Exploiting Variable Physical Damping in Rapid Movement Tasks

Andreea Radulescu, Matthew Howard, David J. Braun and Sethu Vijayakumar

Abstract— Until now, design of variable physical impedance actuators (VIAs) has been limited mainly to realising *variable stiffness* while other components of impedance shaping, such as damping, are either fixed (e.g., with the addition of fixed passive dampers) or modulated with active feedback control schemes. In this work we introduce an actuator that is capable of *simultaneous and independent* physical damping and stiffness modulation. Using optimal control techniques, we explore how variable physical damping can be exploited in such an actuator in the context of rapid movement. Several numerical simulation results are presented, in addition to an experiment realised on variable impedance robotic hardware.

I. INTRODUCTION

In recent years, the robotics community has looked for inspiration in the capabilities of humans to design a generation of actuators that have variable impedance [1], [2], [3]. The benefits of such actuators include high dynamic range (e.g. due to the ability to store energy in spring-like actuators) [4] and a stable and fast response (since compliance is built into the actuator mechanically, sensory feedback is not required to respond to perturbations).

However, a common deficiency of such actuators is that impedance modulation is restricted to altering the physical stiffness alone, while other impedance terms (such as the damping) are fixed. Currently, many passive variable stiffness actuators still rely on active control schemes to realise damping requirements [5], and there are very few actuator designs that are able to passively vary damping [6]. For those that do, it remains an open question as to how to exploit this added degree of actuation redundancy.

Research in the control of variable physical stiffness has shown optimal control to be an effective methodology for designing control strategies that can exploit variable passive compliance. For example, optimal control approaches have been shown to be highly effective in exploiting the elastic properties of variable stiffness actuators in explosive tasks such as throwing and hitting [4], [7] as well as in periodic tasks [8], [9], [10]. In addition, stochastic optimal control with model adaptation has also been exploited in order to cope with model uncertainty and perturbations [11], [12], [13]. This is also consistent with observations regarding the way in which humans adapt their impedance under similar conditions [14].

In this paper we investigate how variable physical damping can be (i) implemented in robotic hardware and (ii) exploited according to the principles of optimal control. We propose a mechanism based on using electrical inductance to cause a variable mechanical damping effect, derive a minimalistic

A. Radulescu, M. Howard, D. J. Braun and S. Vijayakumar are with the Institute of Perception Action and Behaviour, University of Edinburgh, Scotland, UK. E-mail: `A.Radulescu@sms.ed.ac.uk`

model to identify this damping effect, and illustrate how it can be integrated into an established variable stiffness device to achieve *simultaneous and independent* stiffness and damping modulation. Having validated our model of the damping mechanism, we then show how variable damping can be exploited in the context of rapid movement tasks. We present simulation results that illustrate that variable damping can outperform fixed damping schemes in the context of fast reaching to a target. Finally, we illustrate the effectiveness of our approach with an experiment realised on an implementation of our passive variable impedance actuator in hardware.

II. MOTIVATION AND RELATED WORK

There is currently an abundance of designs for mechanisms that can vary *stiffness* (for a review, see [15]), but relatively few explicitly consider damping requirements for efficient operation [5], [6]. In this section we discuss the issue of damping in the context of variable physical impedance actuation.

A. Damping in Variable Stiffness Actuators

To investigate the effects of damping modulation in the above mentioned framework, we selected the simplest system available - a one-link joint system. The equation of motion for such a system is:

$$J\ddot{q} = \tau + \tau_f, \tag{1}$$

where J represents the inertia, $\ddot{q} \in \mathbb{R}$ the joint acceleration, $\tau \in \mathbb{R}$ is the torque provided by the actuators (defined as $\tau = \tau_k + \tau_b$ with τ_k and τ_b the stiffness and damping torques, respectively) and $\tau_f \in \mathbb{R}$ is the torque due to friction.

Intuitively, an appropriate selection of the damping is crucial for responsive, accurate control of a VIA. If the system is too lightly damped oscillations arise when rapid movements are attempted. In contrast, if the damping is too heavy its response will be sluggish, in which case many of the positive effects of (variable) passive stiffness, such as energy storage, may not be exploited. One may also say that, in general, it is natural to consider adding *variable damping* to a system with variable stiffness in order to achieve an appropriate response (e.g., under-, over-, or critical damping) as the stiffness varies. In the following, we explore these issues with respect to a simple example of a variable stiffness actuator (VSA).
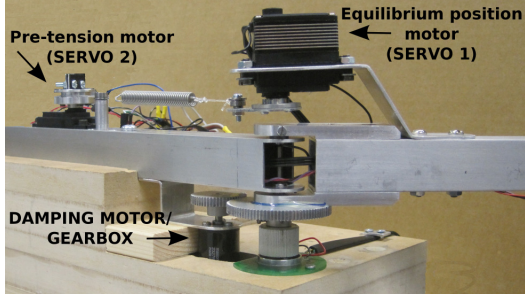
B. Example: Mechanically Adjustable Compliance and Controllable Equilibrium Position Actuator (MACCEPA)

As an example of a VSA that is easy to build and intuitive to operate we apply our variable damping scheme to the
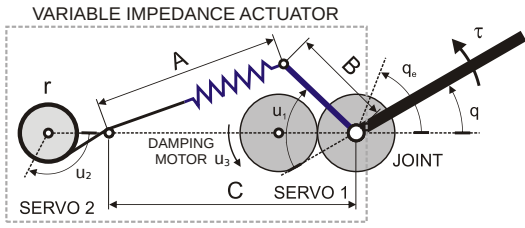
(a) Model.

(b) Robot.

Fig. 1. Mechanically Adjustable Compliance and Controllable Equilibrium Position Actuator (MACCEPA) [16] with Variable Damping. Parameters: r=0.01m, B=0.03m, C=0.13m, link mass: 0.125kg, link length: 0.295m, centre of mass location on the link: 0.1475m

Mechanically Adjustable Compliance and Controllable Equilibrium Position Actuator (MACCEPA) [16]. An illustration of the device is provided in Fig. 1 and a brief description of its operation follows.

The MACCEPA is designed to give simultaneous equilibrium position and joint stiffness control through use of two independently controlled servomotors. The first servomotor (attached to the free-moving link) adjusts the position of a lever of length B, while the second (attached to the base) adjusts the position of a winding drum of radius r attached via a cord to the spring (see Fig. 1(a)).

Torque around the joint is produced when a discrepancy is introduced between the first servomotor angle u_1 and the joint angle q (i.e., $\alpha := u_1 - q \neq 0$) causing the spring to be extended away from its rest length $L_0 = C - B$. The restorative spring force is $F = -\kappa(L - L_0)$ where κ is the spring constant and L is the resultant length of the spring. The latter is given by the length $A = \sqrt{B^2 + C^2 - 2BC \cos \alpha}$ (see Fig. 1(a)) plus any additional extension due to the pre-tensioning servo u_2, i.e., $L = A + ru_2$.

Considering the moment arm, the torque around the joint is given by

$$\tau_k(q, u_1, u_2) = \kappa BC \sin \alpha \left(1 + \frac{ru_2 - (C - B)}{\sqrt{B^2 + C^2 - 2BC \cos \alpha}}\right) \tag{2}$$

and the stiffness ($k := -\partial \tau / \partial q$) is

$$k(q, u_1, u_2) = \kappa BC \cos \alpha \left(1 + \frac{r\, u_2 - (C - B)}{\sqrt{B^2 + C^2 - 2BC \cos \alpha}}\right)$$
$$- \frac{\kappa B^2 C^2 \sin^2 \alpha \,(ru_2 - (C - B))}{(B^2 + C^2 - 2BC \cos \alpha)^{\frac{3}{2}}}. \tag{3}$$

(For full details of the derivation of the torque relationship (2) we refer the reader to [16].) Note that, as with many

VSA designs, the joint torque and stiffness have a *non-linear* dependence on the joint and motor configurations.

In addition to this existing design, we add a third control variable u_3, corresponding to damping modulation. Thus for the combined system (MACCEPA with Variable Damping, MACCEPA-VD) the torque relationship corresponding to τ from (1) is:

$$\tau(q, \dot{q}, u_1, u_2, u_3) = \tau_k(q, u_1, u_2) + \tau_b(\dot{q}, u_3), \tag{4}$$

where $\tau_k(q, u_1, u_2)$ is given by (2). The details of the mechanical design used are presented in Sec. III. For details of the system identification and estimation of the damping please see the Appendix.

C. Selecting the System Response

For the efficient control and operation of a VSA such as the MACCEPA, it is necessary to select an appropriate system response dependant on the task. For example, in periodic tasks, such as running, one may wish to have an *underdamped* response, encouraging oscillation through the natural dynamics. In tasks such as tracking, on the other hand, one may require something akin to *critical damping*, to ensure the target reference is reached as rapidly as possible without overshoot.

For one-link *linear-time-invariant* (LTI) systems this is relatively straight-forward since the damping response can be characterised through the *damping ratio*

$$\xi = b/2\sqrt{kJ}, \tag{5}$$

where J represents the inertia, b the damping and k the stiffness. In an LTI system selecting, for example, a critically damped response simply requires choosing b such that $\xi = 1$.

Systems with variable stiffness (e.g. $k = k(q, u_1, u_2)$), however, are not LTI by definition. Selection of the damping response for such systems is non-trivial, and will in general require modulation of b in a non-trivial way. This is evident in the case of the MACCEPA, where there is a complex, non-linear dependence of the stiffness (3) on the joint and motor configurations. Furthermore, since such systems are not LTI, the damping response is not well-characterised through ξ as defined by (5) and so, for example, ensuring $\xi = 1$ (e.g., by coupling the stiffness and damping through (5)) may not result in the fastest non-oscillatory response (i.e., critical damping). Indeed, this expectation is confirmed in our experiments (ref. Sec. V).

D. Related Work

Active, passive and semi-active damping are all viable ways to modulate the dissipative properties of a dynamical system. Among the three, *active damping*, realised through velocity feedback, is the most frequently used in practice. While implementation of this method can be done solely through software control, its stable dissipative operation is limited by the bandwidth of the feedback loop. Indeed, due to noise and phase lag issues, the non-ideal velocity feedback may lead to instability. In addition to this effect active damping control may lead to a considerable energy cost.

As an alternative purely passive damping can be realised by dissipative mechanical elements (i.e., dampers, brakes) that are guaranteed to be passive regardless of the excitation. Such elements have no bandwidth limitation however, their dissipative properties cannot be modulated during task execution. The compromise between these two methods is provided by semi-active damping. Conceptually semi-active damping is *active modulation of the passive damping*.

In practice semi-active damping is mainly realised through magnetorheological (MR) dampers and frictional dampers (FD). MR dampers are filled with a suspension of iron particles and oil as a base fluid (i.e., MR fluid). The damping properties of such devices can be directly controlled with a magnetic field, that solidifies the MR fluid [17], [18]. In the context of robotic applications, MR dampers require additional space and introduce fluid in the system, which is often not desirable. A different mechanism is employed by frictional dampers that use Coulomb friction between two moving bodies. In this case modulation of the frictional force is achieved with control of the normal contact force. A popular way to realise the latter is through piezo-electric actuators [6], [18]. These semi-active damping methods require additional elements in mechanical design, but have considerably higher bandwidth and spend less energy compared to active damping.

In this paper we consider an alternative design that exploits a DC motor damping effect. The basic concept of this approach is to attach a *back-drivable motor/gearbox* unit to the joint (see Fig. 1) that operates as a generator during the motion. In this way one converts mechanical energy to the electrical domain where it is dissipated through a resistor. In previous work the same effect has been used in the context of designing adjustable dampers in haptic interfaces [19], [20]. To the authors knowledge, the present paper is the first to combine variable physical stiffness and damping in a single actuator with this approach. In the following, we outline the details of the variable damping mechanism in our actuator.

III. MECHANISM DESIGN

A. Variable Passive Electrical Damping

When the output terminals of electric motors are connected together (i.e., shorted) there is a motor braking effect whereby rotations of the output shaft due to an externally applied torque are opposed by a torque associated with the current induced within the motor. In the present paper, we utilise this electrical effect to control the *passive mechanical damping* in our variable impedance system.

The electronics and equivalent circuit representing our variable passive damper is depicted in Fig. 2. According to this circuit, applying a voltage V will cause current I according to

$$L\dot{I} + R_e I + V_{emf} = V, \qquad (6)$$

where L is the inductance of the circuit, R_e the effective resistance and V_{emf} is the back electromotive force (e.m.f.) generated when the motor armature rotates. The relations between (i) the angular velocity of the output shaft $\dot{\theta}$ and
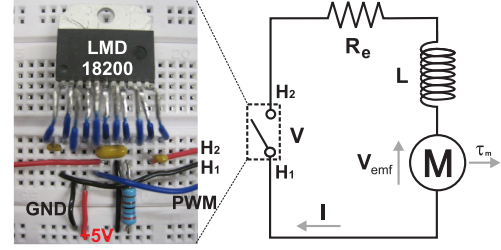

Fig. 2.   H-bridge damping: circuit (left) and equivalent circuit (right).

V_{emf} and (ii) the motor torque τ_m (reflected through the gear reduction) and the current are given by

$$V_{emf} = n\kappa_{\dot{q}}\dot{\theta}, \qquad \tau_m = n\kappa_{\tau}I, \qquad (7)$$

where $\kappa_{\dot{q}}$ and κ_{τ} are the speed and torque constants of the motor and $n : 1$ defines the gear reduction.

We now look to establish the relation between the motor torque and the operation of the electrical circuit. It is clear from (7) that by opening the electrical circuit (i.e., $I = 0$) the motor torque becomes zero $\tau_m = 0$, while by closing the circuit, the torque generated by the motor is related to V and $\dot{\theta}$. Furthermore, if we assume that the time constant $\tau_{RL} = L/R$ of the equivalent (RL) circuit is small[1] we can make a steady-state approximation of the current: $I \approx I_{ss} = (V - V_{emf})/R_e$, and analytically approximate the torque relation

$$\tau_m \approx n\kappa_{\tau}I_{ss} = \frac{n\kappa_{\tau}}{R_e}V - \frac{n^2\kappa_{\tau}\kappa_{\dot{q}}}{R_e}\dot{\theta}. \qquad (8)$$

This relation shows that τ_m is composed of a voltage dependent driving torque (first term) and a motion dependent braking torque (second term). Accordingly, by shorting the terminals of the circuit together (i.e., $V = 0$), we can generate viscous damping at the motor output shaft

$$\tau_b(\dot{\theta}) = -b_e\dot{\theta} = -\frac{n^2\kappa_{\tau}\kappa_{\dot{q}}}{R_e}\dot{\theta}. \qquad (9)$$

Note that this is a purely passive damping provided by the motor without using any velocity feedback.

From (9) it is clear that the damping constant b_e depends on the resistance of the equivalent circuit R_e and, if we are able to modulate this resistance, we can achieve a variable passive damping effect. This can be done in various ways, for example, one might add a mechanical variable resistor, digital potentiometer, or other integrated circuit across the motor terminals to adjust R_e. In this case the resistance of the equivalent circuit is the sum of the internal resistance of the motor and the variable resistance component.

In our design, variable resistance is not used directly, but instead its effect is emulated using an H-bridge (NSC: LMD18200) circuit. This circuit opens and closes the connection between the motor terminals according to a PWM (pulse-width modulation) signal with controllable duty cycle

---

[1]According to (6), the dynamics of the current follows an exponential tendency given by $e^{-t/\tau_{RL}}$. For our set-up, $\tau_{RL} = L/R = 7 \times 10^{-5}$ (i.e., $L = 1.37 \times 10^{-3}H$, and $R_e = 20.2\Omega$) which makes the steady-state approximation of the current and the associated algebraic relation of the torque (8) feasible.

$u_3 \in [0, 1]$. In this way, the damping is modulated according to the proportion of time that the connection is open or closed: if the connection is open the damping effect vanishes (i.e., $u_3 = 0 \Rightarrow \tau_b = 0$), while if the connection is closed maximal damping is achieved (i.e., $u_3 = 1 \Rightarrow \tau_b = -b_e \dot{\theta}$). Experiments characterising the damping effect in our robotic device are reported in Sec. V-C.

The above variable damping mechanism can be integrated with numerous VSAs simply by rigidly connecting the DC motor shaft to the rotating part of the joint. In our implementation we use a spur gear rigidly attached to the free link to pick up the joint rotation of the MACCEPA (see Fig. 1(b)).

IV. CONTROL FRAMEWORK

In general the non-linear, time-varying nature of many VIAs makes the appropriate selection of the damping non-trivial (see the discussion in Sec. II-C). To avoid this difficulty in this paper we take an optimal control approach to design strategies that are best able to exploit this additional degree of freedom. In the following we briefly describe the formulation of our approach, and the method employed for its solution.

A. Optimal Control Formulation

To define our optimal control problem, we consider a state-space representation of the controlled system dynamics as

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}) \quad \in \mathbb{R}^\nu, \tag{10}$$

where $\mathbf{x} \in \mathbb{R}^\nu$ denotes the system state and $\mathbf{u} \in \mathbb{R}^\mu$ denotes the control inputs. The latter are drawn from the set of admissible controls

$$U = \{\mathbf{u} \in \mathbb{R}^\mu : \mathbf{u}_{min} \preceq \mathbf{u} \preceq \mathbf{u}_{max}\}, \tag{11}$$

where $\mathbf{u}_{min}$ and $\mathbf{u}_{max}$ are the lower and the upper bounds on $\mathbf{u}$, respectively.

Under these constraints, we seek to minimise the cost

$$c = h(\mathbf{x}(T)) + \int_0^T l(\mathbf{x}(t), \mathbf{u}(t), t)\, \mathrm{d}t \in \mathbb{R}, \tag{12}$$

over finite time interval $t \in [0, T]$, where $h(\cdot)$ and $l(\cdot)$ are the terminal and running costs, respectively. The latter are selected according to the application to encode the task goals: in Sec. IV-C we define an example cost function to encode the objective of rapid movement.

B. Solution Method

While there exist various tools for the solution of problems as defined by (10)-(12) [21], our method of choice is to use the Iterative Local Quadratic Regulator (ILQR) method [22], as a fast, approximate solver of optimal control problems. The method is based on approximating the optimal control problem as linear-quadratic and performing iterative improvement of the solutions around a nominal trajectory.

In brief, the algorithm starts with the initial state $\mathbf{x}_0$ and user-supplied nominal initial control sequence $\bar{\mathbf{u}}$, and the corresponding state trajectory $\bar{\mathbf{x}}$. The dynamics (10) are locally approximated as linear around the nominal trajectory

$$\delta \dot{\mathbf{x}} = \mathbf{f_x} \delta \mathbf{x} + \mathbf{f_u} \delta \mathbf{u} \tag{13}$$

(where $\delta \mathbf{x} := \mathbf{x} - \bar{\mathbf{x}}$ and $\delta \mathbf{u} := \mathbf{u} - \bar{\mathbf{u}}$) and the cost is approximated as quadratic

$$\delta c = h_{\mathbf{x}}^\top \delta \mathbf{x}(T) + \delta \mathbf{x}^\top(T) h_{\mathbf{xx}} \delta \mathbf{x}(T) + \int_0^T l_{\mathbf{x}}^\top \delta \mathbf{x} + l_{\mathbf{u}}^\top \delta \mathbf{u}\, \mathrm{d}t$$
$$+ \int_0^T \delta \mathbf{x}^\top l_{\mathbf{xx}} \delta \mathbf{x} + \delta \mathbf{x}^\top l_{\mathbf{ux}} \delta \mathbf{u} + \delta \mathbf{u}^\top l_{\mathbf{uu}} \delta \mathbf{u}\, \mathrm{d}t. \tag{14}$$

Here the 'subscript' notation denotes partial differentiation of the term with respect to the subscript (so, for example, $h_{\mathbf{x}}$ is the Jacobian of h with respect to $\mathbf{x}$ and $h_{\mathbf{xx}}$ is the Hessian of h with respect to $\mathbf{x}$).

Taking the two equations (13), (14), we can then form a local LQR subproblem that can be solved efficiently via a modified Ricatti-system [22]. The latter computes the approximate cost-to-go function in the vicinity of the nominal trajectory, which is then used to update the nominal command trajectory $\bar{\mathbf{u}} \leftarrow \bar{\mathbf{u}} + \delta \mathbf{u}$. The new nominal state trajectory $\bar{\mathbf{x}}$ is computed via numerical simulation, and the whole process is then repeated until convergence (i.e., $\delta c \approx 0$ achieved numerically). A pseudocode of the ILQR method is given in Algorithm 1; for full details, we refer the reader to [22].

C. Application to Rapid Movement on the MACCEPA-VD

In order to apply the above control framework to the MACCEPA-VD, we define the state and control inputs as $\mathbf{x} = (q, \dot{q})^\top = (x_1, x_2)^T$ and $\mathbf{u} = (u_1, u_2, u_3)^\top$, respectively. In this representation, by obtaining $\ddot{q}$ from (1) and replacing it in (10), we obtain

$$\dot{\mathbf{x}} = \begin{pmatrix} \dot{q} \\ \ddot{q} \end{pmatrix} = \mathbf{f}(\mathbf{x}, \mathbf{u}) = \begin{pmatrix} x_2 \\ (\tau(\mathbf{x}, \mathbf{u}) + \tau_f)/J \end{pmatrix}. \tag{15}$$

Due to the mechanical stops in the servos and PWD restrictions, the admissible command range is $(-\pi/2, 0, 0)^\top \preceq \mathbf{u} \preceq (\pi/2, \pi/2, 1)^\top$.

To encode the objective of rapid, accurate reaching to some target, we define the cost function

$$c = \int_0^T (q(t) - q^*)^2 + \varepsilon \mathbf{u}(t)^\top \mathbf{u}(t)\, \mathrm{d}t, \tag{16}$$

where q^* is the target point in joint space, and T is the maximum time duration allowed for the movement. Note that, the first term in (16) influences the solution in several ways. First, it encourages rapid movement by penalising solutions for time spent away from the target point q^*. Second, it ensures the system comes to a halt at the target (if the target is reached before the final time T, the optimal solution is to remain there to avoid incurring further costs). Finally, it penalises overshoot (i.e., deviations for which $q(t) - q^* > 0$). The second term acts to regularise solutions, with ε set to a small constant value (in our experiments $\varepsilon = 10^{-8}$).

V. EXPERIMENTS

In this section we look at how variable damping can be exploited in the context of control of rapid movements on VIAs. We first present numerical results illustrating the benefits of

Algorithm 1 Pseudocode of the ILQR algorithm

GIVEN:
- Model of plant dynamics $\mathbf{f}$ (10)
- Encoding of task through cost function (12)

INITIALISE:
- discretise time as $p = 1 \ldots P$ with $\triangle t = T/(P-1)$
- $\bar{\mathbf{u}}_p = \bar{\mathbf{u}}_0$, corresponding $\bar{\mathbf{x}}_p = \bar{\mathbf{x}}_0$
- Levenberg-Marquardt constant $\lambda = \lambda_{init}$

REPEAT:
- Linearise dynamics around $\bar{\mathbf{x}}_p$ (13)
- Find quadratic approximation of the cost (14)
- solve the obtained local LQR problem $\rightarrow \delta\mathbf{u}_p$

––––––––––––––––––––––––––––––––––––––

Solution to local LQR problem
- for $p = 1 \ldots P$

compute the terms $\mathbf{H}$, $\mathbf{G}$ and $\mathbf{g}$

$$\mathbf{g} := (l_{\mathbf{u}})_p + (\mathbf{f_u}^T)_p (h_{\mathbf{x}})_{p+1}$$

$$\mathbf{G} := (l_{\mathbf{uu}})_p + (\mathbf{f_u}^\top)_p (h_{\mathbf{xx}})_{p+1}(\mathbf{I} + (\mathbf{f_x})_p)$$

$$\mathbf{H} := (l_{\mathbf{uu}})_p + (\mathbf{f_u}^\top)_p (h_{\mathbf{xx}})_{p+1}(\mathbf{f_u})_p$$

compute $\mathcal{H}^{-1}$ (modified Cholesky decomposition)

$$[\mathbf{V}, \mathbf{D}] = eig(\mathbf{H})$$

(the eigenvectors $(\mathbf{V})$ and eigenvalues $(\mathbf{D})$ of $\mathbf{H}$)

$$\mathbf{If}\ (\mathbf{D}(i,i) < 0) : \mathbf{D}(i,i) = 0\ \vee i$$

$$\mathbf{D} = \mathbf{D} + \mathbf{I}\lambda$$

$$\mathcal{H}^{-1} = \mathbf{V}\mathbf{D}^{-1}\mathbf{V}^T$$

compute $\delta\mathbf{u}_p = -\mathcal{H}^{-1}(\mathbf{g} + \mathbf{G}\delta\mathbf{x}_p)$

––––––––––––––––––––––––––––––––––––––

- apply control law forward in time to the linearised system ($\delta\mathbf{x}_{p+1} = \mathbf{f_x}\delta\mathbf{x}_p + \mathbf{f_u}\delta\mathbf{u}_p$ with $\delta\mathbf{x}_1 = 0$)
- compute $\tilde{\mathbf{u}}_p = \bar{\mathbf{u}}_p + \delta\mathbf{u}_p$, corresponding $\tilde{\mathbf{x}}_p$ and δc
- **If** (convergence - i.e. $\delta c \approx 0$)
  STOP
  Else - (update Levenberg-Marquardt constant λ)
    If $(\delta c > 0) : \bar{\mathbf{u}}_p = \tilde{\mathbf{u}}_p, \bar{\mathbf{x}}_p = \tilde{\mathbf{x}}_p$, decrease λ
    Else : increase λ
    End If
  End If
UNTIL convergence

---

variable damping over (optimal) fixed damping schemes as the stiffness varies due to the movement and/or the control. We then present an experiment in hardware, verifying the effectiveness of variable damping in the MACCEPA-VD.

A. Optimal Exploitation of Variable Damping

The goal of our first numerical experiment is to highlight the benefits of optimal damping modulation compared to traditional fixed damping schemes in the context of rapid movements. For this, we compare reaching behaviour on the MACCEPA-VD with (i) variable damping (with the damping command optimised with respect to (16)), (ii) fixed damping (also optimised with respect to (16)) and (iii) damping coupled to the stiffness through (5) according to $\xi = 1$.

To keep the analysis simple, we initially compare behaviours for which the equilibrium position command is fixed at the target ($u_1 = q^*$) and the spring pre-tension command is fixed to $u_2 = \pi/8\,rad$. Note that, due to the
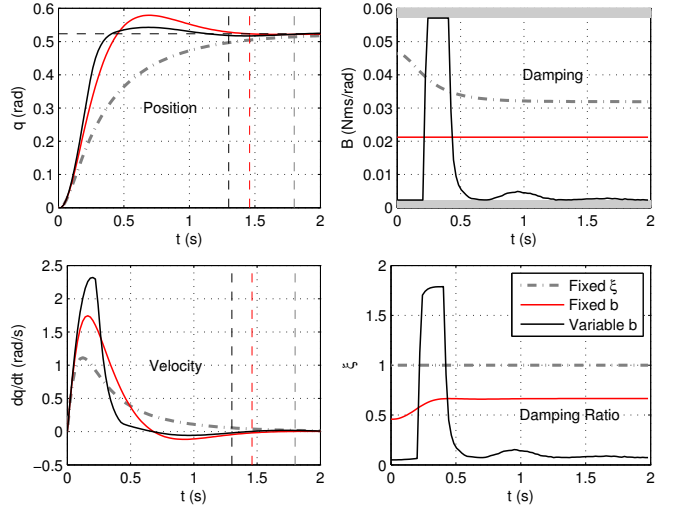

Fig. 3. Joint position, velocity and damping profiles for reaching with (i) critical (ii) optimal fixed and (iii) optimal variable damping. The vertical dashed lines indicate the settling times for the three behaviours and the grey shaded area the damping bounds. The damping ratio over the duration of the movement is shown in the bottom right panel.

dependence of the stiffness k on q (see (3)), the movement causes a natural (uncontrolled) variation of stiffness.

To evaluate performance, for each system we compute the cost incurred under (16), and the settling time which we define as the time at which (i) the absolute distance to the target is less than a threshold value of $|q^* - q(t)| < 0.01\,rad$, and (ii) the velocity is $|\dot{q}(t)| < 0.02\,rad/s$. The results are presented in Fig. 3 and Table I.

	Fixed ξ	Fixed b	Variable b
Cost	0.068	0.048	0.042
Settling time (s)	1.800	1.460	1.300

TABLE I

COST AND SETTLING TIME FOR REACHING WITH FIXED ξ, FIXED DAMPING AND VARIABLE DAMPING.

Looking at Table I, we first observe that the scheme in which there is coupling between stiffness and damping ($\xi = 1$) produces the slowest response and incurs the greatest cost. This confirms the discussion in Sec. II-C: since the stiffness and damping are state and control dependent, $\xi = 1$ *does not ensure critical damping*, and in fact appears to result in an *overdamped response* in this case. For the case of optimal fixed damping the response is faster and (with a small overshoot) settles well before the $2\,s$ time limit. It is interesting to note that for the MACCEPA actuator (see (2)) the joint stiffness drops as the link approaches equilibrium ($q \rightarrow q_0$) resulting in a decreasing stiffness torque τ_k and some inherent stabilisation effect.

Turning again to Table I, however, we observe that the best performance in terms of cost and settling time is achieved by the optimal damping modulation scheme. Looking at the behaviour generated with this approach (see Fig. 3), we see that the device achieves faster convergence to the target, with less overshoot. The optimal damping profile exploits the full range of damping (see Fig. 3, top right) starting with low damping during the initial phase of movement in order to
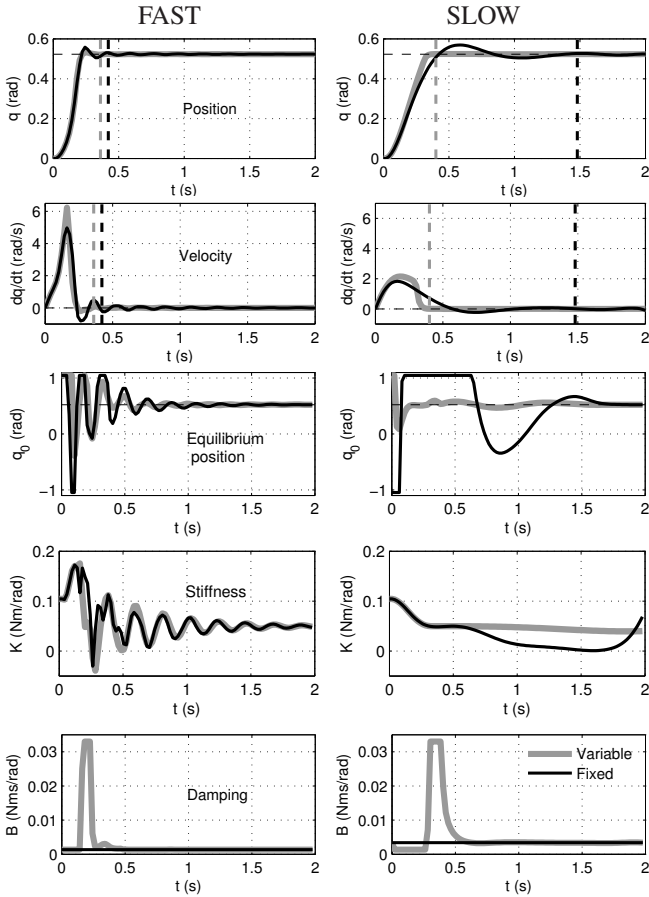
Fig. 4. Joint position, velocity, equilibrium position, stiffness and damping profiles for reaching with (i) optimal fixed (black) and (ii) optimal variable damping (grey) for variable equilibrium position and spring pre-tension. Results for fast (left column) and slow (right column) motor dynamics.

achieve a high initial velocity (see Fig. 3, bottom left), and then switching to the maximum damping as the target is approached in order to brake. The small overshoot results in the damping being again lowered (in order to facilitate return to the target), before finally settling to a low level when the device has halted. The mentioned small overshoot is due to our choice of cost function (16) which does not take into account the end velocity. A short investigation, not presented in here due to space restrictions, indicated that by adding an appropriate terminal cost penalising velocity the system is able to exploit variable damping and achieve the best performance without overshooting.

B. Variable Damping and Speed of the Motor Dynamics

Having shown the benefit of optimal damping modulation when the equilibrium position and spring pre-tension are fixed, we now move on to investigate its effect when they vary according to control. We therefore repeated the experiment from the preceding section, this time allowing the the equilibrium position and stiffness (pre-tension) to be modulated. The latter were also optimised with respect to (16), within servo command ranges $u_1 \in [-\pi/3, \pi/3]$ and $u_2 \in [0, \pi/4]$. The results are presented in the left column of Fig. 4.

From Fig. 4, it can be observed that, when equilibrium position and stiffness are allowed to vary in addition to damping, having the ability to vary damping yields no major improvement over having it fixed: the settling time for variable damping is $0.36\,s$ compared to $0.42\,s$ for fixed damping, and the costs are 0.0365 and 0.0368, respectively.

The reason for the similarity in performance lies in the fact that, in this redundantly actuated system, the equilibrium position and stiffness can be also be used to accelerate and brake the system, even if the damping is fixed. When presented with a constant damping command (Fig. 4, left column, thin black line) the system is able to accelerate by setting the position command beyond the target and increasing the stiffness. Upon approaching the target, the command for the position is brought back in order to brake. When the system overshoots the procedure is repeated until a stable level is achieved. In this way the variable stiffness system can compensate for having fixed damping.

One of the problems with relying on such a strategy, however, is that it assumes the ability to very rapidly vary the equilibrium position and stiffness. In practice, this is not the case with the MACCEPA-VD: while damping modulation (i.e., modulating the PWM signal) is near instantaneous, the servomotors employed in position and stiffness modulation present much slower dynamics.

To test this more realistic case we repeated the experiment with a more accurate model that incorporates these delays in the servomotor response. Empirically we have found a third-order filter model (as described in [23]) to be a good model of the real servo responses. Repeating the optimisation with this model of the MACCEPA-VD, we obtain the results displayed in the right column in Fig. 4.

As can be seen, there is now a much bigger difference in performance. While the use of the equilibrium position and stiffness commands is the same, the slower motor dynamics prevent the fixed damping system from achieving the required braking for rapid settling at the target (ref. Fig. 4, right column, thin black line). In contrast, the variable damping system achieves comparable performance as with the fast motor dynamics (settling time $0.40\,s$) and much better than that of the fixed damping scheme (difference in settling time is $0.52\,s$ and improvement with respect to the cost is cca. 8.7%).

C. Hardware Experiment

Finally, we report experiments verifying the effectiveness of our proposed damping mechanism and show its performance in a rapid reaching task with a hardware implementation of the MACCEPA-VD.

Our implementation of the MACCEPA-VD consists of two servomotors (Hitec HSR-5990TG) controlled with 50 Hz PWM signals from a micro-controller (Atmel ATmega328). Simultaneously, the micro-controller also controls an H-bridge circuit (NSC:LMD18200, see Sec. III-A) that modulates the effective resistance between the terminals of a DC motor/gearbox unit (Maxon A-max 22/110125) - (see Fig. 1). The servomotor positions are measured from their internal
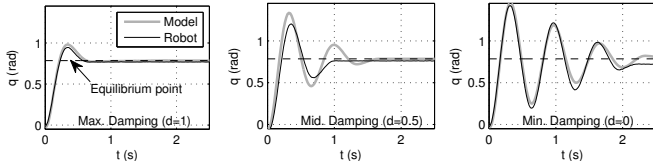
Fig. 5. System response (black) and model predictions (grey) for maximum ($d = 1$), mid-range ($d = 0.5$) and minimum ($d = 0$) damping commands.
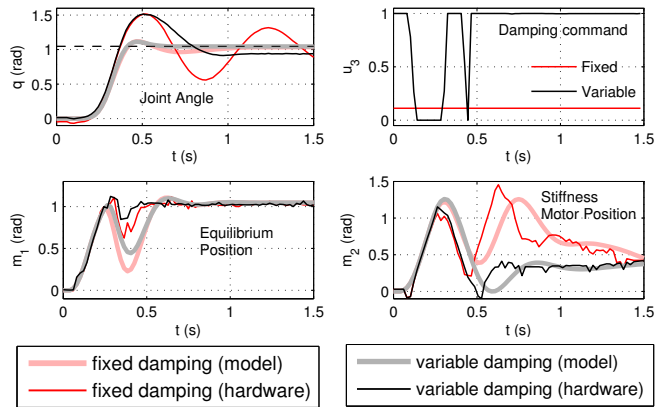

Fig. 6. Model predictions (light, thick lines) and the real robot response (sharp, thin lines). Shown are joint angle, damping command and motor positions when reaching to a target of $q^* = 45°$ (black dashed line) on the MACCEPA-VD actuator with optimal (i) fixed (red lines) and (ii) variable damping (dark lines).

potentiometers and the joint angle is measured with a rotary encoder (Melexis MLX90316GO).

1) Passive Response Test: We first verify the performance of our variable damping mechanism by testing its passive response with different damping commands. Specifically, we conducted tests whereby the position and stiffness commands were fixed ($u_1 = \pi/3 \, rad$ and $u_2 = \pi/2 \, rad$) and the actuator link was held (at rest) at $q = 0$ and then released. The recorded joint trajectories are plotted in Fig. 5 for damping commanded at the maximum ($u_3 = 1$), mid-range ($u_3 = 0.5$) and minimum ($u_3 = 0$) values. Overlaid are predictions of a model based on a system-identification of the device (for details of the latter, please see the Appendix).

As it can be seen, a significant damping range is achieved with the proposed mechanism, approaching critical damping at its maximum (Fig. 5, left) and an underdamped response at its minimum (Fig. 5, right). There is also evidently a close fit between the model and the hardware response at each of the damping levels.

2) Rapid Reaching in Hardware: Finally, we present a comparison of the reaching behaviour of the MACCEPA-VD with fixed versus variable damping. In this experiment the target is set to $q^* = 45°$ and the maximum reaching duration $T = 1.5 \, s$. ILQR was used to to devise the optimal (open loop) command sequence for the three control inputs $(u_1, u_2, u_3)^\top$, starting at rest with $q = 0 \, rad$ and $u_1 = u_2 = 0 \, rad$. We compare the cases where the damping is (i) fixed at the optimal constant value and (ii) allowed to vary freely according to the optimisation. The results are plotted in Fig. 6. (A video recording of the experiment is available in the supplementary material).

As can be seen, (ref. Fig. 6, top left), there is a significant difference in the ability of the system to reach and stabilise at the goal, depending on whether or not the damping is permitted to vary. In both cases we see that all of the control degrees of freedom available are utilised (see Fig. 6, motor positions). Nevertheless, only the system with variable damping is able to stabilise within the $1.5 \, s$ time period. Note also that, as predicted in Sec. V-B, the fixed damping scheme tries to compensate by using the equilibrium position and pre-tension motors to brake. However, in the real device, the relatively slow dynamics of the servomotors prevent it from doing so successfully.

VI. CONCLUSION

In conclusion, we have presented a mechanism for implementing variable damping in variable physical impedance devices. Our approach is based on exploiting electrical inductance effects in a DC motor to create a mechanical damping effect. We have shown (i) how this effect can be modelled with an electromechanical analysis and (ii) the correspondence of our model to measured data collected from a prototype actuator. We have outlined how variable physical damping can be exploited through optimal control in the context of rapid movement tasks. Simulation and hardware experiments suggest a significant benefit of (i) variable damping over fixed damping schemes and (ii) variable damping over variable stiffness when considering the relative speed at which the two can be modulated.

In future work, we intend to extend our experiments to illustrate the role of damping in multi-link variable impedance control and explore the relative benefits of using variable damping versus variable stiffness in the context of energy-optimal behaviour.

APPENDIX

While (9) gives an approximation of the damping torque as a function of the equivalent circuit resistance R_e, in our realisation of the MACCEPA-VD we use PWM-modulated switching to emulate variation of the resistance. This allows the full damping range to be exploited, however, it causes non-linearities in the relation between damping and the commanded duty cycle u_3 due to various effects (e.g., transients when switching between the open and closed configurations of the circuit). We therefore use a data-driven approach for modelling the damping function $b = b(u_3)$.

For this we set up a test rig (Fig. 7) to collect data about the damping response of the motor used in the MACCEPA-VD. Specifically, a DC motor (Maxon A-max 22/110125) was interconnected with the damper motor and used to drive the system as illustrated in Fig. 7 (inset). The equation of motion of this system is given by

$$J_m\ddot{\theta} + \underbrace{(b_{md} + b_m + b_e(u_3))}_{b(u_3)}\dot{\theta} = \underbrace{n_d\kappa_{\tau d}I_d}_{\tau_{md}}, \quad (17)$$

where J_m is the moment of inertia of the interconnected motor-gearbox system, $b(u_3)$ is the total damping of the
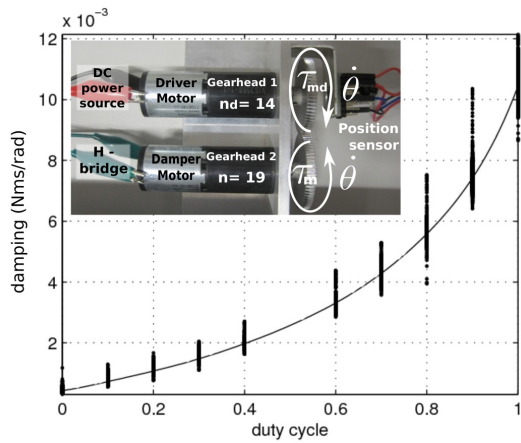
Fig. 7. Damping measurement rig (inset) and model fit.

rig (including the mechanical damping of the two motors $b(0) = b_{md} + b_m$ when $u_3 = 0$), and the additional electrically generated damping $b_e(u_3)$ when $u_3 \in (0, 1]$. On the right hand side of the above equation there are the torques generated by the driver motor[2] according to (7).

During the experiments the driving motor is supplied with a constant voltage $V_d = 24V$ while the duty cycle of the PWM signal $u_3 \in \{0, 0.1, ...1\}$ was changed by ten percent at every $5\,s$. At each of these stages, the system quickly converged to a steady rotation speed ($\ddot{\theta} = 0$). We measured the rotational speed ($\dot{\theta} \neq 0$) and the current of the driver motor I_d. Using this data, we can compute the total damping $b_i(u_3) = n_d \kappa_{\tau d} I_{di} / \dot{\theta}_i$ (based on (17) with $J_m \ddot{\theta} = 0$), for every data point $i \in \{1, 2, ...N\}$ for a given u_3.

From these values the *total mechanical damping* of the system is estimated as a mean of the computed damping constants at $u_3 = 0$: $\hat{b}(0) = \frac{1}{N} \sum_i^N b_i(0)$. The latter is then used to estimate the electrical damping effect on the remaining data: $\hat{b}_{ei}(u_3) = b_i(u_3) - \hat{b}(0)$. Finally, using the quadratic scaling of the damping constant with the gear ratio (see (9)), we define $\hat{b}_m = n^2/(n^2 + n_d^2)\hat{b}(0)$. The latter is used to estimate the total damping coefficient as

$$\hat{b}_i(u_3) = \hat{b}_m + \hat{b}_{ei}(u_3). \tag{18}$$

Using the estimates from (18) as data, we fit a model with normalised Gaussian radial basis functions[3] $\boldsymbol{\phi}(u_3) \in \mathbb{R}^M$,

$$\hat{b}(u_3) = \mathbf{w}^\top \boldsymbol{\phi}(u_3) \tag{19}$$

where $\mathbf{w} \in \mathbb{R}^M$ is a vector of weights. In Fig. 7 we plot the damping as a function of duty cycle u_3 obtained through this process.

Finally, re-inserting the damper motor into the MACCEPA-VD, the joint damping torque is estimated as

$$\hat{\tau}_b(\dot{q}, u_3) = -n_q^2 \hat{b}(u_3)\dot{q} \tag{20}$$

where $n_q = 1.8$ is the gear reduction between the link and the motor shaft.

[2] The quantities: n_d, $\kappa_{\tau d}$ and I_d define the gear ratio, torque constant, and the current of the driver motor respectively.

[3] where $\phi_\mu(u_3) = K_\mu(u_3)/\sum_{\nu=1}^M K_\nu(u_3)$, $K_\mu(u_3) = \exp(-(u_3 - c_\mu)^2/2\sigma^2)$ and c_μ, $\mu = 1 \ldots M$.

REFERENCES

[1] C. English and D. Russell, "Implementation of variable joint stiffness through antagonistic actuation using rolamite springs," *Mechanism and Machine Theory*, vol. 34, no. 1, pp. 27–40, 1999.
[2] G. Tonietti, R. Schiavi, and A. Bicchi, "Design and control of a variable stiffness actuator for safe and fast physical human/robot interaction," in *IEEE Int. Conf. Robotics & Automation*, 2005.
[3] N. Vitiello, T. Lenzi, J. McIntyre, S. Roccella, E. Cattin, F. Vecchi, and M. Carrozza, "Characterization of the neurarm bio-inspired joint position and stiffness open loop controller," in *IEEE Int. Conf. Biomedical Robotics and Biomechatronics*, 2008.
[4] D. J. Braun, M. Howard, and S. Vijayakumar, "Exploiting variable stiffness in explosive movement tasks," in *RSS*, 2011.
[5] F. Petit and A. Albu-Schffer, "State feedback damping control for a multi dof variable stiffness robot arm," in *IEEE Int. Conf. Robotics & Automation*, 2011.
[6] M. Laffranchi, N. G. Tsagarakis, and D. G.Caldwell, "A variable physical damping actuator (VPDA) for compliant robotic joints," in *IEEE Int. Conf. Robotics & Automation*, 2010.
[7] M. Howard, D. Mitrovic, and S. Vijayakumar, "Transferring impedance control strategies between heterogeneous systems via apprenticeship learning," in *Humanoid Robots (Humanoids), 2010 10th IEEE-RAS International Conference on*, pp. 98 –105, dec. 2010.
[8] J. Nakanishi, K. Rawlik, and S. Vijayakumar, "Stiffness and temporal optimization in periodic movements: An optimal control approach," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pp. 718–724, IEEE, 2011.
[9] J. Nakanishi and S. Vijayakumar, "Exploiting passive dynamics with variable stiffness actuation in robot brachiation," in *Robotics: Science and Systems*, 2012.
[10] K. Rawlik, M. Toussaint, and S. Vijayakumar, "An approximate inference approach to temporal optimization in optimal control," 2011.
[11] D. Mitrovic, S. Klanke, M. Howard, and S. Vijayakumar, "Exploiting sensorimotor stochasticity for learning control of variable impedance actuators," in *IEEE Int. Conf. Humanoid Robots*, 2010.
[12] D. Mitrovic, S. Klanke, and S. Vijayakumar, "Learning impedance control of antagonistic systems based on stochastic optimization principles," *IJRR*, vol. 30, no. 5, pp. 556–573, 2011.
[13] D. Mitrovic, S. Klanke, and S. Vijayakumar, "Adaptive optimal feedback control with learned internal dynamics models," *From Motor Learning to Interaction Learning in Robots*, pp. 65–84, 2010.
[14] D. Mitrovic, S. Klanke, R. Osu, M. Kawato, and S. Vijayakumar, "A computational model of limb impedance control based on principles of internal model uncertainty," *PLoS ONE*, vol. 5:10, p. 13601, 2010.
[15] R. Ham, T. Sugar, B. Vanderborght, K. Hollander, and D. Lefeber, "Compliant actuator designs," *Robotics Automation Magazine, IEEE*, vol. 16, pp. 81 –94, september 2009.
[16] R. V. Ham, B. Vanderborght, M. V. Damme, B. Verrelst, and D. Lefeber, "MACCEPA, the mechanically adjustable compliance and controllable equilibrium position actuator: Design and implementation in a biped robot," *Robotics & Autonomous Sys.*, vol. 55, no. 10, pp. 761–768, 2007.
[17] E. Guglielmino, C. W. Stammers, K. A. Edge, T. Sireteanu, and D. Stancioiu, "Damp-by-wire: Magnetorheological vs friction dampers," in *16th IFAC World Congress*, 2005.
[18] M. Unsal, C. Niezrecki, and I. Crane, C., "Two semi-active approaches for vibration isolation: piezoelectric friction damper and magnetorheological damper," in *IEEE Int. Conf. Mechatronics*, 2004.
[19] J. Mehling, J. Colgate, and M. Peshkin, "Increasing the impedance range of a haptic display by adding electrical damping," in *Eurohaptics*, 2005.
[20] M. 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.), vol. 5024 of *LNCS*, pp. 53–62, 2008.
[21] D. E. Kirk, *Optimal Control Theory: An Introduction*. Prentice-Hall, 1970.
[22] W. Li and E. Todorov, "Iterative linear-quadratic regulator design for nonlinear biological movement systems," in *Int. Conf. on Informatics in Control, Automation and Robotics*, vol. 1, pp. 222–229, 2004.
[23] T. Wada, M. Ishikawa, R. Kitayoshi, I. Maruta, and T. Sugie, "Practical modeling and system identification of R/C servo motors," in *IEEE Int. Conf. Control & Automation*, 2009.