

International Journal of Humanoid Robotics

Vol. 12 (2014) 1550034 (34 pages)

© World Scientific Publishing Company

DOI: 10.1142/S0219843615500346



Optimization-Based Whole-Body Control of a Series Elastic Humanoid Robot

Michael A. Hopkins*, Alexander Leonessa[†] and Brian Y. Lattimer[‡]

*Terrestrial Robotics, Engineering and Controls Lab (TREC),
Mechanical Engineering, Virginia Tech
Blacksburg, VA 24061, USA*

**michael.hopkins@vt.edu*

[†]*leonessa@vt.edu*

[‡]*lattimer@vt.edu*

Dennis W. Hong

*Robotics and Mechanisms Laboratory (RoMeLa),
Mechanical and Aerospace Engineering, UCLA
Los Angeles, CA 90095, United States*

dennishong@ucla.edu

Received 15 January 2015

Accepted 3 June 2015

Published

As whole-body control approaches begin to enter the mainstream of humanoid robotics research, there is a real need to address the challenges and pitfalls encountered in hardware implementations. This paper presents an optimization-based whole-body control framework enabling compliant locomotion on THOR, a 34 degree of freedom humanoid featuring force-controllable series elastic actuators (SEAs). Given desired momentum rates of change, end-effector accelerations, and joint accelerations from a high-level locomotion controller, joint torque setpoints are computed using an efficient quadratic program (QP) formulation designed to solve the floating-base inverse dynamics (ID). Constraints on the centroidal dynamics, frictional contact forces, and joint position/torque limits ensure admissibility of the optimized joint setpoints. The control approach is supported by an electromechanical design that relies on custom linear SEAs and embedded joint controllers to accurately regulate the internal and external forces computed by the whole-body QP. Push recovery and walking tests conducted using the THOR humanoid validate the effectiveness of the proposed approach. In each case, balancing is achieved using a planning and control approach based on the time-varying divergent component of motion (DCM) implemented for the first time on hardware. We discuss practical considerations that led to the successful implementation of low-impedance whole-body control on our hardware system including the design of the robot's high-level standing and stepping behaviors and low-level joint-space controllers. The paper concludes with an application of the presented approach for a humanoid firefighting demonstration onboard a decommissioned US Navy ship.

Keywords: Humanoid locomotion; whole-body control; series elastic actuator; impedance control; divergent component of motion; quadratic program.

[†] This material is based upon work supported by (while serving at) the National Science Foundation.

M. A. Hopkins et al.

1. Introduction

Torque-controlled bipeds and quadrupeds are becoming increasingly prevalent in the robotics community as researchers continue to search for control strategies that match the speed and adaptability of low-impedance behaviors found in nature.^{1–8} This has created an interest in inverse dynamics (ID)-based whole-body control frameworks that are robust to the unexpected forces and unmodeled dynamics encountered by hardware systems operating in real environments. In the field of humanoid robotics, whole-body control strategies are often required to implement multi-objective behaviors, such as maintaining balance while aiming a fire hose or operating a drill.

Khatib *et al.*⁹ and Sentis and Khatib¹⁰ introduced pioneering whole-body control approaches utilizing null-space projection methods to compute admissible joint torques given a set of desired motion tasks and dynamic constraints. Moro *et al.*¹¹ proposed an alternative approach utilizing a linear combination of joint torques computed from a set of task-space attractors. Recently, a number of researchers have proposed convex optimization techniques to solve the ID and whole-body control problem given multi-contact constraints.^{12–20} These approaches compute joint torque setpoints that minimize tracking errors for multiple motion tasks including desired momentum rates of change, end-effector accelerations, and joint accelerations relating to whole-body motions. In general, these formulations can serve as the basis for any locomotion, manipulation, or generic multi-contact behavior.

Stephens and Atkeson¹³ were among the first to implement optimization-based ID on a torque-controlled humanoid. In order to achieve desired task-space forces for balancing and walking, the authors computed optimal contact forces using a quadratic program (QP) formulation based on the centroidal dynamics. These forces were used to obtain joint torque and acceleration setpoints via a least squares solution based on the full rigid body dynamics. Lee and Goswami¹⁴ proposed an alternative optimization-based framework to compute joint accelerations given whole-body momentum rate of change objectives. This approach was extended by Koolen *et al.*¹⁵ to simultaneously optimize contact forces using a model-based QP. Multiple acceleration-based motion constraints were included in the proposed formulation to enable task-space control, and a Newton–Euler constraint was introduced to ensure admissibility of the ground reaction forces and joint accelerations. A similar approach was implemented by Kuindersma *et al.*,²⁰ who included a quadratic cost term to stabilize the center of mass (CoM) dynamics. Feng *et al.*¹⁸ optimized accelerations, torques, and contact forces using a QP formulation that minimizes the weighted error of various motion tasks.

Several authors have proposed extensions to task-space control and ID solvers to support hierarchical prioritization of motion tasks. De Lasa and Hertzmann¹² solved a sequence of linearly constrained QPs to enforce strict prioritization of multiple task-space objectives. A similar approach was implemented by Saab *et al.*,¹⁹ who included joint torque and position limits in the optimization. Wensing and Orin¹⁷

proposed a hierarchical whole-body controller using an alternate conic optimization formulation, and Herzog *et al.*¹⁶ proposed a hierarchical ID solver using reduced floating-base dynamic constraints.

In humanoid locomotion, dynamic balancing is typically implemented through the design of feedback controllers to stabilize a robot's centroidal dynamics. In ID-based approaches, this is often achieved by controlling the robot's centroidal momentum rate of change via direct optimization of external contact forces. Several authors have implemented Cartesian PD controllers to track desired CoM trajectories using linear momentum rate of change objectives.^{14,16,17} Kuindersma *et al.*²⁰ proposed minimizing a time-varying LQR cost function for the zero moment point (ZMP) dynamics. Pratt *et al.*²¹ introduced the capture point (CP) transformation which allows the horizontal centroidal dynamics to be stabilized using a simple control law based on the centroidal moment pivot (CMP).^{22,23} This controller has been shown to be equivalent to the best CoM-ZMP regulator,²⁴ known to maximize the horizontal stability margins for CoM tracking.²⁵ Koolen *et al.*¹⁵ implemented CP tracking using a momentum-based whole-body controller, and Engelsberger *et al.*²⁶ introduced the three-dimensional (3D) divergent component of motion (DCM), which constitutes an extension of the CP transformation that simplifies locomotion planning and control on uneven terrain.

Implementing model-based whole-body control on real hardware platforms is exceptionally challenging due to numerous issues arising from communication delays, actuator friction, mechanical bias, structural deformation, and a variety of other electromechanical phenomena. The accuracy of the dynamic model and quality of the low-level torque controller are often limiting factors affecting the performance of ID-based approaches. Although some authors have demonstrated compliant balancing on purely torque-controlled hardware platforms,¹⁶ most implementations have required some form of joint-space position or velocity feedback to compensate for unmodeled dynamics and joint torque errors.^{27,28} While the inclusion of joint stiffness and damping can improve acceleration tracking, high feedback gains can result in unpredictable interaction forces due to reduced compliance. Although the presence of model error typically prevents joint stiffness and damping from being eliminated entirely, high-performance joint torque control can help reduce the reliance on joint PD feedback by enabling accurate tracking of internal and external forces commanded by the whole-body controller.

From a mechanical design perspective, series elastic actuators (SEAs) offer a number of benefits for ID-based control strategies. First proposed by Pratt *et al.* as a way to achieve improved force control of legged robots, SEAs are characterized by a spring element in series with the actuator load.²⁹ Electromechanical SEAs often feature a brushless DC motor and linear or rotary speed reducer coupled with a linear or torsional spring. The output force is typically estimated by measuring the spring deflection subject to Hooke's law. By reducing mechanical impedance, this configuration has been demonstrated to improve force fidelity and stability while increasing shock tolerance.^{29,30}

M. A. Hopkins et al.

Although most compliant whole-body control strategies have been implemented on hydraulic humanoids,^{13,16,27,28} the excellent torque tracking capabilities of electromechanical SEAs make series elastic humanoids well-suited for this type of control. A number of force and torque control strategies have been proposed for linear and rotary SEAs, typically relying on some combination of PID and feedforward control to accelerate the motor in response to the estimated spring deflection.^{31,32} For this purpose, several authors have proposed the use of an inner velocity loop based on encoder or hall effect feedback in order to overcome disturbance forces resulting from backlash and actuator stiction.^{31,32} Others have proposed the use of a disturbance observer (DOB) based on the closed-loop force controller dynamics.^{33,34}

As whole-body control approaches begin to enter the mainstream of humanoid robotics research, there is a real need to address the many challenges encountered in hardware implementations. This paper discusses the high-level, mid-level, and low-level algorithms enabling optimization-based whole-body control of THOR, a new series elastic humanoid.³⁵ Inspired by the work of Koolen *et al.* and Herzog *et al.*,^{15,16} we propose an efficient QP formulation based on the centroidal dynamics, including frictional contact constraints and joint position/torque limits. In the presented approach, the centroidal dynamics are stabilized using a novel momentum controller based on the time-varying DCM dynamics.³⁶ These components are assembled into a general framework for compliant locomotion and manipulation on uneven terrain.

The presented framework serves as the basis for THOR's standing and stepping behaviors and is validated through push recovery and compliant walking experiments. The proposed controls approach is supported by custom linear SEAs spanning the robot's hip, knee, and ankle joints. We include an overview of the robot's state estimator and discuss the embedded joint impedance controller used to track torque and acceleration setpoints computed by the whole-body QP. Compliant trajectory tracking is achieved using a cascaded joint torque controller with low gain velocity feedback. Similar to Kong *et al.* and Paine *et al.*,^{33,37} we employ an inner actuator force loop combining feedforward and PID control with a model-based DOB to achieve suitable force tracking using the THOR SEA.

This paper is an extension of previously published work by the authors.^{38,39} We include additional material related to the low-level control and state estimation implementation for the THOR humanoid as well as experimental results and analysis related to locomotion on grass and gravel. The paper concludes with a recent application of the presented control framework for a humanoid firefighting demonstration onboard the ex-USS Shadwell, a decommissioned Navy ship.

2. Humanoid Dynamics

In this section, we present a summary of humanoid dynamics and provide an overview of the time-varying DCM transformation used to stabilize the centroidal dynamics during locomotion.

2.1. Rigid body dynamics

The configuration of an articulated humanoid with n actuated degrees of freedom (DOF) can be expressed by the vector, $\mathbf{q} = [\mathbf{q}_0^T \mathbf{q}_n^T]^T \in \mathbb{R}^{6+n}$, where $\mathbf{q}_0 \in \mathbb{R}^6$ encodes the six DOF translation and orientation of the floating-base frame and $\mathbf{q}_n \in \mathbb{R}^n$ represents the n DOF vector of actuated joint positions. The full rigid-body equations of motion are given by

$$\begin{bmatrix} \mathbf{0} \\ \boldsymbol{\tau} \end{bmatrix} = \mathbf{H}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) - \sum_c \mathbf{J}_c^T \mathbf{f}_c, \quad (1)$$

where $\boldsymbol{\tau} \in \mathbb{R}^n$ is the vector of actuated joint torques, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$ is the vector of centrifugal, Coriolis and gravity torques, $\mathbf{H}(\mathbf{q})$ is the joint-space inertia matrix, and \mathbf{J}_c and $\mathbf{f}_c \in \mathbb{R}^3$ are the point Jacobians and corresponding reaction forces at each contact point, $\mathbf{r}_c \in \mathbb{R}^3$.⁴⁰ Here, we define the contact subscript, $c = 1 : N$, given N frictional contact points.

2.2. Centroidal dynamics

Alternatively, the centroidal dynamics define the reduced equations of motion for the center of mass, $\mathbf{x} = [x_{\text{com}}, y_{\text{com}}, z_{\text{com}}]^T$, and linear and angular momentum of the system, $\mathbf{l} \in \mathbb{R}^3$ and $\mathbf{k} \in \mathbb{R}^3$, as illustrated in Fig. 1. The total momentum rate of change about the CoM is given by

$$\dot{\mathbf{h}} = \begin{bmatrix} \dot{\mathbf{l}} \\ \dot{\mathbf{k}} \end{bmatrix} = \sum_c \mathbf{W}_c \mathbf{f}_c + \mathbf{w}_g, \quad (2)$$

where $\mathbf{W}_c = [\mathbf{I} \quad \mathbf{r}_{xc}^T]^T \in \mathbb{R}^{6 \times 3}$ maps contact forces to wrenches acting about the CoM and $\mathbf{w}_g = [0, 0, -mg, 0, 0, 0]^T$ encodes the force of gravity.^{14,15} Here $\mathbf{r}_{xc} \in \mathbb{R}^{3 \times 3}$




Fig. 1. Left: Time-varying DCM dynamics. Right: Centroidal dynamics with four contact points.

M. A. Hopkins et al.

is defined as the cross product matrix of the contact-CoM vector, $\mathbf{r}_c - \mathbf{x}$. The momentum rate of change is related to the joint velocities and accelerations by the equality,

$$\dot{\mathbf{h}} = \dot{\mathbf{A}}\dot{\mathbf{q}} + \mathbf{A}\ddot{\mathbf{q}}, \quad (3)$$

where $\mathbf{A} \in \mathbb{R}^{6 \times 6+n}$ represents the centroidal momentum matrix (CMM).⁴¹ Given the current configuration, \mathbf{A} can be computed using the efficient recursive algorithm proposed by Orin *et al.*⁴²

2.3. Divergent component of motion dynamics

The 3D DCM, $\boldsymbol{\xi} = \mathbf{x} + \frac{1}{\omega_0}\dot{\mathbf{x}}$, is a linear transformation of the CoM state that divides the linear centroidal dynamics into stable and unstable first-order components, where $\omega_0 = \sqrt{\frac{g}{\Delta z}}$ represents the natural frequency of the linear inverted pendulum.^{26,43} For a constant CoM height, Δz , the horizontal projection of the DCM is equivalent to the two-dimensional CP,²¹ defined as the point at which the CMP⁴⁴ must be placed at any time to allow the CoM to come to a complete rest. In a previous publication, the present authors proposed a time-varying extension of the 3D DCM to improve dynamic planning and control of the linear centroidal dynamics for locomotion on uneven terrain.³⁶ Here, we briefly summarize those results.

The time-varying DCM is defined as

$$\boldsymbol{\xi} = \mathbf{x} + \frac{1}{\omega(t)}\dot{\mathbf{x}}, \quad (4)$$

where $\omega(t) > 0$ is the time-varying natural frequency of the DCM. Solving (4) for $\dot{\mathbf{x}}$ yields the following constraint on the CoM velocity,

$$\dot{\mathbf{x}} = \omega(\boldsymbol{\xi} - \mathbf{x}). \quad (5)$$

Here, we have introduced the relaxed notation, $\omega := \omega(t)$, omitting the natural frequency's explicit dependence on time. This equality implies that the CoM will converge to the DCM with first-order dynamics. The transformed CoM dynamics (5) are asymptotically stable with respect to an equilibrium point at $\boldsymbol{\xi}$. As a result, the linear centroidal dynamics can be indirectly stabilized by regulating the DCM position using an appropriate control law.

In order to design a suitable DCM controller, we first derive similar equations of motion for the time-varying DCM in terms of the virtual repellent point (VRP),³⁶

$$\dot{\boldsymbol{\xi}} = \left(\omega - \frac{\dot{\omega}}{\omega} \right) (\boldsymbol{\xi} - \mathbf{r}_{\text{vrp}}). \quad (6)$$

First introduced by Engelsberger *et al.*,²⁶ the VRP, \mathbf{r}_{vrp} , represents the unstable equilibrium point of the DCM dynamics and can be viewed as a 3D analog to the

CMP. The time-varying extension of the VRP is defined in terms of the linear momentum rate of change and CoM position,

$$\mathbf{r}_{\text{vrp}} = \mathbf{x} - \frac{\dot{\mathbf{l}}}{m(\omega^2 - \dot{\omega})}, \quad (7)$$

where m is the total mass of the robot. Taking the VRP as a control input, the time-varying DCM dynamics (6) are controllable assuming $\omega \neq 0$ and $\omega^2 - \dot{\omega} \neq 0$.

From (2) we can define the linear momentum rate of change as $\dot{\mathbf{l}} = \sum \mathbf{f}_c - m\mathbf{g}$ where $\mathbf{g} = [0 \ 0 \ g]^T$ is the gravity vector. Thus, we can alternatively express the VRP as

$$\mathbf{r}_{\text{vrp}} = \mathbf{x} - \frac{\sum \mathbf{f}_c - m\mathbf{g}}{m(\omega^2 - \dot{\omega})} = \mathbf{r}_{\text{ecmp}} + \frac{\mathbf{g}}{\omega^2 - \dot{\omega}}, \quad (8)$$

where

$$\mathbf{r}_{\text{ecmp}} = \mathbf{x} - \frac{\sum \mathbf{f}_c}{m(\omega^2 - \dot{\omega})}, \quad (9)$$

is the time-varying extension of the enhanced centroidal moment pivot (eCMP).²⁶ The eCMP encodes the direction and magnitude of all contact forces acting on the robot and lies on the line passing through the CoM, parallel to the net contact force, $\sum \mathbf{f}_c$. When the eCMP intersects the ground plane it is equivalent to the canonical CMP.⁴⁴ In comparison, the VRP encodes the direction and magnitude of all external forces acting on the robot, including gravitational forces.

As illustrated in Fig. 1, the horizontal position of the VRP is equivalent to the eCMP, while the vertical position varies depending on the natural frequency trajectory. If the eCMP lies in the base of support, it is possible to avoid generating a horizontal moment about the CoM by placing the center of pressure (CoP) at the eCMP.^a Thus, through appropriate planning of the eCMP trajectory, the linear centroidal dynamics can be stabilized using a control law based on the VRP without generating significant angular momentum during locomotion.

3. Whole-Body Control

In this section, we provide an overview of the proposed optimization-based whole-body controller implementation. First, we describe the task-space formulation used to represent motion tasks and frictional contact constraints. Then, we discuss the specific QP formulation used to optimize joint accelerations and contact forces given multiple task-space objectives.

^aThe CoP is a 3D reference point that describes the distribution of normal forces acting on the base of support. The CoP has been shown to be equivalent to the ZMP on flat terrain.⁴⁴

M. A. Hopkins et al.

3.1. Task-space formulation

Arbitrary motion tasks such as the acceleration of a Cartesian frame or contact point can be expressed as

$$\dot{\mathbf{v}}_t = \dot{\mathbf{J}}_t \dot{\mathbf{q}} + \mathbf{J}_t \ddot{\mathbf{q}}, \quad (10)$$

where \mathbf{J}_t is the associated task-space Jacobian matrix. In this work, we model contact points using Coulomb friction constraints. Tipping and slipping of supporting surfaces is prevented by constraining the reaction forces at each contact point to lie within a corresponding friction cone, i.e., $\mathbf{f}_c \in \mathcal{C}_c$, as illustrated in Fig. 2.⁴⁵

As in a number of previous approaches,^{14,15,20,46} we employ a polyhedral approximation of the friction cone, $\mathcal{P}_c \subset \mathcal{C}_c$, at each contact point, $c = 1 : N$, allowing contact forces to be computed using a linear generating function,

$$\mathbf{f}_c = \beta_c \boldsymbol{\rho}_c. \quad (11)$$

Here, the columns of $\beta_c \in \mathbb{R}^{3 \times 4}$ are unilateral bases that span the volume of admissible forces in inertial coordinates, and $\boldsymbol{\rho}_c \in \mathbb{R}^4$ is a vector of nonnegative coefficients determined via the whole-body optimization. Additional task-space objectives are included in the whole-body controller to enforce a soft constraint on the acceleration of each contact point, $\dot{\mathbf{v}}_c = \dot{\mathbf{J}}_c \dot{\mathbf{q}} + \mathbf{J}_c \ddot{\mathbf{q}}$, assuming no-slip conditions.

3.2. Model-based optimization (quadratic program)

Given multiple motion tasks, $\dot{\mathbf{v}}_t$, the goal of whole-body control is to compute joint torques, $\boldsymbol{\tau}$, that minimize the tracking error for each task while satisfying dynamic constraints accounting for the limited control authority and range of motion of the robot. Motion tasks may include desired momentum rates of change, spatial accelerations of coordinate frames attached to individual links, or joint-space accelerations computed using an arbitrary high-level controller. Inspired by the work of Koolen *et al.* and Herzog *et al.*,^{15,16} the proposed whole-body controller optimizes desired joint accelerations, $\ddot{\mathbf{q}}$, and generalized contact forces, $\boldsymbol{\rho} = [\boldsymbol{\rho}_1^T \dots \boldsymbol{\rho}_N^T]^T$, using a linearly constrained QP in the form,

$$\min_{\ddot{\mathbf{q}}, \boldsymbol{\rho}} \|\mathbf{C}_b(\mathbf{b} - \dot{\mathbf{J}}\ddot{\mathbf{q}} - \mathbf{J}\ddot{\mathbf{q}})\|^2 + \lambda_{\ddot{\mathbf{q}}}\|\ddot{\mathbf{q}}\|^2 + \lambda_{\boldsymbol{\rho}}\|\boldsymbol{\rho}\|^2, \quad (12)$$




Fig. 2. Friction cone and friction pyramid approximations for foot contacts.

subject to

$$\mathbf{A}\dot{\mathbf{q}} + \mathbf{A}\ddot{\mathbf{q}} = \sum_c \mathbf{W}_c \mathbf{f}_c + \mathbf{w}_g, \quad (13)$$

$$\underline{\mathbf{q}} \leq \mathbf{q} + T\dot{\mathbf{q}} + \frac{1}{2}T^2\ddot{\mathbf{q}} \leq \bar{\mathbf{q}}, \quad (14)$$

$$\underline{\boldsymbol{\tau}} \leq \boldsymbol{\tau} \leq \bar{\boldsymbol{\tau}}, \quad (15)$$

$$\mathbf{0} \leq \boldsymbol{\rho}, \quad (16)$$

where \mathbf{b} is the vector of desired motion tasks and \mathbf{J} is the corresponding matrix of stacked Jacobians, i.e.,

$$\mathbf{b} = \begin{bmatrix} \dot{\mathbf{h}}_d \\ \ddot{\mathbf{q}}_d \\ \dot{\mathbf{v}}_{d_1} \\ \vdots \end{bmatrix}, \quad \mathbf{J} = \begin{bmatrix} \mathbf{A} \\ \mathbf{I} \\ \mathbf{J}_1 \\ \vdots \end{bmatrix}. \quad (17)$$

The QP cost function and constraint equations are described in the subsections below.

3.3. Optimization costs

The cost function (12) is designed to minimize the weighted quadratic error of the task-space objectives, \mathbf{b} , given a semi-positive definite weighting matrix, $\mathbf{Q}_b = \mathbf{C}_b^T \mathbf{C}_b$. Joint acceleration and contact force regularization terms are also included to ensure that the QP is strictly convex given $\lambda_{\ddot{\mathbf{q}}}, \lambda_{\boldsymbol{\rho}} > 0$. Although this approach does not allow a strict hierarchy of task prioritizations as in recent works,^{16,17} soft prioritization of motion tasks can be achieved using an appropriate weighting matrix. High weights are assigned to contact accelerations to ensure dynamic accuracy, while low weights are often assigned to angular momentum rate of change objectives to improve dynamic stability. In general, we have found that highly-weighted costs can be more forgiving than hard constraints in scenarios where one or more tasks are ill-conditioned. By permitting task-space errors, feasible solutions can typically be achieved without large accelerations or contact forces.

3.4. Optimization constraints

Constraints (13)–(16) are linear in the decision variables and ensure admissibility of the optimized joint accelerations and contact forces. Equation (13) follows from (2) and (3) enforces the centroidal dynamic constraints.¹⁵ The final inequality (16) enforces Coulomb friction constraints at each contact point. As proposed by Saab *et al.*,¹⁹ joint range of motion constraints are implemented using (14) where $\underline{\mathbf{q}}$ and $\bar{\mathbf{q}}$ represent the joint position limits and T represents a time constant determining the maximum rate of convergence to either limit. For the experiments presented in Sec. 6 we use $T \approx 0.15$ s.

M. A. Hopkins et al.

A number of QP formulations include joint torques, $\boldsymbol{\tau}$, as additional decision variables to permit explicit torque limits in the optimization.^{17–19} In these approaches, the full rigid body equations of motion (1) are included as linear equality constraints. As discussed by Herzog *et al.*,¹⁶ these additional decision variables and constraints can be eliminated, noting that the torque vector is a linear function of the joint acceleration and contact force vectors. From (1) and (11), we have

$$\begin{bmatrix} \mathbf{0} \\ \boldsymbol{\tau} \end{bmatrix} = [\mathbf{H}(\mathbf{q}) \quad - \sum_c \mathbf{J}_c^T \boldsymbol{\beta}_c \mathbf{S}_c] \begin{bmatrix} \ddot{\mathbf{q}} \\ \boldsymbol{\rho} \end{bmatrix} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}), \quad (18)$$

where \mathbf{S}_c is a selection matrix that projects $\boldsymbol{\rho}$ to $\boldsymbol{\rho}_c$. Joint torque limits are implemented by substituting the lower n equations of (18) into (15) where $\underline{\boldsymbol{\tau}}$ and $\bar{\boldsymbol{\tau}}$ represent lower and upper torque limits. Admissibility of the joint torques, joint accelerations, and contact forces is ensured by the centroidal dynamics constraint (13). This is a departure from the work of Herzog *et al.*¹⁶ and Kuindersma *et al.*²⁰ who constrain the floating-base dynamics using the upper six equations of (1); however, this is an equivalent formulation that also constrains the net momentum rate of change.

During contact switches, the optimized joint torques can exhibit discontinuities due to the addition or removal of contact constraints. In the experiments in Sec. 6, the instantaneous step is on the order of 10 Nm in the hip roll, hip pitch, and knee pitch joints (approximately 10% of the maximum torque) during transitions between single and double support phases. Despite introducing additional jerk, this control discontinuity does not appear to significantly impact the performance during walking.

3.5. QP implementation

At each time step, admissible joint accelerations, $\ddot{\mathbf{q}}_a$, and contact forces, $\boldsymbol{\rho}_a$, are computed using the proposed QP formulation (12). Corresponding joint torque setpoints, $\boldsymbol{\tau}_a$, are computed from (18). The QP is solved using an Eigen implementation of QuadProg++ which is based on an active set method for strictly convex problems.⁴⁷ Given eight active contact points and 30 actuated DOF, the optimization executes at 800 Hz on a quad-core i7 processor with OpenMP support. This enables real-time whole-body control for tasks such as standing manipulation and dynamic walking. The task-space Jacobians, CMM, and joint-space inertia matrix required by the optimization are computed using a floating-base rigid body model of the robot. The vector of centrifugal, Coriolis and gravity torques is computed using the Recursive Newton Euler algorithm.⁴⁰

4. Whole-Body Locomotion

This section provides an overview of the proposed locomotion framework developed using the presented whole-body controller. For reference, a high-level block diagram is shown in Fig. 3. Dynamic behaviors such as standing and stepping are

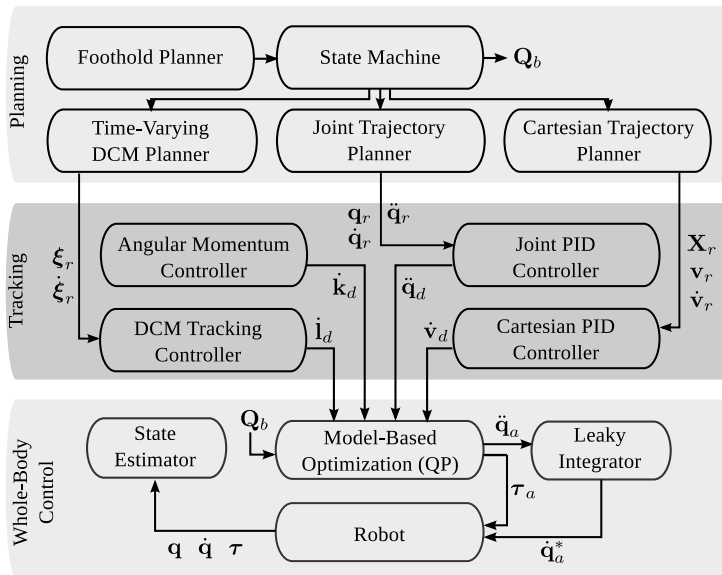


Fig. 3. Control block diagram of locomotion subsystem.

implemented using state machines that respond to external events such as toe-off and heel-strike. Following each event, state-specific parameters such as desired end-effector waypoints and step durations are passed to planning modules to generate joint-space and task-space trajectories for dynamic locomotion.

At each time step, desired motion tasks are computed using DCM, angular momentum, end-effector, and joint-space tracking controllers. Optimal joint torques are computed using the whole-body controller framework presented in Sec. 3, and compliant joint-space control is implemented to track the resulting joint trajectories. Finally, a Kalman filter-based state estimator is implemented to compute the necessary states for whole-body control. The locomotion framework executes at 150 Hz on a quad-core i7 computer. In the current unoptimized implementation, the state estimator serves as the limiting factor on the overall sample rate.

The following subsections present a detailed overview of the compliant standing and stepping behaviors implemented using the proposed framework. Note that the presented approach can be applied to a large variety of tasks including climbing, fall recovery, and whole-body manipulation.

4.1. Standing

The stand controller enables the robot to maintain balance in double or single support while tracking upper body joint trajectories for generic manipulation tasks.

Momentum control. Dynamic balancing is implemented by tracking a DCM reference trajectory, $\xi_r(t)$, that nominally lies above the center of the base of support.

M. A. Hopkins et al.

As discussed in Sec. 2.3, the time-varying DCM dynamics can be stabilized by defining an appropriate control law for the VRP. Engelsberger *et al.*²⁶ proposed a proportional tracking controller for the time-invariant DCM dynamics assuming $\omega(t) = \omega_0$, and Morisawa *et al.*²⁴ introduced integral action into the CP/DCM controller to compensate for steady state errors due to model uncertainties.

We define the following control law based on the time-varying DCM dynamics (6),

$$\mathbf{r}_{\text{vrp}} = \boldsymbol{\xi} - \frac{1}{\omega - \dot{\omega}} \left(\dot{\boldsymbol{\xi}}_r + \mathbf{k}_\xi (\boldsymbol{\xi}_r - \boldsymbol{\xi}) + \mathbf{k}_\Xi \int (\boldsymbol{\xi}_r - \boldsymbol{\xi}) dt \right) \quad (19)$$

Here, $\boldsymbol{\xi}_r$ and $\dot{\boldsymbol{\xi}}_r$ represent the reference DCM position and velocity. The first term cancels the DCM drift dynamics, and the second term implements a proportional-integral controller with unity feedforward. The nonnegative feedback gains, \mathbf{k}_ξ and \mathbf{k}_Ξ , determine the bandwidth and steady-state characteristics of the DCM controller. Note that the natural frequency of the DCM is assumed to satisfy $\omega > 0$ and $\omega^2 - \dot{\omega} > 0$. The natural frequency trajectory can be derived from a desired CoM height trajectory designed to satisfy this condition given $z_{\text{com}} - z_{\text{ecmp}} > 0$ and $\ddot{z}_{\text{com}} > -g$.³⁶ Note that the latter constraint limits the proposed approach to locomotion behaviors without a flight phase.

Given the commanded VRP setpoint from the DCM tracking controller (19), the desired linear momentum rate of change is derived from the VRP definition (7) such that

$$\dot{\mathbf{h}}_d = \begin{bmatrix} \dot{\mathbf{i}}_d \\ \dot{\mathbf{k}}_d \end{bmatrix} = \begin{bmatrix} m(\omega^2 - \dot{\omega})(\mathbf{x} - \mathbf{r}_{\text{vrp}}) \\ \mathbf{0} \end{bmatrix}. \quad (20)$$

The DCM and VRP reference trajectories are designed such that the desired CoP lies inside the base of support whenever the horizontal moment about the CoM is equal to zero. As such, the angular momentum rate of change objective is set to zero to discourage significant centroidal torque during locomotion. While this objective tends to preserve angular momentum in the system, other acceleration-based tasks included in the whole-body controller tend to dissipate angular momentum. In particular, acceleration objectives associated with the upper body and pelvis link tend to prevent excessive rotation while balancing by encouraging nonzero momentum rates of change. In the experiments presented in Sec. 6, the task weighting is tuned to achieve a suitable tradeoff between angular momentum and angular acceleration tracking.

Upper body control. Joint-space acceleration objectives are used to track upper-body joint trajectories. The acceleration objectives are computed using a standard PD controller in the form,

$$\ddot{\mathbf{q}}_d = \ddot{\mathbf{q}}_r + k_q(\mathbf{q}_r - \mathbf{q}) + b_q(\dot{\mathbf{q}}_r - \dot{\mathbf{q}}), \quad (21)$$

where \mathbf{q}_r , $\dot{\mathbf{q}}_r$, and $\ddot{\mathbf{q}}_r$ are the reference joint position, velocity, and acceleration vectors. The PD gains, $k_q \geq 0$ and $b_q \geq 0$, can be adjusted online depending on the

current task. The reference trajectories are typically specified by an external planner to allow standing manipulation, in which case high optimization weights are assigned to support accurate tracking. Low optimization weights may also be assigned to allow the arms to assist in balancing by generating angular momentum and/or shifting the CoM.

Lower body control. Pelvis and swing foot trajectories are generated using piecewise fifth order minimum jerk polynomials. The desired state of each Cartesian frame is expressed by a transform $\mathbf{X}_r = \begin{bmatrix} \mathbf{R}_r & \mathbf{r}_r \\ \mathbf{0} & 1 \end{bmatrix}$, twist $\mathbf{v}_r = [\dot{\mathbf{r}}_r^T \dot{\boldsymbol{\omega}}_r^T]^T$, and spatial acceleration $\ddot{\mathbf{v}}_r = [\ddot{\mathbf{r}}_r^T \ddot{\boldsymbol{\omega}}_r^T]^T$. The six DOF trajectories are tracked using a Cartesian PD controller in the form,

$$\ddot{\mathbf{v}}_d = \begin{bmatrix} \ddot{\mathbf{r}}_d \\ \ddot{\boldsymbol{\omega}}_d \end{bmatrix} = \begin{bmatrix} \ddot{\mathbf{r}}_r + \mathbf{K}_r(\mathbf{r}_r - \mathbf{r}) + \mathbf{B}_r(\dot{\mathbf{r}}_r - \dot{\mathbf{r}}) \\ \ddot{\boldsymbol{\omega}}_r + \mathbf{K}_R(\theta\hat{\mathbf{e}}) + \mathbf{B}_R(\boldsymbol{\omega}_r - \boldsymbol{\omega}) \end{bmatrix}, \quad (22)$$

where $\theta\hat{\mathbf{e}}$ is the axis-angle representation of the rotational error. Here \mathbf{K}_r , \mathbf{K}_R , \mathbf{B}_r , and \mathbf{B}_R are diagonal stiffness and damping matrices. Note that the pelvis linear accelerations are uncontrolled to avoid over-constraining the optimization.

4.2. Stepping

The step controller implements single-step and multi-step behaviors given desired foothold poses and step durations from a high-level footstep planner.

State machine. The step behavior employs the same subset of controllers as the stand behavior. However, an additional finite state machine is implemented to handle transitions between various contact phases as illustrated in Fig. 4. During the double support phase, both feet remain in contact with the support surface. Eight active contact constraints are added to the whole-body controller optimization corresponding to the four corners of each foot. Following the preplanned double support duration, the state machine transitions to the single support phase, breaking contact with the swing foot and configuring appropriate optimization weights for swing foot tracking. Heel-strike events are detected using a force-torque sensor in the

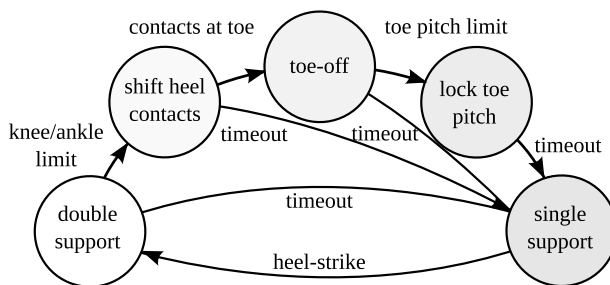


Fig. 4. Step controller state machine.

M. A. Hopkins et al.

sole of the swing foot. When the swing foot is sufficiently loaded following ground contact, the state machine transitions back to double support.

In order to increase the maximum stride length, we adopt a reactive toe-off strategy to compensate for range of motion limits. If the measured knee or ankle pitch of the swing leg reaches a soft position limit during double support, the two contact points on the heel of the swing foot are shifted to the toe in order to appropriately bound the CoP. Next the heel contacts are disabled to enable toe-off by allowing the foot to pivot about the two toe contacts. If the toe pitch reaches a critical limit, the swing foot tracking controller is configured to prevent the foot from rotating further during the double support phase.

Trajectory planning. The time-varying DCM reference trajectory is updated at the onset of each step using a discrete-time dynamic planner described in a previous publication.³⁶ First, nominal CoP and CoM height trajectories are computed from the desired foothold poses and phase durations. Next, reverse-time integration is used to plan a dynamically feasible DCM trajectory over a multi-step time horizon. The final boundary conditions ensure that the DCM lies above the base of support following the final step. Finally, model predictive control (MPC) is used to compute an admissible reference trajectory over a predefined preview window. Although not strictly necessary, the initial boundary conditions of the MPC solution eliminate any discontinuities in the nominal DCM plan.

The proposed time-varying DCM planner shares a number of similarities with MPC and LQR-based CoM planners derived from the time-varying LIPM dynamics.^{20,48} One advantage of CoM-based planners is the ability to handle flight phases during locomotion. On the other hand, a DCM-based formulation offers the ability to compute nominal reference trajectories through direct integration of the first-order dynamics. While CoM planners often compute LQR solutions over the entire multi-step preview window (relying on a jerk regularization term for stability), the proposed DCM-based approach uses a short preview window of approximately 0.5 s to drive the DCM trajectory to the nominal reference trajectory.

In comparison to analytical DCM trajectory generators, the presented planner relaxes several underlying assumptions. Whereas analytical approaches have traditionally assumed a piecewise constant CoP/CMP trajectory to obtain closed-form solutions for the reverse-time DCM dynamics,^{23,26} the proposed discrete-time planner uses a numerical integrator to compute dynamically feasible DCM trajectories given generic CoP and angular momentum trajectories. As discussed in Sec. 2.3, it is also possible to achieve generic vertical CoM trajectories within the presented framework by allowing the natural frequency of the DCM to vary over time. The benefits are most apparent when walking over uneven terrain, where careful planning of the CoP and vertical CoM trajectories can help avoid issues related to the robot's finite range of motion.

The discrete-time planner computes admissible DCM trajectories in approximately 0.002 s on a standard i7 processor given a 6.0 s time horizon, 0.5 s MPC

preview window, and 0.005 s time step. This is significantly slower than an analytical trajectory generator based on the time-invariant DCM dynamics. However, because the DCM trajectory is only replanned at the beginning of double support, the additional computation time does not have a significant effect on performance.

5. Hardware Implementation

The proposed whole-body locomotion framework is implemented on THOR, a 34 DOF series elastic humanoid developed as part of the DARPA Robotics Challenge and ONR SAFFiR programs.³⁵ The THOR humanoid is equipped with custom linear SEAs in the lower body,⁴⁹ and stiff rotary actuators in the upper body. Compliant actuation in the hip, knee, and ankle joints enables direct regulation of contact forces during locomotion, allowing the robot to adapt to uncertain terrain features. The hip yaw-roll and ankle pitch-roll joints are **parallelly** actuated using a two DOF mechanism consisting of a universal joint driven by a pair of linear actuators.⁵⁰ The hip pitch and knee pitch are serially actuated using a linear to rotary Hoeken's mechanism.⁵¹ A comprehensive overview of the THOR hardware platform is provided by Lee,³⁵ and a detailed overview of the THOR SEA is provided by Knabe *et al.*⁴⁹ Here, we provide a brief summary of the relevant actuator specifications.

Each linear SEA is equipped with a 48 V power stage and Maxon EC 4 pole motor driving a rotary to linear transmission composed of a belt reduction and precision ball screw. Universal joints at either end of the actuator constrain the ball nut rotation, and a custom titanium beam in series with the load provides the desired elasticity when subjected to moment loading. Table 1 lists the mechanical design specifications of the THOR SEA.⁴⁹ Spring force estimates are obtained using a Futek LCM-200 load cell mounted in between the spring element and the motor housing, while actuator position and velocity estimates are obtained using a pre-transmission incremental encoder mounted to the motor shaft. Joint position sensing is also implemented using a post-transmission absolute encoder mounted at each joint.

The whole-body controller implementation uses a rigid body model of the THOR humanoid created from a detailed CAD model. Despite accurate modeling of the link inertias, estimation and control errors arise from a variety of sources including sensor noise, communication delays, and structural deformation. Because the proposed ID formulation assumes a rigid body model with ideal torque sources, reliable

Table 1. THOR SEA Specifications.

Ball screw pitch	2 mm
Motor pulley ratio	3
Motor torque constant	0.0255 Nm/A
Spring stiffness	655 N/mm
Maximum velocity	0.2 m/s
Maximum force	2225 N
Continuous force	640 N

M. A. Hopkins et al.

joint-space estimation and control are critical for the accuracy of the controller. The remainder of this section discusses the low-level controller and state estimator developed for the THOR humanoid.

5.1. Low-level control

We begin by describing the low-level control approach developed to track optimized torque and acceleration setpoints obtained via the whole-body QP (12). As demonstrated by Herzog *et al.*,¹⁶ it is possible to achieve ID-based locomotion using a purely torque-controlled hardware interface. In practice, however, low-level velocity feedback is often used to improve motion tracking by introducing active damping at a high sample rate. In the THOR hardware system, upper body joint trajectories are tracked using a high-gain velocity controller, while lower body trajectories are tracked using a low-gain impedance controller.

Figure 5 shows a block diagram of the lower body joint controller. Similar to Paine *et al.* and Focchi *et al.*,^{37,52} we implement a cascaded joint impedance controller with an inner actuator force controller. The inner and outer loops execute at 2 kHz on a custom dual-axis motor controller featuring an ARM Cortex M4 processor. Each embedded device is networked on a high-speed CAN bus to enable distributed control of generic two DOF joints driven by SEAs. A detailed overview of the THOR motor controller is provided by Ressler.⁵³ In the proposed implementation, each drive receives torque and velocity setpoints from a central control computer executing the whole-body locomotion framework. The position feedback path is disabled during standing and stepping behaviors to maintain a high degree of compliance in the lower body joints.




Fig. 5. Left: Proposed dual-axis joint controller featuring an outer joint impedance loop and two inner actuator force loops. Right: Rendering of the THOR lower body featuring two DOF parallel joint mechanisms driven by linear SEAs (shaded dark gray). The outer impedance loop is designed to track joint velocity, position, and torque setpoints, $\dot{\mathbf{q}}_a^*$, \mathbf{q}_a^* , and $\boldsymbol{\tau}_a^* \in \mathbb{R}^2$, by commanding appropriate actuator forces, f_a , using an inverse statics solution based on the mechanism Jacobian, $\mathbf{J}_\ell \in \mathbb{R}^2$. The inner force loops rely on PID feedback and unity feedforward control to achieve high-performance force tracking using a DOB based on an empirical model, P_n , of the open-loop SEA plant, P .

5.1.1. Torque and velocity setpoints

A second-order notch filter is applied to the optimized joint torques, τ_a , to obtain the final torque setpoints for the low-level controller, τ_a^* . The notch filter suppresses a resonance near the high force bandwidth of the actuators causing the legs to oscillate at approximately 20 Hz when the pitch actuators approach their peak output force. The filter response is given by

$$Q_n(s) = \frac{s^2 + \omega_n^2}{s^2 + \frac{\omega_n}{Q}s + \omega_n^2}, \quad (23)$$

where the desired notch frequency and Q-factor are set to $\omega_n = 2\pi \cdot 20$ rad/s and $Q = 10$ for the experiments presented in Sec. 6.

Similar to Johnson *et al.*,²⁷ we obtain low-level joint velocity setpoints, $\dot{\mathbf{q}}_a^*$, by integrating the optimized joint accelerations, $\ddot{\mathbf{q}}_a$. The integrator is expressed by the first-order ODE,

$$\ddot{\mathbf{q}}_a^* = \alpha(\dot{\mathbf{q}} - \dot{\mathbf{q}}_a^*) + \ddot{\mathbf{q}}_a, \quad (24)$$

where $\dot{\mathbf{q}}_a^* = \int \ddot{\mathbf{q}}_a^* dt$. The leaking rate, $\alpha \geq 0$, determines the rate at which the integral drifts towards the estimated velocity. For the experiments presented in Sec. 6, we chose $\alpha = 75$ for the lower body joints and $\alpha = 0$ for the upper body joints to account for the different actuator impedances.

5.1.2. Actuator force controller

The THOR SEA is modeled as a second-order mass-spring-damper system.³⁰ This simple model characterizes the actuator dynamics in scenarios where the output impedance is high, e.g., when the end-points are fixed to ground and the actuator length remains constant. The open-loop transfer function is given by

$$\frac{F}{I} = \frac{\gamma \cdot k_S}{m_L \cdot s^2 + b_L \cdot s + k_S}, \quad (25)$$

where f is the spring force output, i is the motor current input, γ is the DC gain, k_S is the spring stiffness, b_L is the transmission damping coefficient, and m_L is the lumped sprung mass including the reflected rotor inertia and combined mass of the ball screw and motor housing.

As illustrated in Fig. 5, each actuator force loop relies on a PID controller with unity feedforward and model-based DOB to track a desired reference force, f_a . The PID controller shapes the overall response of the closed-loop plant, while the DOB attempts to reject unmodeled disturbances due to nonlinear friction forces and time-varying output impedances through feedforward control. A notch filter, Q_n , is applied to the commanded motor current to suppress a mechanical resonance near 300 Hz. The combined approach results in a force control bandwidth of approximately 30 Hz at 200 N (10% of the peak force amplitude) with excellent

M. A. Hopkins et al.

low-frequency force control. This is despite observing approximately 200 N of stiction in the transmission.

The DOB model is estimated via experimental system identification. First, the SEA is placed into a test fixture and excited by a predefined current signal to obtain the corresponding force response. Next, the frequency response function is computed using an H1 estimator, $H(s) = \frac{S_{yu}(s)}{S_{uu}(s)}$, given the cross-spectrum of the response and excitation signals, $S_{yu}(s)$, and the auto-spectrum of the excitation signal, $S_{uu}(s)$. Finally, the unknown parameters of the open-loop transfer function (25) are estimated by fitting a second-order model to the frequency domain data.

As illustrated in Fig. 5, the disturbance signal, \hat{d} , is estimated by comparing the commanded motor current with the estimated motor current obtained from an inverse model of the empirical open-loop plant, $1/P_n = H^{-1}(s)$. To ensure that the inverse plant is realizable, a second-order low-pass filter, Q_d , is applied to the disturbance estimate with a break frequency of 30 Hz. In contrast, Kong *et al.*³³ and Paine *et al.*³⁴ utilize a DOB to reject disturbances based on a nominal model of the closed-loop SEA plant. One advantage of the proposed approach is that the actuator DOB is effectively decoupled from the PID and feedforward paths, allowing the force controller gains to be modified online without updating the nominal plant.

5.1.3. Joint impedance controller

Joint torque setpoints are tracked by commanding equivalent actuator forces obtained from an inverse statics solution for the corresponding joint mechanism. Figure 6 shows the torque tracking performance for the hip pitch joint with and without the use of a DOB. In this experiment, the robot's pelvis link was mounted to a fixed platform to allow the legs to swing freely. The torque reference trajectory was computed using a fixed base recursive Newton–Euler algorithm given minimum jerk trajectories designed to emulate the motion of the swing leg during single support. The results show that the DOB feedback significantly reduces torque tracking errors due to stiction and unmodeled dynamics. We found that the corresponding improvement in torque control was critical to the success of the whole-body locomotion implementation.




Fig. 6. Hip pitch torque tracking with and without actuator DOB.

In practice, ID and task-space control approaches can be sensitive to model error, despite excellent joint torque control. Several authors have proposed cascade controller architectures featuring an outer impedance loop designed to introduce joint position and velocity feedback into the inner force or torque loop.^{31,37,52,54–57} In order to achieve suitable joint trajectory tracking, the authors experimented with two variations of a “simple” joint-space impedance controller.

The first controller introduces joint-space stiffness and damping using position and velocity estimates from the post-transmission absolute encoders at each joint. The desired actuator forces are given by

$$\mathbf{f}_a = \mathbf{J}_\ell^{-T}(\boldsymbol{\tau}_a^* + k_q(\mathbf{q}_a^* - \mathbf{q}) + b_q(\eta\dot{\mathbf{q}}_a^* - \dot{\mathbf{q}})), \quad (26)$$

where \mathbf{J}_ℓ is the mechanism Jacobian that maps angular joint velocities, $\dot{\mathbf{q}}$, to linear actuator velocities, $\dot{\boldsymbol{\ell}}$. The joint stiffness and damping coefficients, k_q and $b_q \geq 0$, determine the position and velocity feedback gains, and the scalar, $\eta \in [0, 1]$, biases the joint velocity setpoint towards zero in order to improve the stability of the derivative action. If $k_q = b_q = 0$, (26) reduces to a simple torque controller. If $\eta = 0$, the velocity feedback term introduces viscous damping.

The second controller, illustrated in Fig. 5, introduces inner-loop damping based on actuator velocity estimates from the pre-transmission incremental encoders at each motor, i.e.,

$$\mathbf{f}_a = \mathbf{J}_\ell^{-T}(\boldsymbol{\tau}_a^* + k_q(\mathbf{q}_a^* - \mathbf{q}) + b_\ell(\eta\dot{\boldsymbol{\ell}}_a^* - \dot{\boldsymbol{\ell}})). \quad (27)$$

In this scenario, joint-space damping is indirectly achieved by tracking desired linear actuator velocities computed via forward kinematics, $\dot{\boldsymbol{\ell}}_a^* = \mathbf{J}_\ell\dot{\mathbf{q}}_a^*$. Although the effective joint damping varies depending on the configuration, this approach significantly improves the stability of the derivative action, presumably due to the collocation of the motor and velocity sensor. In the experiments presented in the following section, we use $k_q = 0$ Nm/rad, $b_\ell = 10000$ N/m/s, and $\eta = 0.625$. As discussed in Sec. 5.1, joint position feedback is disabled during locomotion.

5.2. State estimation

In order to improve the whole-body controller performance, a minimal state estimator was implemented to suppress sensor noise resulting from mechanical vibration, electrical interference, and digital quantization. The estimated states include the current robot configuration, \mathbf{q} , and generalized velocities, $\dot{\mathbf{q}}$. The n DOF joint position and velocity estimates are obtained from the post-transmission absolute encoders mounted at each joint. The authors found that the use of absolute encoder derivatives for velocity estimation resulted in stable whole-body control as long as active joint damping was implemented using the proposed inner-loop velocity feedback (27). This approach was observed to reduce errors in the CoM and Cartesian velocities compared to an estimator utilizing the pre-transmission encoder measurements.

M. A. Hopkins et al.

In our system, the floating-base frame is defined at the robot's pelvis. The pelvis link is equipped with a Microstrain 3DM-GX3-25 attitude and heading reference system (AHRS) to measure the angular velocity and orientation relative to the inertial frame. Similar to Xinjilefu *et al.* and Fallon *et al.*,^{58,59} we compute linear position and velocity estimates for the floating-base frame using a combination of leg kinematics and inertial measurements. Assuming that the active contact points remain stationary during support, the estimated pelvis frame is given by

$${}^0\mathbf{X}_p = \frac{1}{N} \sum_{c=1}^N {}^0\mathbf{X}_c {}^c\mathbf{X}_p, \quad (28)$$

where ${}^c\mathbf{X}_p$ is the pelvis transform with respect to each contact point (obtained using forward kinematics), and ${}^0\mathbf{X}_c$ is the contact point transform relative to the inertial frame. The translational component of ${}^0\mathbf{X}_c$ is constant during each support phase and is only updated following heel-strike. The rotational component is defined such that the orientation of the pelvis transform, ${}^0\mathbf{X}_p$, is equal to the measured AHRS orientation in order to avoid errors due to foot rotation during the support phase.

The signal-to-noise ratio of the position measurements is much larger than the differentiated velocity measurements. As a result, we do not filter the position states in our system. To compensate for noise in the velocity estimates, we implement a kinematic Kalman filter using a process model based on the admissible joint-space and task-space acceleration inputs from the whole-body controller. At every time-step, the filtered velocity estimates are computed using SISO prediction and correction equations based on a Tustin approximation of the continuous-time dynamics,

$$\begin{aligned} v_k^- &= Av_{k-1}^+ + Bu_k, \\ v_k^+ &= v_k^- + K(y_k - (Cv_k^- + Du_k)). \end{aligned} \quad (29)$$

The state-space coefficients are given by $A = 1$, $B = \Delta T$, $C = 1$, and $D = 0.5\Delta T$. Here, y_k represents the velocity measurement, u_k represents the corresponding acceleration input obtained via the whole-body controller QP, and v_k^- and v_k^+ represent the *a priori* and *a posteriori* velocity estimates, respectively. In this formulation the Kalman gain is given by

$$K = PC^T(CPC^T + R)^{-1} = P(P + R)^{-1}, \quad (30)$$

where R is the measurement covariance and P is the steady-state *a priori* error covariance.

The omission of position states greatly simplifies the design of the estimator and enables a relatively efficient real-time implementation. In addition to filtering the floating-base and joint-space velocity estimates, we also filter task-space velocity and centroidal momentum estimates using an identical Kalman filter formulation. In each case, the task-space measurements are based on the estimated floating-base and joint-space states. Although this additional stage of filtering is redundant, it is useful in some scenarios to adjust the individual covariances for each state. For the

results presented in Sec. 6, the process and measurement covariances were selected experimentally.

6. Experimental Results

This section presents experimental results obtained using the THOR hardware platform. A comprehensive list of whole-body controller weights and feedback gains is provided in Table 2. Note that the linear momentum weights are an order of magnitude larger than the angular momentum weights. This is meant to improve dynamic stability by permitting large restoring forces when it is not possible to simultaneously satisfy both linear and angular momentum rate of change objectives. In all experiments, the DCM proportional gain is set to $k_{\xi} = 2.5$ in the horizontal axes and $k_{\xi} = 7.5$ in the vertical axis. We found that a higher stiffness in the vertical axis was necessary to achieve comparable performance due to the effect of gravity. The DCM integral gain is set to $k_{\Xi} = 0.75$ in each axis in order to compensate for steady-state errors. We also include integral feedback in the swing foot and pelvis controllers to improve tracking performance.

6.1. Balancing

Figure 7 shows THOR balancing on one foot following a external impulse of approximately 25 Ns applied to the back. The foot dimensions are 0.25×0.125 m and the CoM height is approximately 0.9 m during stance. The 60 kg robot is able to recover from the unexpected disturbance by pitching its hip and shoulder joints in order to generate the necessary reaction forces to stabilize the DCM. This is made feasible by compliant joint control in the lower body. Low damping coefficients allow the leg joints to immediately accelerate in response to the impulse without shifting the CoP to the edge of the foot.

The controller setpoints and estimates for a single push are shown in Fig. 8. The x -axis of the inertial frame is oriented to the front of the robot and the y -axis is oriented to the left. The region between the virtual toe and heel contacts is marked for reference. Note that the whole-body controller formulation ensures the optimized

Table 2. Whole-body controller weights and PID gains.

Motion Task	Units	Weight	P-Gain	I-Gain	D-Gain
\mathbf{i}	N	10	—	—	—
\mathbf{k}	Nm	1	0	0	0
$\ddot{\mathbf{q}}_{\text{arms}}$	rad/s ²	6	15	0	10
$\ddot{\mathbf{q}}_{\text{waist}}$	rad/s ²	100	40	0	20
$\ddot{\mathbf{r}}_{\text{contact points}}$	m/s ²	10000	0, 0, 0	0, 0, 0	0, 0, 0
$\ddot{\mathbf{r}}_{\text{swing foot}}$	m/s ²	1000	150, 150, 200	100, 100, 100	20, 20, 20
$\dot{\boldsymbol{\omega}}_{\text{swing foot}}$	rad/s ²	1000	100, 100, 100	100, 100, 100	10, 10, 20
$\dot{\boldsymbol{\omega}}_{\text{pelvis}}$	rad/s ²	30	70, 70, 30	50, 50, 30	30, 30, 15

M. A. Hopkins et al.




Fig. 7. THOR balancing on one foot. DCM tracking is used to stabilize the CoM dynamics following a disturbance impulse applied to the back.

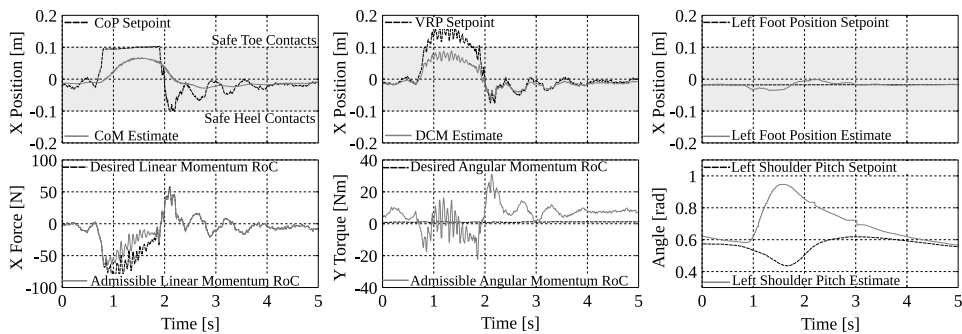


Fig. 8. Controller response following a disturbance impulse applied to the back of the robot while balancing on the right foot. The noise occurring around $t = 1$ s is related to a mechanical resonance in the lower body joint encoders. Here, the x -axis is oriented to the front of the robot, and the y -axis is oriented to the left. Admissible values correspond to the outputs of the whole-body controller optimization.

CoP setpoint does not exceed this safety region, thereby preventing the support foot from tipping. On the other hand, the VRP setpoint is allowed to leave the base of support in order to stabilize the DCM dynamics. Consequently, a net angular momentum rate of change is induced along the y -axis.

Given the relatively low weight of the arm acceleration objectives, the left shoulder pitch joint diverges significantly from the desired reference trajectory in order to generate the necessary spin angular momentum. Contrarily, the high weights of the swing foot acceleration objectives allow accurate reference tracking throughout the disturbance. Once the linear dynamics are stabilized, the pelvis and arm objectives dominate, allowing the robot to return to its original pose.

6.2. Walking

Figure 9 shows the robot walking on various surfaces including cement, gravel, and soft grass. The desired and estimated DCM and VRP trajectories are included for the

Optimization-Based Whole-Body Control

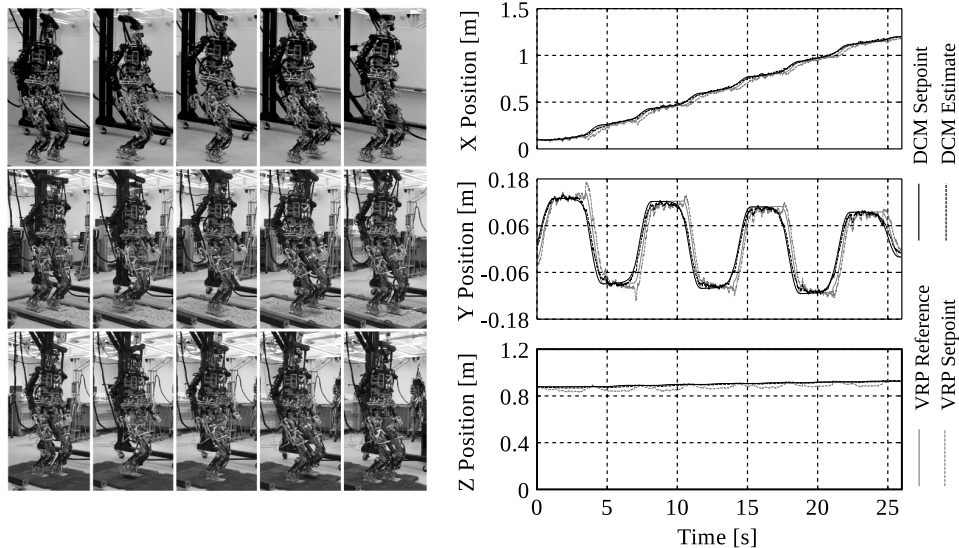


Fig. 9. Left: THOR hardware platform walking on cement, gravel, and grass. The optimization weights, control gains, and estimation parameters are identical in each test. Right: DCM and VRP trajectories while walking on cement.

cement case. Each step is characterized by a total duration of 3.5 s, stride length of 0.175 m, and swing foot apex of 0.09 m. In this experiment, the double support phase accounts for 37.5% of the step timing, resulting in a double support duration of 1.3125 s and single support duration of 2.1875 s. The optimization weights, controller gains, and estimation parameters are identical in each test (see Table 2).

Note that, although the whole-body QP formulation assumes a rigid body contact model, the step controller is able to maintain stability on compliant and unstable terrain, partially due to the low impedance of the support leg. As the results demonstrate, the robot is able to track the DCM reference trajectory using the VRP-based momentum controller defined in (19) with only minor deviations from the nominal VRP trajectory on flat terrain. The largest errors occur during heel-strike, when the DCM and VRP plans have just been updated, and the swing foot impact forces are difficult to predict.

As discussed in Sec. 5, compliant actuation in the lower body enables accurate control of the ground reaction forces during locomotion. Figure 10 compares the optimized and estimated reaction forces imparted on the left foot while taking a single step on cement. Here, the foot wrenches were estimated using a six-axis force-torque sensor mounted in between the ankle and the sole. In this scenario, the locomotion framework is able to achieve the desired foot wrench using a purely model-based control approach, i.e., without introducing a feedback path based on the force-torque data. Although significant transient errors appear in the tangential

M. A. Hopkins et al.




Fig. 10. Left foot reaction forces during a single 3.5 s step on cement. For reference, the single support phase is marked in gray. Admissible values correspond to the outputs of the whole-body controller optimization.

reaction forces, the DCM tracking controller appears robust to the corresponding variations in linear momentum rate of change.

In ID-based approaches, the performance of the low-level torque controller is often a limiting factor affecting the accuracy of the optimized contact forces. As discussed in Sec. 5.1, high-performance torque control combined with low-gain velocity feedback can offer a reasonable tradeoff between accurate force and acceleration tracking on real hardware systems. Figure 11 shows the optimized and estimated joint torques for the stance leg during stepping. Note that the actuator damping feedback contributes to the observed torque tracking errors since the whole-body controller optimization assumes ideal torque sources at each joint.

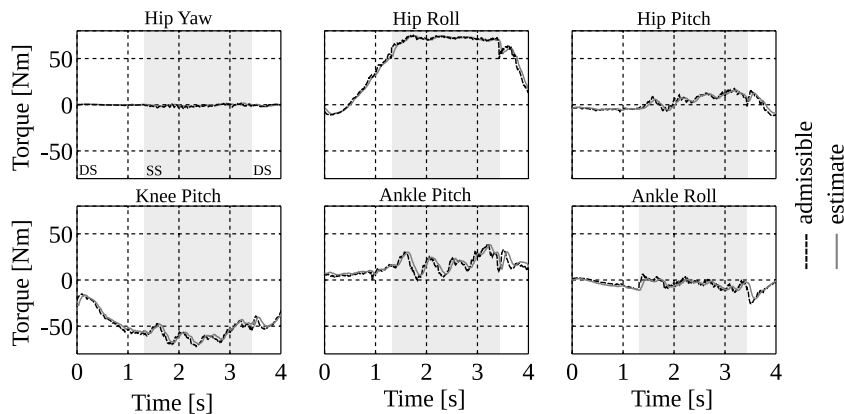


Fig. 11. Left leg joint torques while taking a single 3.5 s step on cement. Admissible values correspond to the outputs of the whole-body controller optimization. Note that the actuator damping feedback described in Sec. 5.1.3 contributes to the observed torque tracking errors.




Fig. 12. Estimated swing foot and CoM trajectories while walking on cement using the THOR hardware platform.

While precise tracking of external contact forces tends to result in improved momentum control, accurate swing foot tracking can also increase dynamic stability by ensuring proper foot placement during walking. Figure 12 shows the estimated 3D swing foot and CoM trajectories for the THOR humanoid while walking on cement. Despite significant compliance in the swing leg joints, the whole-body controller is able to track desired foothold positions to within approximately 2 cm based on the estimated pelvis trajectory and leg kinematics. The proposed DCM trajectory planner is able to adapt to deviations in the nominal foothold position and swing phase durations by modifying the nominal VRP and CoP trajectories at the beginning of double support. However, the stability of the walking controller is improved when the associated tracking errors are minimized.

Figure 13 shows the corresponding horizontal CoP trajectory during the walking experiment. The average point of pressure was estimated from the normal forces and tangential torques acting on each foot. Note that the horizontal CoP reference was defined to be equivalent to the VRP/eCMP reference shown in Fig. 9. As demonstrated in the balancing experiment, the optimized CoP is often collocated with the optimized VRP/eCMP in order to minimize horizontal angular momentum rates of change during locomotion. CoP excursions in the xy -plane are typically correlated with restoring forces generated by the momentum controller, which shift the eCMP setpoint away from the reference trajectory.

Figure 14 shows the norm horizontal CoP error for the walking experiments depicted in Fig. 9. Larger CoP errors appear in the gravel and grass data due to the

M. A. Hopkins et al.




Fig. 13. Estimated CoP trajectory while walking on cement using the THOR hardware platform.




Fig. 14. Norm horizontal CoP error while walking on cement, gravel, and grass.

effects of unmodeled surface compliance. The gravel is free to shift underfoot, while the grass deflects significantly under load. In the grass experiment, CoP errors exceeding 7.5 cm occur during the transition from single to double support. Deviations from the nominal CoP trajectory can limit control authority of the linear and angular momentum rates of change given a finite base of support. Although the stability margins decrease on compliant surfaces, the robot is able to successfully traverse the varying terrain using the same control gains and optimization weights employed on stiff surfaces.

6.3. Robustness to unmodeled terrain

The use of compliant joint control in the lower body results in behaviors that are inherently robust to uncertain and unstable terrain. Figure 15 shows the robot standing and stepping onto various terrain, initially modeled as a flat surface. In the




Fig. 15. Left: THOR standing on a balance board that rotates about a pivot. Right: Stepping onto unexpected debris and adapting to the terrain. Images used with permission of Eric Hahn.

leftmost image, the robot stands on a balance board that rotates about a pivot, and in the images to the right, the robot steps onto unexpected debris. In each scenario, the controller is able to safely adapt to uncertain terrain using low-impedance task-space feedback given identical control gains and optimization weights.

7. Shipboard Firefighting Application

The proposed whole-body locomotion framework was developed as part of the Office of Naval Research (ONR) Shipboard Autonomous **Fire Fighting** (SAFFiR) program. The goal of the ONR SAFFiR project is to develop an autonomous humanoid robot to assist in routine maintenance and damage control tasks onboard a US Navy ship. In November 2014, the THOR humanoid successfully extinguished a controlled fire onboard the ex-USS Shadwell using the presented control approach. As part of a major demonstration, the robot was required to traverse a ship corridor, grasp a standard Navy handline hose, approach a compartment containing a fire, aim the nozzle, and suppress the fire using live water.

Figure 16 shows THOR completing the fire suppression task. The robot is wearing a water resistant suit to protect the internal electronics from airborne particulates and spray emitted by the nozzle. During the demo, a human operator selected desired waypoints in order to navigate the robot through the corridor. A high-level footstep planner was used to compute appropriate foothold poses and step durations to properly position the robot. Once aligned with the compartment door, the fire was

Please change "Fir



Fig. 16. THOR suppressing a controlled fire onboard the ex-USS Shadwell in Mobile, AL. Images used with permission of Virginia Tech/Logan Wallace.

M. A. Hopkins et al.

extinguished by teleoperating the upper body joints while the whole-body controller maintained balance.

The deck was coated in a relatively low-friction epoxy that varied in slope by approximately five degrees due to warping from multiple fire tests conducted on the ship. Additionally, the fire nozzle exerted a significant wrench on the grasping hand due to the weight and pressure of the water-filled hose. For safety purposes, the hose was partially supported by an assistant firefighter (a common practice in human firefighting), resulting in an effective load of approximately 2 kg carried by the robot. The unmodeled wrench was effectively counteracted by the integral action of the pelvis and DCM controllers. Despite these challenges, the robot was able to walk to the compartment door while grasping the nozzle and maintain balance while extinguishing the fire.

The use of a task-space control formulation greatly reduced the complexity of the whole-body behaviors used to carry and aim the hose. The stand and step behaviors defined in Sec. 4 were repurposed for these tasks by assigning large whole-body controller weights to the upper body acceleration objectives to provide accurate trajectory tracking for manipulating the hose. Overall, the proposed approach adapted well to the unmodeled surface features and disturbance forces encountered during the required locomotion and manipulation tasks. We believe that the high degree of compliance achieved using low impedance SEAs greatly improved the robustness of the hardware platform under these circumstances.

8. Conclusion

This paper presented a successful implementation of optimization-based whole-body control on the THOR humanoid, enabling compliant balancing and walking. In the proposed control framework, joint torque setpoints are computed using an efficient QP formulation given a set of motion tasks and a respective weighting matrix. Frictional contact constraints and joint position/torque limits ensure admissibility of the optimized joint torques and accelerations, which are tracked using a cascaded impedance controller for joints with series elastic actuation. The introduction of low-gain joint velocity feedback using leaky integration of the optimized joint accelerations enables improved tracking of the task-space objectives while maintaining joint-space compliance.

The proposed approach was verified through push recovery and walking experiments on various terrain, culminating in the successful completion of a firefighting demo onboard a Navy ship. We are currently working to improve the speed and versatility of the walking controller to increase mobility in these types of environments. Future research efforts will focus on improved humanoid state estimation, as the accuracy of the floating base odometry is believed to be a limiting factor in the current implementation. We are also investigating improved joint-space torque control approaches using more accurate models for linear SEAs. As whole-body control approaches become increasingly robust, new opportunities are arising for

optimization-based locomotion planners that incorporate novel task-space and joint-space cost terms and constraints. We are currently interested in methods to handle uncertainty and failure during locomotion, including the application of adaptive control methods for dynamic step recovery.

Acknowledgments

This work was supported by ONR through grant N00014-11-1-0074 and by DARPA through grant N65236-12-1-1002. We would like to thank Bryce Lee, Steve Ressler, Derek Lahr, Viktor Orekhov, Coleman Knabe, Jacob Webb, Jack Newton, Joe Holler, and the many others who contributed to the development of THOR. We would also like to thank Robert Griffin, John Seminatore, Jason Ziglar, and the remainder of the TREC team who contributed to the SAFFiR demo.

References

1. J. Pratt and B. Krupp, Design of a bipedal walking robot, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 6962 (2008).
2. J. Engelsberger, A. Werner, C. Ott, B. Henze, M. A. Roa, G. Garofalo, R. Burger, A. Beyer, O. Eiberger, K. Schmid and A. Albu-Schäffer, Overview of the torque-controlled humanoid robot TORO, *14th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)* (Madrid, Spain, 2014), pp. 916–923.
3. C. Semini, N. G. Tsagarakis, E. Guglielmino, M. Focchi, F. Cannella and D. G. Caldwell, Design of HyQ — a hydraulically and electrically actuated quadruped robot, *J. Syst. Control Eng.* **225**(6) (2011) 831–849.
4. M. Hutter, C. Gehring, M. Bloesch, M. Hoepflinger, C. D. Remy and R. Siegwart, StarETH: A compliant quadrupedal robot for fast, efficient, and versatile locomotion, *15th Int. Conf. Climbing and Walking Robot (CLAWAR)* (Baltimore, USA, 2012).
5. N. Tsagarakis, S. Morfey, G. Cerda, L. Zhibin and D. Caldwell, COMpliant huMANoid COMAN: Optimal joint stiffness tuning for modal frequency control, *IEEE Int. Conf. Robotics and Automation (ICRA)* (2013), pp. 673–678.
6. M. Slovich, N. Paine, K. Kemper, A. Metger, A. Edinger, J. Weber and L. Sentis, Building HUME: A bipedal robot for human-centered hyper-agility, *Dynamic Walking Meeting* (Pensacola, USA, 2012).
7. D. Lahr, V. Orekhov, B. Lee and D. Hong, Development of a **parallely** actuated humanoid, SAFFiR, *ASME Int. Design Engineering Technical Conf.* (Portland, USA, 2013).
8. N. A. Radford, P. Strawser, K. Hambuchen, J. S. Mehling, W. K. Verdeyen, S. Donnan, J. Holley, J. Sanchez, V. Nguyen, L. Bridgwater, R. Berka, R. Ambrose, C. McQuin, J. Yamokoski, S. Hart, R. Guo, A. Parsons, B. Wightman, P. Dinh, B. Ames, C. Blakely, C. Edmonson, B. Sommers, R. Rea, C. Tobler, H. Bibby, B. Howard, L. Nui, A. Lee, M. Conover, L. Truong, D. Chesney, R. Platt Jr., G. Johnson, C.-L. Fok, N. Paine, L. Sentis, E. Cousineau, R. Sinnet, J. Lack, M. Powell, B. Morris and A. Ames, Valkyrie: NASA's first bipedal humanoid robot, *J. Field Robot.* **32**(3) (2015) 397–419.
9. O. Khatib, L. Sentis, J. Park and J. Warren, Whole-body dynamic behavior and control of human-like robots, *Int. J. Humanoid Robot.* **1**(1) (2004) 29–43.

Please change "

M. A. Hopkins et al.

10. L. Sentis and O. Khatib, A whole-body control framework for humanoids operating in human environments, *IEEE Int. Conf. Robotics and Automation (ICRA)* (Orlando, USA, 2006), pp. 2641–2648.
11. F. L. Moro, M. Gienger, A. Goswami, N. G. Tsagarakis and D. G. Caldwell, An attractor-based whole-body motion control (WBMC) system for humanoid robots, *13th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)* (Atlanta, USA, 2013).
12. M. de Lasa and A. Hertzmann, Prioritized optimization for task-space control, *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (St. Louis, USA, 2009), pp. 5755–5762.
13. B. Stephens and C. Atkeson, Dynamic balance force control for compliant humanoid robots, *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (Taipei, Taiwan, 2010), pp. 1248–1255.
14. S.-H. Lee and A. Goswami, Ground reaction force control at each foot: A momentum-based humanoid balance controller for non-level and non-stationary ground, *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (Taipei, Taiwan, 2010), pp. 3157–3162.
15. T. Koolen, J. Smith, G. Thomas, S. Bertrand, J. Carff, N. Mertins, D. Stephen, P. Abeles, J. Engelsberger, S. McCrory, J. van Egmond, M. Griffioen, M. Floyd, S. Kobus, N. Manor, S. Alsheikh, D. Duran, L. Bunch, E. Morphis, L. Colasanto, K.-L. Ho Hoang, B. Layton, P. Neuhaus, M. Johnson and J. Pratt, Summary of team IHMC’s virtual robotics challenge entry, *13th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)* (Atlanta, USA, 2013).
16. A. Herzog, L. Righetti, F. Grimminger, P. Pastor and S. Schaal, Balancing experiments on a torque-controlled humanoid with hierarchical inverse dynamics, in *Intelligent Robots and Systems (IROS)*, *IEEE/RSJ International Conference on* (Chicago, USA, 2014), pp. 981–988.
17. P. M. Wensing and D. E. Orin, Generation of dynamic humanoid behaviors through task-space control with conic optimization, *IEEE Int. Conf. Robotics and Automation (ICRA)* (Karlsruhe, Germany, 2013), pp. 3103–3109.
18. S. Feng, X. Xinjilefu, W. Huang and C. G. Atkeson, 3D walking based on online optimization, *13th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)* (Atlanta, USA, 2013).
19. L. Saab, O. Ramos, F. Keith, N. Mansard, P. Soueres and J. Fourquet, Dynamic whole-body motion generation under rigid contacts and other unilateral constraints, *IEEE Trans. Robot.* **29**(2) (2013) 346–362.
20. S. Kuindersma, F. Permenter and R. Tedrake, An efficiently solvable quadratic program for stabilizing dynamic locomotion, *IEEE Int. Conf. Robotics and Automation (ICRA)* (Hong Kong, China, 2014).
21. J. Pratt, J. Carff, S. Drakunov and A. Goswami, Capture point: A step toward humanoid push recovery, *6th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)* (Genova, Italy, 2006), pp. 200–207.
22. T. Koolen, T. De Boer, J. Rebula, A. Goswami and J. Pratt, Capturability-based analysis and control of legged locomotion, part 1: Theory and application to three simple gait models, *Int. J. Robot. Res.* **31**(9) (2012) 1094–1113.
23. J. Engelsberger, C. Ott, M. Roa, A. Albu-Schaffer and G. Hirzinger, Bipedal walking control based on Capture Point dynamics, *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (San Francisco, USA, 2011), pp. 4420–4427.
24. M. Morisawa, S. Kajita, F. Kanehiro, K. Kaneko, K. Miura and K. Yokoi, Balance control based on Capture Point error compensation for biped walking on uneven terrain, *12th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)* (Osaka, Japan, 2012), pp. 734–740.


25. T. Sugihara, Standing stabilizability and stepping maneuver in planar bipedalism based on the best COM-ZMP regulator, *IEEE Int. Conf. Robotics and Automation (ICRA)* (Kobe, Japan, 2009), pp. 1966–1971.
26. J. Engelsberger, C. Ott and A. Albu-Schaffer, Three-dimensional bipedal walking control using Divergent Component of Motion, *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (Tokyo, Japan, 2013), pp. 2600–2607.
27. M. Johnson, B. Shrewsbury, S. Bertrand, T. Wu, D. Duran, M. Floyd, P. Abeles, D. Stephen, N. Mertins, A. Lesman *et al.*, Team ihmc’s lessons learned from the darpa robotics challenge trials, *J. Field Robot.* **32**(2) (2015) 192–208.
28. S. Feng, E. Whitman, X. Xinjilefu and C. G. Atkeson, Optimization-based full body control for the Atlas robot, *14th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)* (Madrid, Spain, 2014).
29. G. Pratt and M. Williamson, Series elastic actuators, *IEEE/RSJ Int. Conf. ‘Human Robot Interaction and Cooperative Robots’ Intelligent Robots and Systems (IROS)* Vol. 1 (1995), pp. 399–406.
30. D. Robinson, J. Pratt, D. Paluska and G. Pratt, Series elastic actuator development for a biomimetic walking robot, *IEEE/ASME Int. Conf. Advanced Intelligent Mechatronics* (Atlanta, USA, 1999), pp. 561–568.
31. G. Pratt, P. Willisson, C. Bolton and A. Hofman, Late motor processing in low-impedance robots: Impedance control of series elastic actuators, *American Control Conf.*, Vol. 4 (2004), pp. 3245–3251.
32. H. Vallery, R. Ekkelenkamp, H. van der Kooij and M. Buss, Passive and accurate torque control of series elastic actuators, *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (San Diego, USA, 2007), pp. 3534–3538.
33. K. Kong, J. Bae and M. Tomizuka, Control of rotary series elastic actuator for ideal force-mode actuation in human-robot interaction applications, *IEEE/ASME Trans. Mechatron.* **14**(1) (2009) 105–118.
34. N. Paine, S. Oh and L. Sentis, Design and control considerations for high-performance series elastic actuators, *IEEE/ASME Trans. Mechatron.* **19**(3) (2014) 1080–1091.
35. B. Lee, Design of a humanoid robot for disaster response, Master’s thesis, Virginia Polytechnic Institute and State University (2014).
36. M. A. Hopkins, D. W. Hong and A. Leonessa, Humanoid locomotion on uneven terrain using the time-varying Divergent Component of Motion, *14th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)* (Madrid, Spain, 2014).
37. N. Paine, J. S. Mehling, J. Holley, N. A. Radford, G. Johnson, C.-L. Fok and L. Sentis, Actuator control for the NASA-JSC Valkyrie humanoid robot: A decoupled dynamics approach for torque control of series elastic robots, *J. Field Robot.* **32**(3) (2015) 378–396.
38. M. A. Hopkins, D. W. Hong and A. Leonessa, Compliant locomotion using whole-body control and Divergent Component of Motion tracking, *IEEE Int. Conf. Robotics and Automation (ICRA)* (Seattle, USA, 2015).
39. M. A. Hopkins, S. A. Ressler, D. F. Lahr, D. W. Hong and A. Leonessa, Embedded joint-space control of a series elastic humanoid, *IEEE/RSJ Intelligent Conf. Intelligent Robots and Systems (IROS)* (Hamburg, Germany, 2015).
40. R. Featherstone, *Rigid Body Dynamics Algorithms*, Vol. 49 (Springer Berlin, 2008).
41. D. Orin and A. Goswami, Centroidal momentum matrix of a humanoid robot: Structure and properties, *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (Nice, France, 2008), pp. 653–659.
42. D. E. Orin, A. Goswami and S.-H. Lee, Centroidal dynamics of a humanoid robot, *Auton. Robots* **35**(23) (2013) 161–176.

M. A. Hopkins et al.

43. J. Engelsberger, T. Koolen, S. Bertrand, J. Pratt, C. Ott and A. Albu-Schaffer, Trajectory generation for continuous leg forces during double support and heel-to-toe shift based on divergent component of motion, *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (Chicago, USA, 2014).
44. M. B. Popovic and H. Herr, Ground reference points in legged locomotion: Definitions, biological trajectories and control implications, *Int. J. Robot. Res.* **24**(12) (2005) 1013–1032.
45. Y. Abe, M. da Silva and J. Popovic, Multiobjective control with frictional contacts, *Symp. Computer Animation* (San Diego, USA, 2007), pp. 249–258.
46. C. Ott, M. Roa and G. Hirzinger, Posture and balance control for biped robots based on contact force optimization, *11th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)* (Bled, Slovenia, 2011), pp. 26–33.
47. D. Goldfarb and A. Idnani, A numerically stable dual method for solving strictly convex quadratic programs, *Math. Program.* **27**(1) (1983) 1–33.
48. J. Park and Y. Youm, General ZMP preview control for bipedal walking, *2007 IEEE Int. Conf. Robotics and Automation* (Rome, Italy, 2007), pp. 2682–2687.
49. C. Knabe, B. Lee, V. Orekhov and D. Hong, Design of a compact, lightweight, electro-mechanical linear series elastic actuator, *ASME International Design Engineering Technical Conf.* (Buffalo, USA, 2014).
50. B. Lee, C. Knabe, V. Orekhov and D. Hong, Design of a human-like range of motion hip joint for humanoid robots, *ASME International Design Engineering Technical Conf.* (Buffalo, USA, 2014).
51. C. Knabe, B. Lee and D. Hong, An inverted straight line mechanism for augmenting joint range of motion in a humanoid robot, *ASME International Design Engineering Technical Conf.* (Buffalo, USA, 2014).
52. M. Focchi, G. A. Medrano-Cerda, T. Boaventura, M. Frigerio, C. Semini, J. Buchli and D. G. Caldwell, Robot impedance control and passivity analysis with inner torque and velocity feedback loops (2014), arXiv:1406.4047.
53. S. A. Ressler, Design and implementation of a dual axis motor controller for parallel and serial series elastic actuators, Master’s thesis, Virginia Polytechnic Institute and State University (2014).
54. D. Bentivegna, C. Atkeson and J.-Y. Kim, Compliant control of a hydraulic humanoid joint, *7th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)* (Pittsburg, USA, 2007), pp. 483–489.
55. H. Vallery, J. Veneman, E. van Asseldonk, R. Ekkelenkamp, M. Buss and H. van Der Kooij, Compliant actuation of rehabilitation robots, *IEEE Robot. Autom. Maga.* **15**(3) (2008) 60–69.
56. M. Hutter, C. D. Remy, M. A. Hoepflinger and R. Siegwart, High compliant series elastic actuation for the robotic leg ScarLETH, *Proc. Int. Conf. Climbing and Walking Robots (CLAWAR)* (Paris, France, 2011).
57. N. L. Tagliamonte and D. Accoto, Passivity constraints for the impedance control of series elastic actuators, *Proc. Inst. Mech. Eng. J. Syst. Control Eng.* **228**(3) (2014) 138–153.
58. X. Xinjilefu, S. Feng, W. Huang and C. Atkeson, Decoupled state estimation for humanoids using full-body dynamics, *IEEE Int. Conf. Robotics and Automation (ICRA)* (Hong Kong, 2014), pp. 195–201.
59. M. F. Fallon, M. Antone, N. Roy and S. Teller, Drift-free humanoid state estimation fusing kinematic, inertial and LIDAR sensing, *14th IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)* (Madrid, Spain, 2014), pp. 112–119.



Michael Hopkins received his Ph.D. in Computer Engineering from Virginia Tech in 2014. His dissertation focused on Dynamic Locomotion and Whole-Body Control for Compliant Humanoids. He recently contributed to Virginia Tech's DARPA Robotics Challenge team as a research assistant in the Department of Mechanical Engineering. His research interests include legged locomotion, dynamic planning, biologically-inspired control, optimization-based inverse dynamics, and force and impedance control strategies for series elastic actuators.




Alexander Leonessa obtained a Doctoral degree in Aerospace Engineering at Georgia Tech in December 1999. His research focused on nonlinear robust control techniques for general nonlinear systems. His appointment as a faculty member at Virginia Tech started in December 2007, after two previous similar appointments at Florida Atlantic University and the University of Central Florida. His research and contribution include (i) control theory with application to autonomous vehicles guidance and navigation, (ii) nonlinear system identification with application to health monitoring, (iii) real-time embedded control with application to system design of robotic systems, and (iv) functional electrical stimulation of muscles for rehabilitation of stroke survivors and patients with spinal cord injuries. In particular, the dominant idea in his research effort is that most real-world physical systems are too complex to accurately model, hence model uncertainties must be accounted for in the control system design process using some kind of self-learning procedure. Dr. Leonessa has been involved in these areas of research for more than 15 years during which he has published more than 60 papers (all peer reviewed).



Brian Lattimer is a Professor in the Department of Mechanical Engineering at Virginia Tech where he conducts research on disaster response and safety related technology. In the area of robotics, his research interests include developing and fielding humanoid and wheeled robots for disaster response, vision systems for low visibility conditions, fire feature classification, autonomous suppression of fires, and environmental protection for robots. He was the Principal Investigator for the U.S. Navy Shipboard Autonomous Firefighting Robot (SAFFiR) project that developed a humanoid robot to assist sailors as well as a co-principal investigator for one of the DARPA Robotics Challenge Finals teams. He obtained his PhD in 1996 from Virginia Tech in Mechanical Engineering.

M. A. Hopkins et al.



Dennis Hong is a Professor and the Founding Director of RoMeLa (Robotics & Mechanisms Laboratory) of the Mechanical & Aerospace Engineering Department at UCLA. His research focuses on Humanoids, Robot Locomotion, Autonomous Vehicles, Platform Design, Kinematics & Mechanisms. Dr. Hong has been honored with numerous awards including the NSF CAREER award, the SAE Teetor Educational Award, the ASME Freudenstein/GM Young Investigator Award to name a few. Dr. Hong also actively leads student teams for various robotics competitions winning top prizes including the DARPA Urban Challenge (Third place), the RoboCup (Humanoid division, 4-time world champion, Louis Vuitton Cup) and the DARPA Robotics Challenge (Finalist). Washington Post magazine called Dr. Hong “the Leonardo da Vinci of robots.”