to the point that net-zero energy walking would be achievable at elevated walking speeds.

Although the prototype prosthesis already has low intrinsic friction, further reductions in friction would assist in increasing the efficiency of energy harnessing. Note that the improved regenerating efficiency (estimated at 20% from joint to motor leads) is still considerably low relative to other research prototypes recently presented (e.g., estimated around 60% for the device in [25] and around 73% for the device in [26] when operated in “passive” mode). The lead screw used in the crank slide mechanism is likely the main culprit – it is well-known that lead screws have low backdriving efficiency (the devices in [25], [26] employ ball screws) – so the actuator could be redesigned to use a higher efficiency screw or some other transmission architecture. Different transmission elements will trade various design and performance parameters (such as inertia, friction, torque capability, life, cost, and audible noise) so such redesign would not be a straightforward task. Another element of the system that could be altered, and which may be an easier task to accomplish, is the selection of a larger motor along with a lower transmission ratio in the actuator, which may help reduce friction in the actuator and increase efficiency. Somewhat counterintuitively, the current implementation of the prototype in fact benefits from some level of friction because the remaining torque required for swing can be produced by the motor while staying within the band of impedances in which regeneration occurs [3]. Therefore, care must be taken to avoid motor plugging (reverse current braking) if friction is significantly reduced.

Two important limitations of the results presented in this paper are that 1) all experimentation was performed while walking on a treadmill and 2) the experiment was conducted with only one amputee participant. Since the control laws established in this work rely on the same conditions for state transition and demonstrated equivalent gait kinematics relative to the baseline controller, it is reasonable to expect that the optimized control laws would perform just as well as the baseline control laws in overground walking (as reported in [27]). Limiting the study to a single participant facilitated development and testing cycles, but care must be taken to avoid overfitting to a particular individual.

5.6 Conclusion

This work presented a new approach to optimizing energy regeneration in knee prostheses. We build upon the idea of a biomimetic, trajectory-free ballistic swing phase in which the lower leg swings like a pendulum both kinematically and kinetically in accordance to the motion of the user's thigh [15], [89]. In this context, a prosthesis provides some resistive behavior to achieve kinematic goals such as preventing exaggerated flexion and providing a comfortable extension. While many (perhaps infinitely many) profiles of resistance can meet those goals, we demonstrate that behaviors informed by optimal control techniques can additionally realize greater energy regeneration and reduce battery energy demand in a motorized prosthesis.

Experiments with an amputee participant walking on a treadmill showed that the selection of energy-optimized control laws for stance flexion, swing flexion, and swing extension can increase energy regeneration substantially in absolute and relative terms. We report an overall increase of about 1 J in regeneration per level walking stride relative to a baseline control law presented in [27]. Likewise, the efficiency of the system from mechanical energy at the knee joint to motor electrical energy per stride doubled, from about 10% to about 20% for slower walking speeds (around 0.8 m/s), and also increased, albeit not as significantly, for faster walking speeds. This doubling of efficiency was achieved with practically no variation in gait patterns, almost identical mechanical power at the knee joint, and with the participant hardly able to notice any change.

Ultimately, two thirds of the mechanical energy removed by the knee joint is always dissipated as heat due to friction in the drive of the prototype device used in this work. Therefore, the optimized control laws can regenerate a significant portion of the remaining third that reaches the motor. In absolute terms, the motor regenerates between 1.5 and 4 J per stride with the optimized control laws (increasing with walking speed). However, the clutches used to configure the actuator transmission and the embedded system in the prosthesis consume approximately twice that amount due to implementation challenges with clutch slippage, so the prosthesis does not achieve net zero energy as whole yet.

5.7 Appendix I: Optimizing Energy Regeneration in Dynamical Systems

A mechatronic servo system is shown in Fig. 5.14 in bond graph form representing a first-order mass-damper with motor actuation under an external backdriving effort source. As presented, the bond graph can be viewed as a generalization of any typical motor-driven system with regeneration capabilities (e.g., automotive or robotic). In Fig. 5.14, i is the motor current, u is the control signal that applies portion of the battery voltage across the motor leads ($-1 \leq u \leq 1$), \dot{q} is the angular velocity of the motor, and τ is a backdriving torque. Lumped parameters R , k , b , and I represent electrical resistance, motor torque constant, viscous friction, and inertia in the system, respectively. The motor torque originated from the electromotive phenomenon and associated back electromotive force are given by the gyrator constitutive equations, ki and $k\dot{q}$, respectively. Note that b , I and τ are analogous to a Thévenin equivalent structure for the entire mechanical system (reflected onto the motor output shaft). High-level control tasks are achieved by setting a motor current (i.e., torque) reference i_{REF} according to some control law which can depend for instance on velocity \dot{q} , position q , and predefined task goals. Torque control in the motor consists of setting u in order to track the reference current i_{REF} . In the bond graph, that is represented by modulating a transformer between the electrical energy storage system (e.g., battery) and the motor leads; in practice, that is accomplished with a control bridge and commutation, where the control signal is a duty cycle to power semiconductor devices. The motor windings also exhibit inductance which is

ignored for the purposes of this analysis because the corresponding RL circuit time constant is typically orders of magnitude that of the mechanical system. The equations for the two 1-junctions are:

$$uv_B = Ri + k\dot{q} \quad (5.10)$$

$$\tau = -ki + b\dot{q} + I\ddot{q} \quad (5.11)$$

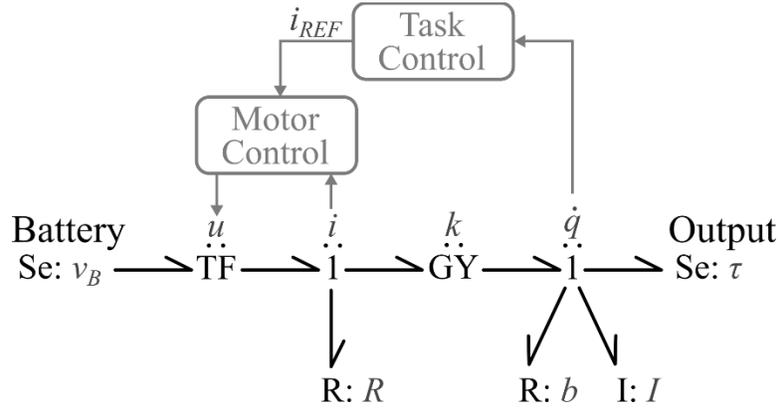


Fig. 5.14. Bond graph of a generalized servo control system with a DC motor capable of energy regeneration into a battery. High-level control tasks are achieved with a cascading feedback control structure in which the kinematic state variable (\dot{q}) is used by a high-level controller to define a motor current (i.e., torque) reference. The current command is then implemented by motor current control (inner loop). The mechanical system is characterized by intrinsic inertia and friction (damping in this linear model) and receives effort inputs from the motor and from an external, independent τ .

We momentarily ignore the effect of inertia I to study the steady-state (i.e., $\ddot{q} = 0$) power transfer and efficiency characteristics of the system between the effort source (τ) and the battery. Power exchange between battery and motor is given by $P_E = uv_B i = (Ri^2 + k\dot{q}i)$ and power between the external effort source and the motor is given by $P_M = k\dot{q}i$. The behavior for which the maximum amount of energy will be regenerated by the motor into the battery can be found by minimizing P_E with respect to the manipulated variable (motor current), resulting in the maximum power transfer condition in (5.12). Likewise, the behavior that leads to peak regeneration efficiency can be found by maximizing the ratio of powers in (5.13).

$$\min_{i(t)} P_E \Rightarrow i_{MPT} = -\frac{k\dot{q}}{2R} \quad (5.12)$$

$$\max_{i(t)} \frac{P_E}{P_M} \Rightarrow i_{EFF} = -\frac{\dot{q}}{k} \left(\sqrt{b^2 + \frac{k^2 b}{R}} - b \right) \quad (5.13)$$

In both cases, the idealized behaviors can be achieved with an impedance control law in the high-level control block without loss of generality:

$$i_{REF} = -\frac{B\dot{q}}{k} \quad (5.14)$$

where B is a damping coefficient that can be configured to either the MPT behavior ($B_{MPT} = B_E/2$) or peak efficiency behavior ($B_{EFF} = \sqrt{b^2 + B_E b} - b$). B_E is the behavior of the motor when leads are short-

circuited ($u = 0$) which corresponds to the reflection of the winding resistance R across the gyrator, $B_E \triangleq k^2/R$. It can be shown by substituting (5.14) into P_E that energy regeneration ($P_E < 0$) only occurs for $0 < B < B_E$ [3], [4], [12], [24].

It is, however, unclear from these results which behavior (B_{MPT} or B_{EFF}) leads to higher energy regeneration in practice in a system in which velocity is subject to internal dynamics and external action, and reasonable arguments can be made for either. Both options have been proposed in works with the goal of optimizing regeneration in prosthetic systems [3], [4], [20], [21], [23], [24] but no clear distinction has been presented in any of those works. As will be presented, both behaviors are in fact useful, and the distinction becomes clear when the dynamics are taken into account. Therefore, we treat the optimization problem under a time integral, optimizing energy into the battery within a finite time horizon, $t = 0$ to T .

$$E = \int_0^T P_E dt \quad (5.15)$$

The optimization goal is to minimize E , defined in (5.15), which maximizes energy regeneration, since P_E is negative in that regime. System dynamics serve as constraints in the optimal control sense and calculus of variations is used to find an optimal trajectory $i^*(t)$ with Pontryagin's principle [94]. A fixed horizon T is used here due to the requirements of coordinated motion involved in prosthesis control (i.e., the controller of a prosthesis does not have agency to affect the duration of each gait phase). The Hamiltonian for this optimization is given in (5.16) with associated co-state variables in (5.17).

$$\begin{aligned} \mathcal{H} &= P_E + p_1 \dot{q} + p_2 \ddot{q} \\ &= Ri^2 + k\dot{q}i + p_1 \dot{q} + \frac{p_2}{I}(ki + \tau - b\dot{q}) \end{aligned} \quad (5.16)$$

$$\begin{cases} \dot{p}_1 = -\frac{\partial \mathcal{H}}{\partial q} = 0 \\ \dot{p}_2 = -\frac{\partial \mathcal{H}}{\partial \dot{q}} = -ki - p_1 + p_2 \frac{b}{I} \end{cases} \quad (5.17)$$

In (5.17) we assume that τ is independent of the system state (i.e., no external controller is commanding torque based on the system state). Then, the optimal control law consists of finding i that minimizes the Hamiltonian with all other variable held constant, which yields the optimal solution i^* in (5.18).

$$\frac{\partial \mathcal{H}}{\partial i} = 0 \Rightarrow i^* = -\frac{k}{2R} \left(\dot{q} + \frac{p_2}{I} \right) \quad (5.18)$$

Substituting i^* into (5.11) and (5.17) produces the optimally controlled system which effectively is a second order linear system (dynamics are nonzero for \dot{q} and p_2). The boundary conditions are defined at $t = 0$ for q and \dot{q} , representing the initial configuration of the system, $q(0) = q_0$ and $\dot{q}(0) = \dot{q}_0$, and at $t = T$ for p_1 and p_2 as a result of the Pontryagin's principle formulation, $p_1(T) = 0$ and $p_2(T) = 0$ [94]. With these boundary conditions, the optimum motor behavior can be found in closed form given in terms of an optimal emulated damping B^* trajectory as shown in (5.19).

$$B^* = -\frac{ki^*}{\dot{q}} = B_{MPT} \left(1 + \frac{p_2}{I\dot{q}} \right) \quad (5.19)$$

Note that B^* varies with time, unlike the previous MPT and maximum efficiency cases in which the behaviors were fixed. The boundary condition $p_2(T) = 0$ guarantees that the optimal damping control

law will converge to B_{MPT} as $t \rightarrow T$. The behavior of B^* is demonstrated in Fig. 5.15, with the time remaining in the optimization window ($T - t$) chosen as the independent variable for plotting B^* . Note that B^* always converges to the MPT behavior at the end of the optimization window, regardless of the length of the optimization window or the level of torque provided by the external force. The value of B^* at $t = 0$, however, depends on both of those terms. In the homogeneous case, $B^*(0)$ converges to B_{EFF} as T grows.

Recall that B_{MPT} was determined in (5.12) by optimizing only electrical power, without regard to mechanical power. As such, the MPT behavior can be interpreted as a greedy strategy, that quickly depletes the mechanical energy in the system in order to get as much electrical energy as possible without regard to how efficiently the energy transfer occurs between domains. Having B_{MPT} as the terminal behavior in the optimization window makes sense; the optimizer starts with operation at high efficiency to limit losses and transitions to MPT to try to extract as much energy as possible out of the system before the optimization window ends.

Another important factor necessary to interpret the optimization results is understanding how fast the prescribed behavior will converge towards B_{MPT} . While closed form expressions quickly become unwieldy and difficult to interpret, we can analyze the eigenvalues λ of the linear system formed by the coupled dynamics of \dot{q} and p_2 , which are presented in (5.20). The real and opposite eigenvalues correspond to a saddle point-style system (instability is not a concern here due to the fixed optimization window). As shown, the magnitude of the eigenvalues is given by a balance between the mechanical system time constant and the motor size constant k_M , where $k_M^2 = 2B_{MPT}$.

$$\lambda = \pm \frac{\sqrt{b}}{I} \sqrt{k_M^2 + b} \quad (5.20)$$

The eigenvalues of the optimized system dictate how quickly the optimized behavior converges to B_{MPT} towards the end of the optimization window. If $|\lambda|$ is large, the prescribed optimal behavior starts close to B_{EFF} and lingers before quickly converging to B_{MPT} at the end of the optimization window; if $|\lambda|$ is small, then the prescribed behavior starts further from B_{EFF} and slowly converges to B_{MPT} . Analyzing the balance between mechanical system time constant and motor size, we find that if the motor is large or powerful relative to the system, then the optimized behavior is that of a large $|\lambda|$ (remain closer to B_{EFF} for longer before quickly converging to B_{MPT}); conversely, if the motor is small or weak relative to the system, then the behavior is that of a reduced $|\lambda|$ (start closer to B_{MPT} and gradually converge to B_{MPT} throughout the optimization window).

In summary, we can subjectively conclude that the strategy taken by other works of maintaining the motor at $B = B_{MPT}$ can be appropriate for cases when the motor is weak relative to the mechanical system or when the time horizon for optimization is too short for the motor to efficiently affect the system. The strategy of optimizing for efficiency with $B = B_{EFF}$ is better suited for cases where the motor is strong relative to the system or when the optimization window is long, but the behavior should always converge to B_{MPT} at the end of the optimization horizon.

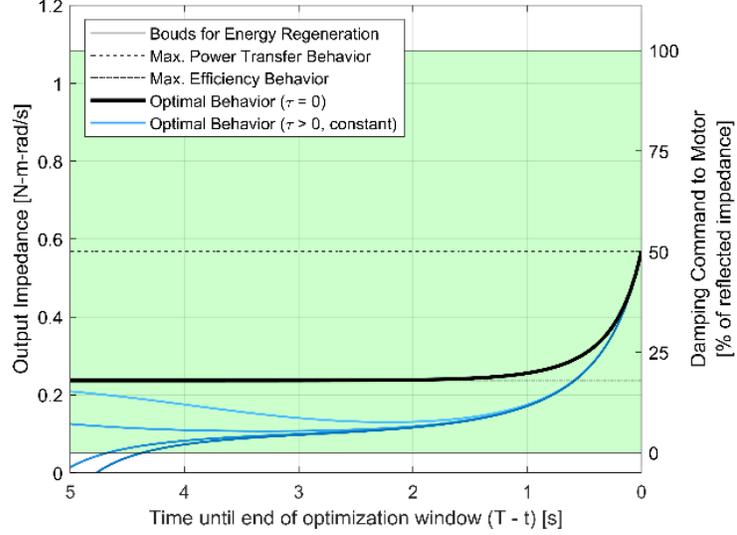


Fig. 5.15. Exemplification of optimized behaviors in terms of system impedance for the homogenous case ($\tau = 0$) and for increasing levels of constant backdriving force. Results are presented as a function of the time remaining in the optimization window (i.e., as a countdown to the end of the optimization window).

5.7.1 Backdriving effort source

In the homogeneous case, the motor regenerates a portion of the initial energy content of the mechanical system (given by \dot{q}_0), however the motor can additionally regenerate energy provided by the external effort source, represented by τ in Fig. 5.14. Recall that the optimal result assumes that the force is external and independent of the state of the system, but the opposite is not true. That is, the optimal motor behavior is not independent of the force. In fact, the solution and its optimality criteria require modulation of the motor behavior according to past and future $\tau(t)$. That is, motor behavior does not adapt in real-time to τ but rather must be calculated ahead-of-time with knowledge of the behavior of τ within the entire optimization horizon (which introduces challenges to real-time implementation of this method). The terminal condition $p_2(T) = 0$ still holds, so regardless of τ , it still holds that $B^*(T) = B_{MPT}$. Through simulation, it can be observed that the optimal behavior B^* with a constant τ that assists the motion starts with a lower magnitude damping reference than in the homogeneous case and converges to the same MPT point. Five of such cases, with varying magnitude of constant τ are presented in Fig. 5.15.

5.7.2 Kinematic Goals

Tracking a desired position or velocity is generally the main objective of motor control in a mechatronic system, so optimizing energy regeneration without concern for kinematic goals would drastically limit the applicability of this analysis. Achieving kinematic goals is very important for the control of prostheses, as is well-known, to ensure user safety and biomimetic behavior. We therefore expand the analysis to allow for a tradeoff between energy regeneration and desired terminal kinematic configuration. Specifically, equation (5.21) shows the new cost function J , which contains the original cost defined in (5.15) with added quadratic terms due to errors in achieving kinematic goals for position and velocity at $t = T$.

$$J = \overbrace{\frac{\sigma_1}{2} (q(T) - q_{DES})^2 + \frac{\sigma_2}{2} (\dot{q}(T) - \dot{q}_{DES})^2}^{\Phi} + \overbrace{\int_0^T P_E dt}^E \quad (5.21)$$

In (5.21), q_{DES} and \dot{q}_{DES} are the desired end position and velocity at $t = T$, respectively, and σ_1 and σ_2 are corresponding adjustable gains that can be chosen to balance the aforementioned tradeoff. The aggregate cost associated to the error to desired kinematic outcomes is expressed as Φ . With this formulation, the only change to the optimized system is in the boundary conditions for p_1 and p_2 [94] as shown in (5.22).

$$\begin{cases} p_1(T) = \frac{\partial \Phi}{\partial q(T)} = \sigma_1(q(T) - q_{DES}) \\ p_2(T) = \frac{\partial \Phi}{\partial \dot{q}(T)} = \sigma_2(\dot{q}(T) - \dot{q}_{DES}) \end{cases} \quad (5.22)$$

Now, since $p_2(T)$ is no longer necessarily zero, the optimum B^* will not necessarily converge to B_{MPT} . Moreover, since $\dot{p}_1 = 0$, we have that $p_1(t) = p_1(T) \forall t$, which affects \dot{p}_2 and therefore also modifies B^* . Those observations can be better visualized with examples, presented in Fig. 5.16 for the cases of position and velocity kinematic goals.

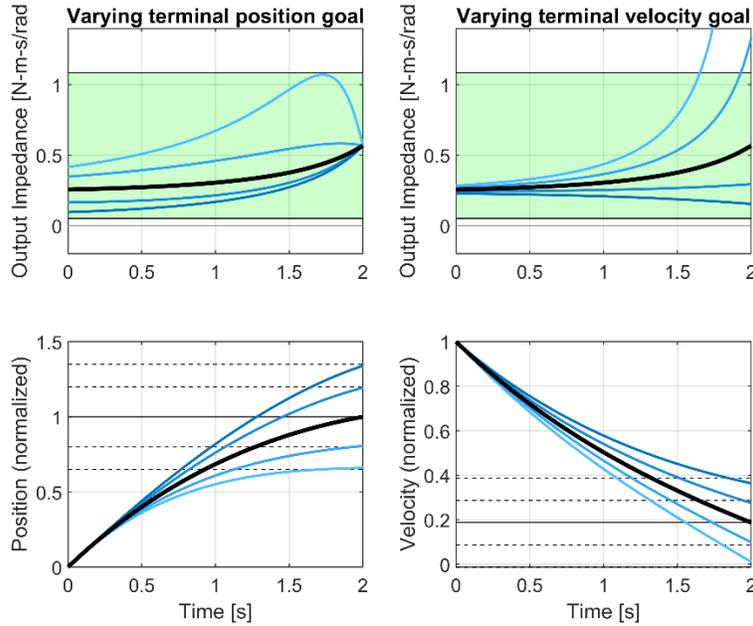


Fig. 5.16. Examples of optimized behaviors with kinematic goals. Varying the terminal position goal (plots on the left) affects the value of $B^*(0)$ but $B^*(T)$ remains at B_{MPT} . In order to achieve more displacement during the regenerating task, one should modulate the system behavior to initiate below B_{EFF} , and vice-versa. Varying the terminal velocity goal (plots on the right) affects $B^*(T)$ but $B^*(0)$ remains unchanged. In order to achieve higher terminal velocities, one should keep the optimized behavior below B_{MPT} , and vice-versa.

5.8 Appendix II: Applying the optimal regeneration framework to the prosthetic system

The discussion of optimizing energy regeneration in a dynamical system with specified kinematic goals in Appendix I was based on a simplified mechatronic system with generalized impedance and a backdriving effort source independent of the system state. In this appendix, the developed framework is applied to the prosthetic system presented in section 5.3.2. Specifically, we treat a few main characteristics of the lower leg system that are not covered by the generalized system (with required adaptations to the cost

function, the Hamiltonian, and the co-state variables). Then, we present in detail the numerical “shooting method” used to perform the optimization and produce the results presented in the main text.

The main differences between the simplified mechatronics system of Appendix I and the system comprised of the ECT prosthesis in swing phase are as follows: 1) the backdriving effort source is not independent of the system state since the application of gravity and inertial forces depend on the lower leg orientation; 2) energy regeneration depends on motor velocity which is given by the knee angular velocity, however knee angular velocity is not a state variable; 3) kinematic goals are expressed in terms of knee angle and angular velocity, which again are not state variables; and 4) the transmission ratio is not constant, depending on knee angle. These differences affect the definitions of the cost function, the Hamiltonian, and the co-state variable dynamics as well as their associated terminal values.

First, the expanded cost function which is presented in terms of the generalized state variable q in (5.21) is rewritten in terms of the knee angle θ_K and P_E according to (5.3). Recall that knee angle is calculated from the subtraction of thigh and shank angles, $\theta_K = \theta_T - \theta_S$, where θ_S and its derivative $\dot{\theta}_S$ form the system state proper and that θ_T and its derivative $\dot{\theta}_T$ are provided as exogenous time-series data. The new cost function and new Hamiltonian are presented in (5.23) and (5.24), respectively. Parameters θ_K^{goal} and $\dot{\theta}_K^{goal}$ represent the desired knee angle and knee angular velocity at the end of each phase of gait (see discussion in section 5.3.3). As in the previous appendix, parameters $\sigma_{\{1,2\}}$ are used to balance the kinematic goals and regeneration. The values used for $\sigma_{\{1,2\}}$ are presented in Table 5.4.

$$J = \overbrace{\frac{\sigma_1}{2} (\theta_K(T) - \theta_K^{goal})^2 + \frac{\sigma_2}{2} (\dot{\theta}_K(T) - \dot{\theta}_K^{goal})^2}^{\Phi} + \overbrace{\int_0^T Ri^2 + kn(\theta_K)\dot{\theta}_K i dt}^E \quad (5.23)$$

$$\mathcal{H} = Ri^2 + kn(\theta_K)\dot{\theta}_K i + p_1\dot{\theta}_S + p_2\ddot{\theta}_S \quad (5.24)$$

Table 5.4. Values of optimization hyperparameters for each phase of swing.

Parameter		Swing flexion	Swing extension
Terminal displacement goal coefficient	σ_1	400	100
Terminal velocity goal coefficient	σ_2	0.001	0.1

For clarity, the cost function and Hamiltonian are rewritten in terms of the state variables explicitly in (5.25) and (5.26).

$$J = \overbrace{\frac{\sigma_1}{2} (\theta_T(T) - \theta_S(T) - \theta_K^{goal})^2 + \frac{\sigma_2}{2} (\dot{\theta}_T(T) - \dot{\theta}_S(T) - \dot{\theta}_K^{goal})^2}^{\Phi} + \overbrace{\int_0^T Ri^2 + kn(\theta_T - \theta_S)(\dot{\theta}_S - \dot{\theta}_T) i dt}^E \quad (5.25)$$

$$\mathcal{H} = Ri^2 + kn(\theta_T - \theta_S)(\dot{\theta}_S - \dot{\theta}_T) i + p_1\dot{\theta}_S + p_2\ddot{\theta}_S \quad (5.26)$$

From the Hamiltonian, we obtain the co-state behavior in (5.27)-(5.28) and terminal boundary condition in (5.29) using the models in (5.1)-(5.4). The function $\delta(\cdot)$ represents the Dirac delta, which is the derivative of the $sgn(\cdot)$ used in the model for Coulomb friction in (5.1).

$$\dot{p}_1 = -\frac{\partial \mathcal{H}}{\partial \theta_S} = -ki(2\alpha_2(\theta_T - \theta_S) + \alpha_1) \left(\dot{\theta}_S - \dot{\theta}_T + \frac{p_2}{I} \right) - \frac{p_2}{I} [lm\ddot{x}_K \sin \theta_S - lm(\dot{y}_K + g) \cos \theta_S] \quad (5.27)$$

$$\dot{p}_2 = -\frac{\partial \mathcal{H}}{\partial \dot{\theta}_S} = n(\theta_T - \theta_S)ki - p_1 + \frac{p_2}{I} [b + (\tau_c + \beta|i|)\delta(\dot{\theta}_S - \dot{\theta}_T)] \quad (5.28)$$

$$\begin{cases} p_1(T) = \frac{\partial \Phi}{\partial \theta_S(T)} = -\sigma_1(\theta_T(T) - \theta_S(T) - \theta_K^{goal}) \\ p_2(T) = \frac{\partial \Phi}{\partial \dot{\theta}_S(T)} = -\sigma_2(\dot{\theta}_T(T) - \dot{\theta}_S(T) - \dot{\theta}_K^{goal}) \end{cases} \quad (5.29)$$

The optimal control law for motor current is defined by finding i^* that minimizes the Hamiltonian. The ECT friction term that depends on motor torque ($\beta|i|$) introduces some complexity here because the derivative of the absolute value function is discontinuous. Therefore, we must find i^* at every time step by calculating the Hamiltonian assuming it to be positive, negative, and zero, and choosing the option that results in the lowest value of the Hamiltonian while remaining consistent (e.g., i^* must be found to be positive if it was assumed positive). The assumed sign of i^* is represented by $\text{sgn}(i_{\pm})$ in (5.30).

$$i^* = -\frac{k}{2R} n(\theta_K) \dot{\theta}_K + \frac{p_2}{2RI} [kn(\theta_K) - \beta \text{sgn}(i_{\pm}) \text{sgn}(\dot{\theta}_K)] \quad (5.30)$$

Due to the complexity of the lower leg system relative to the simple mass-damper used to motivate the optimization approach, optimal solutions can no longer be found in closed form. The shooting method is applied to determine appropriate initial conditions for the generalized co-state variables p_1 and p_2 such that the terminal conditions in (5.29) are met. The method consists of initially guessing $p_1(0)$ and $p_2(0)$, running the optimized system to $t = T$ and determining the error in $p_{\{1,2\}}(T)$ based on the boundary condition in (5.29). The boundary condition error is then used in an iterative approach to adjust new estimates for $p_1(0)$ and $p_2(0)$. The process repeats until the sum-squared error converges to zero, within a numerical tolerance of 10^{-4} .

6 Additional Contributions

6.1 Controllers for Powered Knee Prosthesis

All versions of the Vanderbilt Powered Leg [13] consist of a powered knee joint and a powered ankle joint. Controllers have been developed for the leg prosthesis that coordinate actuation of both joints for multiple tasks, such as level walking, slope walking [95], stair ambulation [16], [96], running [97], and bilateral use [98]. In the most recent version of the leg, the two powered joints are self-contained. As a result, the powered ankle unit has been used independently as a transtibial prosthesis [99]–[101]. The powered knee joint had not, however, been applied independently.

There was interest within the CREATE group to use the powered knee as a testbed to study stumble recovery strategies in amputees. The ankle joint was removed to limit project complexity as the ankle plays the smallest role of the three lower limb joints during stumble recovery. Moreover, most transfemoral amputees use knee prostheses with no ankle articulation. A prosthesis with an articulated, powered ankle would be unfamiliar to study participants and would likely require additional training. The ability of the powered unit to deliver substantial torque (up to 85 Nm [13]) to the knee joint permits study of biomimetic recovery responses without being limited by the power capability of the device.

I led the development of a level walking controller for the powered knee, which was subsequently expanded upon by a colleague to add stumble recovery features. Later, I further developed the controller to add stair ambulation capabilities. A spline tool that I developed to generate parameterized trajectories for the swing phase of level walking was subsequently used to generate stumble recovery trajectories, swing trajectories for stair ascent, and swing trajectories in a semi-powered knee prototype ([78]).

The stair ambulation modes were developed to use the powered knee in a comparative study along with modulated-passive and semi-powered knee prostheses. Our goal was to evaluate the effect of powered stance assistance in stair ascent. Data was collected on three amputee participants. The study was limited to collecting lower-body kinematics and forces on the handrails of the stairs. The main hypothesis was that a powered device would lead to lower magnitude of forces on the handrail. We were not able to find any compelling differences between powered and unpowered devices with the collected data even though the powered prosthesis provided significant assistive torque at the knee (motor current was recorded). It is possible that more training was necessary for the participants to adjust to stance assistance since all of them have modulated passive knees as their daily-use prostheses. It is also possible that a large difference could have been observed in hip torque or metabolic cost (as shown in [16]) but those were not measured.

6.2 Understanding Stumble Recovery Strategy Selection

A classification study was performed on the biomechanical kinematic and kinetic measurements of 188 stumble events on 7 able-bodied participants (see [102]). Our goal was to determine the structure of the recovery strategy selection process. Machine learning techniques were used to find which kinematic and kinetic features had the strongest effect on the selection of each strategy. Results indicated that these biomechanical features can predict recovery strategy more accurately than swing percentage at which the perturbation occurs. Note that swing percentage is not a physical quantity and is not real-time measurable but nevertheless is typically associated to recovery strategy selection in literature. The study

serves to inform stumble detection and recovery features in orthoses and prostheses, which must respond to swing perturbations using only physical, measurable quantities.

7 References

- [1] L. G. Vailati and M. Goldfarb, "A Method for Mass Flow and Displacement Estimation in a Pneumatic Actuation System Using Valve-Based Pressure Sensing," *IEEEASME Trans. Mechatron.*, vol. 26, no. 1, pp. 235–245, Feb. 2021, doi: 10.1109/TMECH.2020.3011348.
- [2] G. Church, "Innovative Motor Designs for Electric Cars Come to Life," *Multiphysics Simulation 2017 (IEEE Spectrum)*, no. 2017, pp. 10–12, 2017.
- [3] B. Seth and W. C. Flowers, "Generalized Actuator Concept for the Study of the Efficiency of Energetic Systems," *J. Dyn. Syst. Meas. Control*, vol. 112, no. 2, pp. 233–238, Jun. 1990, doi: 10.1115/1.2896130.
- [4] R. K. Heinzmann, B. Seth, and J. Turi, "Application of a Generalized Actuator Model to the Study of Energy Regeneration Control Strategies," *J. Dyn. Syst. Meas. Control*, vol. 114, no. 3, pp. 462–467, Sep. 1992, doi: 10.1115/1.2897369.
- [5] Léo G. Vailati and Michael Goldfarb, "On Using a Brushless Motor as a Passive Torque-Controllable Brake," *J. Dyn. Syst. Meas. Control*, vol. 144, no. 9, Jun. 2022, doi: 10.1115/1.4054733.
- [6] B. Laschowski, J. McPhee, and J. Andrysek, "Lower-Limb Prostheses and Exoskeletons With Energy Regeneration: Mechatronic Design and Optimization Review," *J. Mech. Robot.*, vol. 11, no. 040801, May 2019, doi: 10.1115/1.4043460.
- [7] D. A. Winter and S. E. Sienko, "Biomechanics of below-knee amputee gait," *J. Biomech.*, vol. 21, no. 5, pp. 361–367, Jan. 1988, doi: 10.1016/0021-9290(88)90142-X.
- [8] R. Riemer and A. Shapiro, "Biomechanical energy harvesting from human motion: theory, state of the art, design guidelines, and future directions," *J. NeuroEngineering Rehabil.*, vol. 8, no. 1, p. 22, Apr. 2011, doi: 10.1186/1743-0003-8-22.
- [9] J. Thiele, B. Westebbe, M. Bellmann, and M. Kraft, "Designs and performance of microprocessor-controlled knee joints," *Biomed. Tech. Eng.*, vol. 59, no. 1, pp. 65–77, 2014, doi: 10.1515/bmt-2013-0069.
- [10] J. Thiele, C. Schöllig, M. Bellmann, and M. Kraft, "Designs and performance of three new microprocessor-controlled knee joints," *Biomed. Eng. Biomed. Tech.*, vol. 64, no. 1, pp. 119–126, 2019, doi: 10.1515/bmt-2017-0053.
- [11] B. J. Hafner and R. L. Askew, "Physical performance and self-report outcomes associated with use of passive, adaptive, and active prosthetic knees in persons with unilateral, transfemoral amputation: Randomized crossover trial," *J. Rehabil. Res. Dev.*, vol. 52, no. 6, pp. 677–700, 2015, doi: 10.1682/JRRD.2014.09.0210.
- [12] L. G. Vailati and M. Goldfarb, "On Using a Brushless DC Motor as a Passive Torque-Controllable Brake," *J. Dyn. Syst. Meas. Control*, Jun. 2022, doi: 10.1115/1.4054733.
- [13] B. E. Lawson, J. Mitchell, D. Truex, A. Shultz, E. Ledoux, and M. Goldfarb, "A Robotic Leg Prosthesis: Design, Control, and Implementation," *IEEE Robot. Autom. Mag.*, vol. 21, no. 4, pp. 70–81, Dec. 2014, doi: 10.1109/MRA.2014.2360303.
- [14] A. F. Azocar, L. M. Mooney, J.-F. Duval, A. M. Simon, L. J. Hargrove, and E. J. Rouse, "Design and clinical implementation of an open-source bionic leg," *Nat. Biomed. Eng.*, vol. 4, no. 10, pp. 941–953, Oct. 2020, doi: 10.1038/s41551-020-00619-3.
- [15] S. Mochon and T. A. McMahon, "Ballistic walking: an improved model," *Math. Biosci.*, vol. 52, no. 3, pp. 241–260, Dec. 1980, doi: 10.1016/0025-5564(80)90070-X.
- [16] E. D. Ledoux, B. E. Lawson, A. H. Shultz, H. L. Bartlett, and M. Goldfarb, "Metabolics of stair ascent with a powered transfemoral prosthesis," in *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Aug. 2015, pp. 5307–5310. doi: 10.1109/EMBC.2015.7319589.

- [17] M. J. Highsmith, J. T. Kahle, S. L. Carey, D. J. Lura, R. V. Dubey, and W. S. Quillen, "Kinetic Differences Using a Power Knee and C-Leg While Sitting Down and Standing Up: A Case Report," *JPO J. Prosthet. Orthot.*, vol. 22, no. 4, 2010, [Online]. Available: https://journals.lww.com/jpojjournal/Fulltext/2010/10000/Kinetic_Differences_Using_a_Power_Knee_and_C_Leg.7.aspx
- [18] Erik. J. Wolf, V. Q. Everding, A. A. Linberg, J. M. Czerniecki, and C. J. M. Gambel, "Comparison of the Power Knee and C-Leg during step-up and sit-to-stand tasks," *Gait Posture*, vol. 38, no. 3, pp. 397–402, Jul. 2013, doi: 10.1016/j.gaitpost.2013.01.007.
- [19] J. Andrysek and G. Chau, "An Electromechanical Swing-Phase-Controlled Prosthetic Knee Joint for Conversion of Physiological Energy to Electrical Energy: Feasibility Study," *IEEE Trans. Biomed. Eng.*, vol. 54, no. 12, pp. 2276–2283, Dec. 2007, doi: 10.1109/TBME.2007.908309.
- [20] M. R. Tucker and K. B. Fite, "Mechanical damping with electrical regeneration for a powered transfemoral prosthesis," in *2010 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, Jul. 2010, pp. 13–18. doi: 10.1109/AIM.2010.5695828.
- [21] H. Warner, P. Khalaf, H. Richter, D. Simon, E. Hardin, and A. J. van den Bogert, "Early evaluation of a powered transfemoral prosthesis with force-modulated impedance control and energy regeneration," *Med. Eng. Phys.*, vol. 100, p. 103744, Feb. 2022, doi: 10.1016/j.medengphy.2021.103744.
- [22] B. L. Hunter, "Design of a self-contained, active, regenerative computer controlled above-knee prosthesis," 1981.
- [23] K. Tabor, "The Real-Time Digital Control of a Regenerative Above-Knee Prosthesis," Michigan Technological Institute, 1984.
- [24] B. Seth, "Energy regeneration and its application to active above-knee prostheses," Thesis (Ph. D.), Massachusetts Institute of Technology, 1987. Accessed: Sep. 23, 2020. [Online]. Available: <https://dspace.mit.edu/handle/1721.1/14911?show=full>
- [25] M. Tran, L. Gabert, M. Cempini, and T. Lenzi, "A Lightweight, Efficient Fully Powered Knee Prosthesis With Actively Variable Transmission," *IEEE Robot. Autom. Lett.*, vol. 4, no. 2, pp. 1186–1193, Apr. 2019, doi: 10.1109/LRA.2019.2892204.
- [26] Minh Tran, Lukas Gabert, Sarah Hood, and Tommaso Lenzi, "A lightweight robotic leg prosthesis replicating the biomechanics of the knee, ankle, and toe joint," *Sci. Robot.*, vol. 7, no. 72, p. eabo3996, Nov. 2022, doi: 10.1126/scirobotics.abo3996.
- [27] S. C. Culver, L. G. Vailati, and M. Goldfarb, "A Primarily-Passive Knee Prosthesis with Powered Stance and Swing Assistance," in *2022 International Conference on Rehabilitation Robotics (ICORR)*, Jul. 2022, pp. 1–6. doi: 10.1109/ICORR55369.2022.9896545.
- [28] S. C. Culver, L. G. Vailati, and M. Goldfarb, "A Power-Capable Knee Prosthesis With Ballistic Swing-Phase," *IEEE Trans. Med. Robot. Bionics*, vol. 4, no. 4, pp. 1034–1045, Nov. 2022, doi: 10.1109/TMRB.2022.3216475.
- [29] P. Beater, *Pneumatic drives: System design, modelling and control*. 2007. doi: 10.1007/978-3-540-69471-7.
- [30] Bimba, "Intellisense User Manual."
- [31] P. G. Harris, G. E. O'Donnell, and T. Whelan, "Modelling and identification of industrial pneumatic drive system," *Int. J. Adv. Manuf. Technol.*, vol. 58, no. 9, pp. 1075–1086, Feb. 2012, doi: 10.1007/s00170-011-3447-7.
- [32] E. Richer and Y. Hurmuzlu, "A High Performance Pneumatic Force Actuator System: Part I—Nonlinear Mathematical Model," *J. Dyn. Syst. Meas. Control*, vol. 122, no. 3, pp. 416–425, Jun. 1999, doi: 10.1115/1.1286336.

- [33] Q. Ye and J. Chen, "Dynamic analysis of a pilot-operated two-stage solenoid valve used in pneumatic system," *Simul. Model. Pract. Theory*, vol. 17, no. 5, pp. 794–816, May 2009, doi: 10.1016/j.simpat.2009.01.005.
- [34] M. Goldfarb, "Simultaneous Force and Stiffness Control of a Pneumatic Actuator," *J. Dyn. Syst. Meas. Control*, vol. 129, no. 4, pp. 425–434, Jan. 2007, doi: 10.1115/1.2745850.
- [35] E. Richer and Y. Hurmuzlu, "A High Performance Pneumatic Force Actuator System: Part II—Nonlinear Controller Design," *J. Dyn. Syst. Meas. Control*, vol. 122, no. 3, pp. 426–434, Jun. 1999, doi: 10.1115/1.1286366.
- [36] X. Shen, J. Zhang, E. J. Barth, and M. Goldfarb, "Nonlinear Model-Based Control of Pulse Width Modulated Pneumatic Servo Systems," *J. Dyn. Syst. Meas. Control*, vol. 128, no. 3, pp. 663–669, Nov. 2005, doi: 10.1115/1.2232689.
- [37] T. Nguyen, J. Leavitt, F. Jabbari, and J. E. Bobrow, "Accurate Sliding-Mode Control of Pneumatic Systems Using Low-Cost Solenoid Valves," *IEEEASME Trans. Mechatron.*, vol. 12, no. 2, pp. 216–219, Apr. 2007, doi: 10.1109/TMECH.2007.892821.
- [38] D. Ben-Dov and S. E. Salcudean, "A force-controlled pneumatic actuator," *IEEE Trans. Robot. Autom.*, vol. 11, no. 6, pp. 906–911, Dec. 1995, doi: 10.1109/70.478438.
- [39] N. Gulati and E. J. Barth, "A Globally Stable, Load-Independent Pressure Observer for the Servo Control of Pneumatic Actuators," *IEEEASME Trans. Mechatron.*, vol. 14, no. 3, pp. 295–306, Jun. 2009, doi: 10.1109/TMECH.2008.2009222.
- [40] M. Turkseven and J. Ueda, "An Asymptotically Stable Pressure Observer Based on Load and Displacement Sensing for Pneumatic Actuators With Long Transmission Lines," *IEEEASME Trans. Mechatron.*, vol. 22, no. 2, pp. 681–692, Apr. 2017, doi: 10.1109/TMECH.2016.2582222.
- [41] F. Abry, X. Brun, M. Di Loreto, S. Sesmat, and É. Bideaux, "Piston position estimation for an electro-pneumatic actuator at standstill," *Control Eng. Pract.*, vol. 41, pp. 176–185, Aug. 2015, doi: 10.1016/j.conengprac.2015.04.016.
- [42] M. C. Potter and J. F. Foss, *Fluid Mechanics*. Great Lakes Press, inc, 1982.
- [43] Y. Kawakami, J. Akao, S. Kawai, and T. Machiyama, "Some considerations on the dynamic characteristics of pneumatic cylinders," *J. Fluid Control*, vol. 19, no. 2, pp. 22–36, Sep. 1988.
- [44] A. M. Al-Ibrahim and P. D. R. Otis, "Transient air temperature and pressure measurements during charging and discharging processes of an actuating pneumatic cylinder.," in *Proceedings of the 45th International Fluid Power Exposition and Technical Conference*, 1992.
- [45] K. Kawashima, Y. Ishii, T. Funaki, and T. Kagawa, "Determination of Flow Rate Characteristics of Pneumatic Solenoid Valves Using an Isothermal Chamber," *J. Fluids Eng.*, vol. 126, no. 2, pp. 273–279, May 2004, doi: 10.1115/1.1667888.
- [46] Tao Wang and Guangzheng Peng and Toshiharu Kagawa, "Determination of flow-rate characteristics for pneumatic components using a quasi-isothermal tank with temperature compensation," *Meas. Sci. Technol.*, vol. 21, no. 6, p. 065402, 2010.
- [47] ISO 6358-2, "Pneumatic fluid power - Determination of flow-rate characteristics of components using compressible fluids - Part2: Alternative test methods." 2013.
- [48] M. Turkseven and J. Ueda, "Model-Based Force Control of Pneumatic Actuators With Long Transmission Lines," *IEEEASME Trans. Mechatron.*, vol. 23, no. 3, pp. 1292–1302, Jun. 2018, doi: 10.1109/TMECH.2018.2832986.
- [49] G. Church, "Innovative Motor Designs for Electric Cars Come to Life," Oct. 2017.
- [50] N. Hogan, "Impedance Control: An Approach to Manipulation: Part I—Theory," *J. Dyn. Syst. Meas. Control*, vol. 107, no. 1, pp. 1–7, Mar. 1985, doi: 10.1115/1.3140702.
- [51] N. Hogan, "Impedance Control: An Approach to Manipulation: Part II—Implementation," *J. Dyn. Syst. Meas. Control*, vol. 107, no. 1, pp. 8–16, Mar. 1985, doi: 10.1115/1.3140713.

- [52] R. B. Gillespie and M. R. Cutkosky, "Stable User-Specific Haptic Rendering of the Virtual Wall," presented at the International Mechanical Engineering Congress and Exposition, ASME, Nov. 1996, pp. 397–406. [Online]. Available: <https://www.tib.eu/de/suchen/id/BLCP%3ACN018237245>
- [53] J. E. Colgate and J. M. Brown, "Factors affecting the Z-Width of a haptic display," in *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, May 1994, pp. 3205–3210 vol.4. doi: 10.1109/ROBOT.1994.351077.
- [54] D. A. Lawrence, "Impedance control stability properties in common implementations," in *Proceedings. 1988 IEEE International Conference on Robotics and Automation*, Apr. 1988, pp. 1185–1190 vol.2. doi: 10.1109/ROBOT.1988.12222.
- [55] M. I. Awad, A. A. Dehghani-Sanij, D. Moser, and S. Zahedi, "Motor electrical damping for back-drivable prosthetic knee," in *2016 11th France-Japan & 9th Europe-Asia Congress on Mechatronics (MECATRONICS) /17th International Conference on Research and Education in Mechatronics (REM)*, Jun. 2016, pp. 348–353. doi: 10.1109/MECATRONICS.2016.7547167.
- [56] K. Yuan, Q. Wang, and L. Wang, "Energy-Efficient Braking Torque Control of Robotic Transtibial Prosthesis," *IEEEASME Trans. Mechatron.*, vol. 22, no. 1, pp. 149–160, Feb. 2017, doi: 10.1109/TMECH.2016.2620166.
- [57] C. Chen, W. Chi, and M. Cheng, "Regenerative braking control for light electric vehicles," in *2011 IEEE Ninth International Conference on Power Electronics and Drive Systems*, Dec. 2011, pp. 631–636. doi: 10.1109/PEDS.2011.6147317.
- [58] X. Nian, F. Peng, and H. Zhang, "Regenerative Braking System of Electric Vehicle Driven by Brushless DC Motor," *IEEE Trans. Ind. Electron.*, vol. 61, no. 10, pp. 5798–5808, Oct. 2014, doi: 10.1109/TIE.2014.2300059.
- [59] F. Naseri, E. Farjah, and T. Ghanbari, "An Efficient Regenerative Braking System Based on Battery/Supercapacitor for Electric, Hybrid, and Plug-In Hybrid Electric Vehicles With BLDC Motor," *IEEE Trans. Veh. Technol.*, vol. 66, no. 5, pp. 3724–3738, May 2017, doi: 10.1109/TVT.2016.2611655.
- [60] A. Joseph Godfrey and V. Sankaranarayanan, "A new electric braking system with energy regeneration for a BLDC motor driven electric vehicle," *Eng. Sci. Technol. Int. J.*, vol. 21, no. 4, pp. 704–713, Aug. 2018, doi: 10.1016/j.jestch.2018.05.003.
- [61] P. R. Saradalekshmi and A. C. Binojumar, "Combined fuzzy and PI control of regenerative braking system of electric vehicle driven by brushless DC motor," in *American Institute of Physics Conference Series*, in American Institute of Physics Conference Series, vol. 2222. Apr. 2020, p. 040005. doi: 10.1063/5.0003927.
- [62] X. Zhou and J. Fang, "Precise Braking Torque Control for Attitude Control Flywheel With Small Inductance Brushless DC Motor," *IEEE Trans. Power Electron.*, vol. 28, no. 11, pp. 5380–5390, Nov. 2013, doi: 10.1109/TPEL.2013.2244617.
- [63] Y. Cao, T. Shi, X. Niu, X. Li, and C. Xia, "A Smooth Torque Control Strategy for Brushless DC Motor in Braking Operation," *IEEE Trans. Energy Convers.*, vol. 33, no. 3, pp. 1443–1452, Sep. 2018, doi: 10.1109/TEC.2018.2800750.
- [64] Y. Cao, T. Shi, Y. Yan, X. Li, and C. Xia, "Braking Torque Control Strategy for Brushless DC Motor With a Noninductive Hybrid Energy Storage Topology," *IEEE Trans. Power Electron.*, vol. 35, no. 8, pp. 8417–8428, Aug. 2020, doi: 10.1109/TPEL.2020.2964434.
- [65] F. Naseri, E. Farjah, E. Schaltz, K. Lu, and N. Tashakor, "Predictive Control of Low-Cost Three-Phase Four-Switch Inverter-Fed Drives for Brushless DC Motor Applications," *IEEE Trans. Circuits Syst. Regul. Pap.*, vol. 68, no. 3, pp. 1308–1318, Mar. 2021, doi: 10.1109/TCSI.2020.3043468.
- [66] J. Andrysek, T. Liang, and B. Steinnagel, "Evaluation of a Prosthetic Swing-Phase Controller With Electrical Power Generation," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 17, no. 4, pp. 390–396, Aug. 2009, doi: 10.1109/TNSRE.2009.2023292.

- [67] J. S. Mehling, J. E. Colgate, and M. A. Peshkin, "Increasing the impedance range of a haptic display by adding electrical damping," in *First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics Conference*, Mar. 2005, pp. 257–262. doi: 10.1109/WHC.2005.79.
- [68] M. B. Srikanth, H. Vasudevan, and M. Muniyandi, "DC Motor Damping: A Strategy to Increase Passive Stiffness of Haptic Devices," in *Haptics: Perception, Devices and Scenarios*, M. Ferre, Ed., Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 53–62.
- [69] Sanjaya Maniktala, "Understanding and Using Discontinuous Conduction Mode," in *Switching Power Supply Design & Optimization, Second Edition*, 2nd ed. New York: McGraw-Hill Education, 2014. [Online]. Available: <https://www.accessengineeringlibrary.com/content/book/9780071798143/chapter/chapter4>
- [70] R. Hamming, "Newton's Method," in *Numerical Methods for Scientists and Engineers (2nd Edition)*, Dover Publications, 1986. [Online]. Available: <https://app.knovel.com/hotlink/pdf/id:kt00B4EHG3/numerical-methods-scientists/newtons-method>
- [71] A. Visioli, Ed., "Anti-windup Strategies," in *Practical PID Control*, London: Springer London, 2006, pp. 35–60. doi: 10.1007/1-84628-586-0_3.
- [72] J. E. Colgate, "The Control of Dynamically Interacting Systems," Thesis (Ph. D.), Massachusetts Institute of Technology, Massachusetts Institute of Technology. Dept. of Mechanical Engineering., 1988. [Online]. Available: <http://dspace.mit.edu/handle/1721.1/14380>
- [73] N. Diolaiti, G. Niemeyer, F. Barbagli, and J. K. Salisbury, "Stability of Haptic Rendering: Discretization, Quantization, Time Delay, and Coulomb Effects," *IEEE Trans. Robot.*, vol. 22, no. 2, pp. 256–268, Apr. 2006, doi: 10.1109/TRO.2005.862487.
- [74] K. Kronander and A. Billard, "Stability Considerations for Variable Impedance Control," *IEEE Trans. Robot.*, vol. 32, no. 5, pp. 1298–1305, Oct. 2016, doi: 10.1109/TRO.2016.2593492.
- [75] A. Radulescu, M. Howard, D. J. Braun, and S. Vijayakumar, "Exploiting variable physical damping in rapid movement tasks," in *2012 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, Jul. 2012, pp. 141–148. doi: 10.1109/AIM.2012.6265889.
- [76] J. E. Colgate, M. C. Stanley, and J. M. Brown, "Issues in the haptic display of tool use," in *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*, Aug. 1995, pp. 140–145 vol.3. doi: 10.1109/IROS.1995.525875.
- [77] V. N. M. Arelekatti and A. G. V. Winter, "Design of a fully passive prosthetic knee mechanism for transfemoral amputees in India," in *2015 IEEE International Conference on Rehabilitation Robotics (ICORR)*, Aug. 2015, pp. 350–356. doi: 10.1109/ICORR.2015.7281224.
- [78] J. T. Lee, H. L. Bartlett, and M. Goldfarb, "Design of a Semipowered Stance-Control Swing-Assist Transfemoral Prosthesis," *IEEE/ASME Trans. Mechatron.*, vol. 25, no. 1, pp. 175–184, Feb. 2020, doi: 10.1109/TMECH.2019.2952084.
- [79] H. L. Bartlett, S. T. King, M. Goldfarb, and B. E. Lawson, "Design and Assist-As-Needed Control of a Lightly Powered Prosthetic Knee," *IEEE Trans. Med. Robot. Bionics*, vol. 4, no. 2, pp. 490–501, May 2022, doi: 10.1109/TMRB.2022.3161068.
- [80] M. Silveira, B. R. Pontes, and J. M. Balthazar, "Use of nonlinear asymmetrical shock absorber to improve comfort on passenger vehicles," *J. Sound Vib.*, vol. 333, no. 7, pp. 2114–2129, Mar. 2014, doi: 10.1016/j.jsv.2013.12.001.
- [81] A. Soliman and M. Kaldas, "Semi-active suspension systems from research to mass-market – A review," *J. Low Freq. Noise Vib. Act. Control*, vol. 40, no. 2, pp. 1005–1023, Jun. 2021, doi: 10.1177/1461348419876392.

- [82] D. V. Zenkov, A. M. Bloch, and J. E. Marsden, "The Lyapunov–Malkin theorem and stabilization of the unicycle with rider," *Syst. Control Lett.*, vol. 45, no. 4, pp. 293–302, Apr. 2002, doi: 10.1016/S0167-6911(01)00187-6.
- [83] D. A. Winter, "Energy generation and absorption at the ankle and knee during fast, natural, and slow cadences.," *Clin. Orthop.*, no. 175, pp. 147–154, May 1983.
- [84] A. Baimyshev, B. Lawson, and M. Goldfarb, "Design and Preliminary Assessment of Lightweight Swing-Assist Knee Prosthesis," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Jul. 2018, pp. 3198–3201. doi: 10.1109/EMBC.2018.8513087.
- [85] R. Fluit, E. C. Prinsen, S. Wang, and H. van der Kooij, "A Comparison of Control Strategies in Commercial and Research Knee Prostheses," *IEEE Trans. Biomed. Eng.*, vol. 67, no. 1, pp. 277–290, Jan. 2020, doi: 10.1109/TBME.2019.2912466.
- [86] M. A. Díaz *et al.*, "Human-in-the-Loop Optimization of Wearable Robotic Devices to Improve Human–Robot Interaction: A Systematic Review," *IEEE Trans. Cybern.*, pp. 1–14, 2022, doi: 10.1109/TCYB.2022.3224895.
- [87] C. G. Welker, A. S. Voloshina, V. L. Chiu, and S. H. Collins, "Shortcomings of human-in-the-loop optimization of an ankle-foot prosthesis emulator: a case series," *R. Soc. Open Sci.*, vol. 8, no. 5, p. 202020, 2021.
- [88] V. J. Harandi *et al.*, "Gait compensatory mechanisms in unilateral transfemoral amputees," *Med. Eng. Phys.*, vol. 77, pp. 95–106, Mar. 2020, doi: 10.1016/j.medengphy.2019.11.006.
- [89] S. Mochon and T. A. McMahon, "Ballistic walking," *J. Biomech.*, vol. 13, no. 1, pp. 49–57, Jan. 1980, doi: 10.1016/0021-9290(80)90007-X.
- [90] Jonathan Camargo, Aditya Ramanathan, Will Flanagan, and Aaron Young, "A comprehensive, open-source dataset of lower limb biomechanics in multiple conditions of stairs, ramps, and level-ground ambulation and transitions," *J. Biomech.*, vol. 119, p. 110320, Apr. 2021, doi: 10.1016/j.jbiomech.2021.110320.
- [91] E. C. Honert and K. E. Zelik, "Foot and shoe responsible for majority of soft tissue work in early stance of walking," *Hum. Mov. Sci.*, vol. 64, pp. 191–202, Apr. 2019, doi: 10.1016/j.humov.2019.01.008.
- [92] K. E. Zelik and P. G. Adamczyk, "A unified perspective on ankle push-off in human walking," *J. Exp. Biol.*, vol. 219, no. Pt 23, pp. 3676–3683, Dec. 2016, doi: 10.1242/jeb.140376.
- [93] F. E. Zajac, R. R. Neptune, and S. A. Kautz, "Biomechanics and muscle coordination of human walking: Part II: Lessons from dynamical simulations and clinical implications," *Gait Posture*, vol. 17, no. 1, pp. 1–17, Feb. 2003, doi: 10.1016/S0966-6362(02)00069-3.
- [94] D. Bertsekas, *Dynamic Programming and Optimal Control: Volume I*. in Athena scientific optimization and computation series. Athena Scientific, 2012.
- [95] F. Sup, H. A. Varol, and M. Goldfarb, "Upslope Walking With a Powered Knee and Ankle Prosthesis: Initial Results With an Amputee Subject," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 19, no. 1, pp. 71–78, Feb. 2011, doi: 10.1109/TNSRE.2010.2087360.
- [96] B. E. Lawson, H. A. Varol, A. Huff, E. Erdemir, and M. Goldfarb, "Control of Stair Ascent and Descent With a Powered Transfemoral Prosthesis," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 21, no. 3, pp. 466–473, May 2013, doi: 10.1109/TNSRE.2012.2225640.
- [97] A. H. Shultz, B. E. Lawson, and M. Goldfarb, "Running With a Powered Knee and Ankle Prosthesis," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 23, no. 3, pp. 403–412, May 2015, doi: 10.1109/TNSRE.2014.2336597.
- [98] B. E. Lawson, B. Ruhe, A. Shultz, and M. Goldfarb, "A Powered Prosthetic Intervention for Bilateral Transfemoral Amputees," *IEEE Trans. Biomed. Eng.*, vol. 62, no. 4, pp. 1042–1050, Apr. 2015, doi: 10.1109/TBME.2014.2334616.

- [99] A. H. Shultz, J. E. Mitchell, D. Truex, B. E. Lawson, and M. Goldfarb, "Preliminary evaluation of a walking controller for a powered ankle prosthesis," in *2013 IEEE International Conference on Robotics and Automation*, May 2013, pp. 4838–4843. doi: 10.1109/ICRA.2013.6631267.
- [100] A. H. Shultz, J. E. Mitchell, D. Truex, B. E. Lawson, E. Ledoux, and M. Goldfarb, "A walking controller for a powered ankle prosthesis," in *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Aug. 2014, pp. 6203–6206. doi: 10.1109/EMBC.2014.6945046.
- [101] S. Culver, H. Bartlett, A. Shultz, and M. Goldfarb, "A Stair Ascent and Descent Controller for a Powered Ankle Prosthesis.," *IEEE Trans. Neural Syst. Rehabil. Eng. Publ. IEEE Eng. Med. Biol. Soc.*, vol. 26, no. 5, pp. 993–1002, May 2018, doi: 10.1109/TNSRE.2018.2819508.
- [102] S. T. King, M. E. Eveld, A. Martínez, K. E. Zelik, and M. Goldfarb, "A novel system for introducing precisely-controlled, unanticipated gait perturbations for the study of stumble recovery," *J. NeuroEngineering Rehabil.*, vol. 16, no. 1, p. 69, Jun. 2019, doi: 10.1186/s12984-019-0527-7.