

Gain-Scheduled Inverse Optimal Satellite Attitude Control

NADJIM M. HORRI, Member, IEEE
PHIL PALMER
MARK ROBERTS
University of Surrey

Despite the theoretical advances in optimal control, satellite attitude control is still typically achieved with linear state feedback controllers, which are less efficient but easier to implement. A switched controller is proposed, based on inverse optimal control theory, which circumvents the complex task of numerically solving online the Hamilton-Jacobi-Bellman (HJB) partial differential equation of the global nonlinear optimal control problem. The inverse optimal controller is designed to minimize the torque consumption pointwise, while imposing the stabilization rate of a linear benchmark controller. The controller is then modified by gain scheduling to achieve a tradeoff enhancement compared with the benchmark controller, while maintaining torque saturation limits. The extent to which performance can be enhanced is shown to be dependent on the controller parameters. A controller tuning analysis shows how a design settling time limit can be achieved, within the problem's constraints on the maximum torque and the total integrated torque. The proposed optimization approach is globally stabilizing and presents low implementation complexity, which is highly desirable given the limited resources onboard satellites.

Manuscript received November 24, 2010; revised April 6 and June 3, 2011; released for publication September 15, 2011.

IEEE Log No. T-AES/48/3/944026.

Refereeing of this contribution was handled by M. Jah.

This work was supported by the European Union Marie Curie Astrodynamics research network ASTRONET.

Authors' addresses: N. M. Horri and P. Palmer, Surrey Space Centre, University of Surrey, Guildford, Surrey, GU2 7XH, United Kingdom, E-mail: (n.horri@surrey.ac.uk); M. Roberts, Department of Mathematics, University of Surrey, Guildford, United Kingdom.

0018-9251/12/\$26.00 © 2012 IEEE

I. INTRODUCTION

Two different viewpoints have so far dominated nonlinear optimal control theory: the local Euler-Lagrange approach from calculus of variations and the global Hamilton-Jacobi-Bellman (HJB) approach based on dynamic programming. The Euler-Lagrange approach only solves the optimization problem locally and requires numerical techniques to solve boundary value problems, which are complex to implement onboard a satellite. The HJB approach is also complex to implement and involves solving partial differential equations using numerical methods (see [16] for example). As a consequence, despite the theoretical superiority of optimal control, the operational attitude control of small satellites is still predominantly based on standard linear stabilizing feedbacks, such as PD-type control laws, which are preferred for being easier to tune, validate and implement. This is for example the case of the satellites developed by Surrey Satellite Technology Limited (SSTL).

In this paper a switched controller is constructed, based on inverse optimization theory, which is used partly because it circumvents the task of numerically solving an HJB equation, but also to ensure the enhancement of a tradeoff between convergence rate and torque consumption compared with a benchmark controller [1, 2, 8, 11].

The starting point of inverse optimal control is to construct a stabilizing feedback controller, known as the benchmark, based on a control Lyapunov function (CLF) V . Once V is known, an inverse optimal controller is designed to analytically solve the HJB equation associated with a meaningful cost function.

Inverse optimal control techniques are being increasingly considered for a large range of control systems applications (see [4], [9], [10], [13]). Inverse optimality was used for satellite attitude control by Krstic and Tsotras in [11], followed by [15] and [17] where continuous state feedback controllers were designed to solve an HJB (or HJ-Isaacs) equation by minimizing a "meaningful" weighted cost function. However, the meaningful cost found to be minimized in these references is a complicated function of the state and control vectors and these mathematically meaningful cost functions are not always clearly worth minimizing in practice. The approach considered here is not to focus on determining the expression of the meaningful cost being minimized (known to exist [2]) but to formulate the optimization tradeoff as a nonlinear program at each state (as in [1], [2]) to impose performance constraints that are not solved by standard controllers as in [17]. We then use performance metrics such as settling time and integrated torque to optimize controller parameters. In this paper we allow the controller to be switched and the continuous inverse optimal controller of [11] is used for comparison purposes.

The most relevant application of inverse optimal control to this work was in [1]. This paper aims to extend the results in [1] and [2], where the minimum norm optimization problem of minimizing a norm of the control effort subject to a constraint on the convergence rate of a Lyapunov function was solved at each state. It was proven in [2] that the resulting switched min-norm controller has a zero torque mode and is inverse optimal with respect to a meaningful cost function of the form $\int_0^\infty \{l(\mathbf{x}) + r(\mathbf{x}, \mathbf{u})\} dt$, where $l(\mathbf{x})$ is a continuous positive definite function of the state vector \mathbf{x} (with $l(0) = 0$) and $r(\mathbf{x}, \mathbf{u})$ is a continuous positive definite and increasing convex function of the norm of the control vector \mathbf{u} . The meaningful cost function is generally unspecified in this approach, where the optimization problem is posed using a nonlinear programming formulation.

A more specific objective of the paper is to exploit inverse optimal minimum norm control to enhance a benchmark controller in the sense of reducing the settling time to a limit that can be prescribed (not usual in inverse optimal control) for a given level of the integrated torque, but within torque saturation constraints (not included in [1]). Time is not claimed to be minimized because the meaningful cost function is not necessarily the one that minimizes time for a prescribed integrated torque, although the latter (as well as the dual problem of minimizing integrated torque subject to a fixed time) also involves a zero torque phase [14].

The switched inverse optimal attitude control approach is applied to a small satellite. The parameters of a typical 100 kg microsatellite equipped with three orthogonal wheels, are considered. The minimum norm control law of [1] was applied to satellite attitude control but the attitude parameterization chosen in that paper was subject to singularities. The Lyapunov function was quadratic with a bilinear coupling term and a similar Lyapunov function is used here, based on quaternion parameters. Standard minimum norm optimization has been shown in [8] to significantly reduce the overall torque expenditure, while slightly degrading settling time compared with the controller used as a benchmark. In [1] slew time was also enhanced, but only because the benchmark controller was inefficiently tuned. It was heavily underdamped.

The following are key contributions of this paper compared with previous work in [1], [8], and [11].

- 1) A gain-scheduled formulation of the inverse optimal minimum norm controller is proposed to achieve a settling time enhancement compared with a PD benchmark controller, within constraints on the maximum instantaneous torque and the integrated torque. The controller avoids the risk of torque saturation, which would be inevitable by tuning the standard minimum norm controller of [1] to

deliver the same integrated torque as the benchmark controller. This controller is also shown to outperform the widely cited inverse optimal controller proposed in [11], in terms of settling time for a given integrated torque and under maximum torque limitations.

- 2) An analytic analysis is provided to link performance metrics to controller parameters. A controller tuning procedure is proposed to achieve a designed settling time limit, within admissible constraints on the integrated and instantaneous torque (Section IV).

The proposed gain-scheduled inverse optimal controller consists of switching between three modes:

- 1) a low gain PD control mode, which is taken to be a benchmark controller,
- 2) a zero torque coasting phase (typical of problems involving an integrated torque/time minimization tradeoff),
- 3) a high gain PD control mode (feasible with this approach and allows for a settling time enhancement).

The paper is organized as follows. Section II describes the attitude dynamic and kinematic models of a satellite controlled by three reaction wheels. Section III describes the switched minimum norm control approach, based on inverse optimality, with application to 3-axis attitude control. Background information is given in Sections III-A and III-B on the standard minimum norm controller. The bilinear Lyapunov function used to design the controller is introduced in Section III-C. A gain-scheduled formulation of the minimum norm controller is proposed in Section III-D. The continuous inverse optimal approach of [11], which is used for comparison, is briefly described in Section III-E. Section III is concluded by a numerical simulation subsection that demonstrates performance enhancement of a benchmark controller by gain-scheduled min norm optimization. The controller is also compared with the inverse optimal controller from [11]. The analysis of the controller tuning is given in Section IV, where linear eigenaxis rotations are assumed for simplicity. Section IV also includes a numerical simulation subsection and is concluded by a discussion of implementation considerations.

II. DYNAMIC AND KINEMATIC MODELS

In the disturbance-free case, the dynamic model of a rigid body satellite is given by Euler's rotational equation of motion:

$$\dot{\mathbf{L}} + \boldsymbol{\omega} \times \mathbf{L} = 0 \quad (1)$$

where \mathbf{L} denotes the total angular momentum and $\boldsymbol{\omega} = [\omega_1, \omega_2, \omega_3]^T$ is the angular velocity vector, both in the body fixed reference frame.

The total angular momentum of the satellite with reaction wheels is given by

$$\mathbf{L} = \mathbf{I}\boldsymbol{\omega} + \mathbf{h} \quad (2)$$

where $\mathbf{h} = [h_1, h_2, h_3]^T$ represents the angular momentum generated by the reaction wheels in the body frame.

By substituting \mathbf{L} from (2), into (1), we obtain the general Euler's rotational equation using three orthogonal reaction wheels, aligned on the body axes of the satellite:

$$\begin{aligned} I_1\dot{\omega}_1 &= (I_2 - I_3)\omega_2\omega_3 + N_1 - \omega_2h_3 + \omega_3h_2 \\ I_2\dot{\omega}_2 &= (I_3 - I_1)\omega_1\omega_3 + N_2 - \omega_3h_1 + \omega_1h_3 \\ I_3\dot{\omega}_3 &= (I_1 - I_2)\omega_1\omega_2 + N_3 - \omega_1h_2 + \omega_2h_1 \end{aligned} \quad (3)$$

where $\mathbf{I} = \text{diag}(I_1, I_2, I_3)$: moment of inertia matrix of the body of the satellite about its centre of mass. We define the control torque of the wheels on the satellite as $N_i = -\dot{h}_i$, $i = 1, 3$.

Using the quaternion parameterization of attitude kinematics, the kinematic model of a satellite is given by

$$\begin{aligned} \dot{\mathbf{q}} &= -\frac{1}{2}\boldsymbol{\omega} \times \mathbf{q} + \frac{1}{2}q_4\boldsymbol{\omega} \\ \dot{q}_4 &= -\frac{1}{2}\boldsymbol{\omega}^T \mathbf{q} \end{aligned} \quad (4)$$

where $\bar{\mathbf{q}} = [\mathbf{q}, q_4]^T = [q_1, q_2, q_3, q_4]^T$ is the attitude quaternion of the satellite.

For the sake of convenience, we represent our system as an affine control system of the form:

$$\begin{aligned} \dot{\mathbf{x}} &= f(\mathbf{x}, \mathbf{h}) + g(\mathbf{x})\mathbf{u} \\ \dot{\mathbf{h}} &= -\mathbf{u} \end{aligned} \quad (5)$$

where

$\mathbf{x} = [\bar{\mathbf{q}}, \boldsymbol{\omega}]^T \in \mathbb{R}^7$ is the state vector to be controlled to $\mathbf{x}_{\text{eq}} = [0, 0, 0, 1, 0, 0, 0]^T$ (at equilibrium, \mathbf{h} equals the total momentum),

$$\mathbf{u} = [N_1, N_2, N_3]^T \in \mathbb{R}^3,$$

$$f(\mathbf{x}, \mathbf{h}) = [\dot{\bar{\mathbf{q}}}, \dot{\boldsymbol{\omega}}]_{\mathbf{u}=0_{3 \times 1}} : \mathbb{R}^7 \rightarrow \mathbb{R}^7 \text{ and}$$

$$g(\mathbf{x}) = \begin{bmatrix} \mathbf{0}_{4 \times 3} \\ \mathbf{I}_{3 \times 3}^{-1} \end{bmatrix}^T : \mathbb{R}^7 \rightarrow \mathbb{R}^{7 \times 3}$$

are, respectively, smooth vector and matrix functions of \mathbf{x} .

This mathematical model of attitude kinematics and dynamics is used for the simulation analysis of the proposed control techniques in the case of three axis attitude control. A simpler model is then used in Section IV-G on the numerical analysis of controller tuning.

III. SWITCHED INVERSE OPTIMAL ATTITUDE CONTROL BY MINIMUM NORM OPTIMIZATION

In this section we introduce minimum norm inverse optimal control theory for the affine control system of (5). Inverse optimal design consists of

determining a stabilizing feedback $\mathbf{u}(\mathbf{x})$, which is optimal in the sense of minimizing a meaningful cost function with a positive penalty on both the control effort and the states. Following Freeman and Kokotovic (see [2]), minimum norm inverse optimization strategies are based on the existence of a CLF and its' associated stabilizing controller $\mathbf{u}(\mathbf{x})$. After the introductory background on the minimum norm control problem, a new gain-scheduled approach to switched (variable structure) minimum norm satellite attitude control is proposed. The continuous inverse optimal control approach of [11], which is used for comparison purposes, is also described in this section.

A. Background on the Standard Minimum Norm Controller [1, 2]

For our affine control system, we introduce the concept of CLF (see [18], [1]). The existence of a CLF for a stabilizing controller $\mathbf{u}(\mathbf{x})$ implies that the system is stabilizable. Stabilizability and controllability are generally implicitly assumed and satisfied for a fully actuated rigid body spacecraft (see [11], [15], [1]). For example, [11] assumes the existence of a stabilizing controller with respect to a Lyapunov function (a stronger condition than stabilizability). Similarly, we assume stabilizability at an equilibrium point $\mathbf{x} = \mathbf{x}_{\text{eq}}$. (A rigorous controllability check is possible by verifying that a matrix of Lie brackets has the full rank of the system). We verify the existence and stability of the equilibrium point \mathbf{x}_{eq} . A CLF (refer for example to [12]) associated with a controller $\mathbf{u}(\mathbf{x})$ is a continuously differentiable positive definite function $V(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}^+$, such that $V(\mathbf{x}_{\text{eq}}) = 0$ and

$$\inf_{\mathbf{u}} [\dot{V} = L_f V(\mathbf{x}, \mathbf{h}) + L_g V(\mathbf{x})\mathbf{u}(\mathbf{x})] < 0, \quad \forall \mathbf{x} \neq \mathbf{x}_{\text{eq}} \quad (6)$$

where the functions $L_f V$ and $L_g V$, respectively, represent the Lie derivatives of the Lyapunov function $V(\mathbf{x})$ in the directions of the vector field f and the matrix g . $L_f V = (\partial V^T / \partial \mathbf{x})f$ is a scalar function of \mathbf{x} and \mathbf{h} and $L_g V = (\partial V^T / \partial \mathbf{x})g$ is a row vector of the same dimension as the control vector \mathbf{u} .

For our system, $L_g V \in \mathbb{R}^{1 \times 3}$. For ease of notation, $L_f V(\mathbf{x}, \mathbf{h})$ and $L_g V(\mathbf{x})$ are denoted $L_f V$ and $L_g V$.

Equilibrium point: It is in general necessary to verify that $f(\mathbf{x}_{\text{eq}}) = 0$. This is the case for our spacecraft model with $n = 7$ and $\mathbf{x}_{\text{eq}} = [0, 0, 0, 1, 0, 0, 0]^T$. The value $q_{4\text{eq}} = 1$ is used rather than $q_{4\text{eq}} = -1$ (same physical orientation) for reasons explained after (15).

For a CLF function, there exists by definition a controller \mathbf{u} that makes $\dot{V}(\mathbf{x}, \mathbf{u}) < 0$, $\forall \mathbf{x} \neq \mathbf{x}_{\text{eq}}$. Equation (6) can be satisfied by more than just one controller. As in [1], [2], a particular inverse optimal controller is constructed here from the solution to a pointwise

minimum norm optimization problem:

$$\begin{aligned} & \text{Minimize } \|\mathbf{u}\|^2 \text{ at each point } \mathbf{x} \\ & \text{subject to } \dot{V}(\mathbf{x}) = L_f V + L_g V \mathbf{u}(\mathbf{x}) \quad (7) \\ & \qquad \qquad \leq -\sigma(\mathbf{x}) \leq 0 \end{aligned}$$

where $\|\cdot\|$ denotes the Euclidian 2-norm. Note that $L_g V$ depends on x only and although $L_f V$ is a function of both \mathbf{x} and \mathbf{h} , the controller is designed in Section III-B to be independent from $L_f V$. Hence, \mathbf{u} is a function of \mathbf{x} only. Also note that the pointwise minimization of $\|\mathbf{u}\|^2$ or $\|\mathbf{u}\|$ leads to the same solution because $\|\mathbf{u}\|^2$ and $\|\mathbf{u}\|$ have the same level sets.

The function $-\sigma(\mathbf{x})$ can be viewed as describing a nonlinear stability margin [1, 2] but also as a way of imposing pointwise convergence rate constraints. Reference [2] showed how (7) solves an HJB equation for a meaningful cost.

The nonlinear program in (7) can be written more compactly as a least norm problem:

$$\text{Minimize } \|\mathbf{u}\|^2 \text{ st. } \langle \mathbf{a}, \mathbf{u} \rangle \leq b \quad (8)$$

with

$$\begin{aligned} b &= -L_f V - \sigma(\mathbf{x}) \\ \mathbf{a} &= L_g V. \end{aligned} \quad (9)$$

As stated in [2], the analytic solution to this least norm problem is the following switched control law, which we refer to as the standard minimum norm controller:

$$\mathbf{u} = \begin{cases} \frac{b\mathbf{a}^T}{\mathbf{a}\mathbf{a}^T} & \text{if } b < 0 \\ 0 & \text{if } b \geq 0 \end{cases}. \quad (10)$$

Note that b is a scalar, while \mathbf{a} is a row vector and so the denominator $\mathbf{a}\mathbf{a}^T$ is a scalar, not a matrix (as also stated in [1] and [2]). The solution is a projection onto the space spanned by the vector \mathbf{a}^T if the constraint is not solved by turning the controller off. In this case, the projection to the boundary of the inequality constraint $\mathbf{a}\mathbf{u} \leq b$ is the solution of the minimum norm problem with the equality constraint $\mathbf{a}\mathbf{u} = b$. Otherwise, when the constraint $b \geq 0$ is satisfied, then the minimum norm solution satisfying the constraint of (7) is simply $\mathbf{u} = \mathbf{0}_{3 \times 1}$.

In the following, the choice of $\sigma(\mathbf{x})$ is derived from the existence of a stabilizing benchmark controller $k(\mathbf{x})$ with respect to a Lyapunov function V . The existence of such a stabilizing controller was also assumed in [1], [11].

B. Standard Minimum Norm Attitude Controller with a PD Benchmark

As in [1], we now apply the above inverse optimization theory to the satellite attitude control problem. A natural choice of the negativity margin

function $-\sigma$ is given by the time derivative of a Lyapunov function (representing a certain convergence rate), under the effect of a benchmark controller $k(\mathbf{x})$:

$$\sigma(\mathbf{x}) = -\dot{V}_{k(\mathbf{x})} = -L_f V - L_g V k(\mathbf{x}). \quad (11)$$

The scalar b in (10) is then given by

$$b = L_g V k(\mathbf{x}). \quad (12)$$

As a benchmark controller, we consider the standard PD law:

$$k(\mathbf{x}) = \mathbf{u}_{\text{PD}} = -k_p \mathbf{I} \mathbf{q} - k_d \mathbf{I} \dot{\omega}. \quad (13)$$

This quaternion feedback controller is chosen as a benchmark for simplicity and because linear or feedback stabilizing quaternion controllers are common both in the literature [3, 5–7] and in the onboard software of most satellites. As in [1], the final expression of the standard minimum norm attitude control law (10) is now given by

$$\mathbf{u}_{\text{opt}} = \begin{cases} \frac{(L_g V \mathbf{u}_{\text{PD}})(L_g V)^T}{L_g V (L_g V)^T} & \text{if } L_g V \mathbf{u}_{\text{PD}} < 0 \\ 0 & \text{if } L_g V \mathbf{u}_{\text{PD}} \geq 0 \end{cases} \quad (14)$$

where the term between square brackets is the scalar b of (10) and the projection is in the direction of $L_g V^T$. Provided that a benchmark controller exists, this controller yields a solution that is optimal pointwise (for a small time horizon) with guaranteed stability. The zero torque mode is typical of problems achieving a tradeoff between integrated torque and maneuver time such as the minimum time fuel problem in [14]. The practically interesting property of the approach is the possibility, in principle, of enhancing a benchmark controller, although modifications to controller (14) are shown in Section III-F.2 to be necessary to make performance enhancement a reality under realistic conditions such as torque limitation constraints.

C. Lyapunov Function with Bilinear Coupling

We apply the above theory to the system of equations (3),(4). For our minimum norm problem, we consider the following Lyapunov function with bilinear coupling to prove stability by the PD benchmark controller of (13):

$$\begin{aligned} V &= (k_p + \gamma k_d)(\mathbf{q}^T \mathbf{q} + (1 - q_4)^2) \\ &+ \frac{1}{2} \boldsymbol{\omega}^T \boldsymbol{\omega} + \gamma \mathbf{q}^T \boldsymbol{\omega} \end{aligned} \quad (15)$$

$$V = 2(k_p + \gamma k_d)(1 - q_4) + \frac{1}{2} \boldsymbol{\omega}^T \boldsymbol{\omega} + \gamma \mathbf{q}^T \boldsymbol{\omega}.$$

The stability proof is outlined in the Appendix. The last term of (15) is introduced to allow for the adjustment of the tradeoff between slew time and torque expenditure via the weighting factor γ . With this choice of V , $L_g V = (\boldsymbol{\omega} + \gamma \mathbf{q})^T \mathbf{I}^{-1}$ (with \mathbf{I}

invertible). The stable equilibrium is $\bar{\mathbf{q}}_{\text{eq}} = [0, 0, 0, 1]^T$, $\omega_{\text{eq}} = [0, 0, 0]^T$. As in [21], [17], [5], we chose $q_{4\text{eq}} = 1$ rather than $q_{4\text{eq}} = -1$, which corresponds to the same physical orientation. As stated in [17], q_4 can be initialized as positive. Also, it is relatively simple to extend the analysis to the case $q_{4\text{eq}} = -1$ (see [3] for example).

Note on the validity of V: In (15), it has to be understood that the positivity of V and negativity of \dot{V} will be satisfied for any state vector \mathbf{x} but only for a certain region in the parameter space (k_p, k_d, γ) that is determined from the stability analysis. This was also the case in [1], [21]. In particular, it cannot be allowed to be too large for V to be a valid Lyapunov function. A very similar Lyapunov function with a bilinear coupling was also proposed in [1] using a different attitude parameterization and in [21], [17] with quaternions, although not for a minimum norm problem. Another difference is that the Lyapunov function in [21], [17], [1] was dependent on \mathbf{I} to make stability easier to prove for a benchmark PD law $\mathbf{u}_{\text{PD}} = -k_p \mathbf{q} - k_d \omega$ that was independent from \mathbf{I} . Conversely, the Lyapunov function in (15) is independent from \mathbf{I} to simplify the stability proof (see Appendix) for the controller $\mathbf{u}_{\text{PD}} = -k_p \mathbf{I} \mathbf{q} - k_d \mathbf{I} \omega$, where \mathbf{I} is the moment of inertia matrix. This allows the gains to be proportionally higher on the axes with higher moments of inertia. The Lyapunov function in [11] was also independent from \mathbf{I} for similar reasons. Also, attitude control in all those references was performed by thrusters rather than wheels. With the Lyapunov function of (15), the control torque given in (14) is typically switched off between two straight lines in the phase space domain: $L_g V = 0$ and $\mathbf{u}_{\text{PD}} = 0$. The first switching curve is the straight line $L_g V^T = \mathbf{I}^{-1}(\omega + \gamma \mathbf{q}) = 0$. In practice, orthogonality between $L_g V$ and \mathbf{u}_{PD} was found to be unlikely by taking γ to be a fraction of k_p/k_d .

D. Switched Control by Gain-Scheduled Minimum Norm Optimization

In this subsection, we propose a modification to the standard minimum norm controller of (14). The standard min-norm controller is shown in the numerical simulation section to enhance settling time when tuned to deliver the same integrated torque expenditure as the benchmark controller. However, the same overall torque expenditure can only be obtained by allowing the maximum instantaneous torque of the min-norm controller during the maneuver to be higher than the maximum torque of the benchmark controller, hence the need for the proposed gain-scheduled approach. To conduct a fair comparison between the standard, gain-scheduled min-norm controllers and the PD law used as a benchmark, we impose the same maneuver to these controllers, tune them to

achieve similar integrated torque expenditure defined as $\int_0^{t_f} \mathbf{u}^T(\tau) \mathbf{u}(\tau) d\tau$ and compare settling times.

Gain-scheduled formulation of minimum norm control: Based on minimum norm optimization, we construct a minimum norm controller that reduces settling time compared with a benchmark controller, not only for the same total integrated torque but also without exceeding the maximum instantaneous torque that would be obtained from the benchmark controller. The proposed switching logic is therefore based on gain scheduling to avoid the risk of torque saturation. The proposed gain scheduling scheme is as follows.

- 1) We first adopt the same gains as the PD law used as a benchmark (k_{p1}, k_{d1}) , when there is a risk of torque saturation. For the case of rest-to-rest maneuvers, these lower gains are used initially when the torque saturation risk is highest.
- 2) The controller is then turned off between two straight lines in the phase space domain (minimum norm approach with bilinear Lyapunov function).
- 3) When the controller is turned on again, higher gains (k_{p2}, k_{d2}) are adopted because there are no torque saturation issues at that stage and the torque has been zero for a significant time, allowing for an integrated torque expenditure margin. The higher gains are selected to deliver the same or a lower overall torque expenditure as the PD benchmark without ever exceeding the initial torque of the benchmark controller. Settling time is reduced precisely because higher gains are achievable for a similar overall torque with this gain-scheduled minimum norm approach.

The use of three control modes can be justified by the fact that the bang-off-bang controller, which solves an optimal tradeoff between time and integrated torque (see [14]), also has three modes: acceleration, coasting, and deceleration. The proposed feedback controller has the robustness advantages of closed-loop control. To demonstrate that gain scheduling is required, simulations in Section III-F.2 show how the use of one set of higher gains in a standard min-norm controller with no gain scheduling (see (14)) can amplify the torque initially with a risk of torque saturation.

To guarantee stability, the gain-scheduled min norm controller is first formulated as an inverse optimal control problem:

$$\text{Minimize } \|\mathbf{u}\| \quad \text{st.} \quad \begin{cases} \dot{V} < \dot{V}(\mathbf{x}, \mathbf{u}_1) & \text{if } \|\mathbf{u}_2\|_\infty \geq \epsilon \\ \dot{V} < \dot{V}(\mathbf{x}, \mathbf{u}_2) & \text{if } \|\mathbf{u}_2\|_\infty < \epsilon \end{cases} \quad (16)$$

The convergence rate constraint of the optimization problem depends on the phase space region. The controller is switched. This inverse optimization problem is solved by a gain-scheduled switched

control law, which is given by

$$\mathbf{u}_{\text{opt}} = \begin{cases} \frac{(L_g V \mathbf{u}_2)(L_g V)^T}{L_g V (L_g V)^T} & \text{if } L_g V \mathbf{u}_2 < 0 \quad \text{and} \quad \|\mathbf{u}_2\|_\infty < \epsilon \\ 0 & \text{if } L_g V \mathbf{u}_2 \geq 0 \\ \frac{(L_g V \mathbf{u}_1)(L_g V)^T}{L_g V (L_g V)^T} & \text{if } L_g V \mathbf{u}_2 < 0 \quad \text{and} \quad \|\mathbf{u}_2\|_\infty \geq \epsilon \end{cases} \quad (17)$$

where ϵ is a positive constant representing a percentage of the maximum torque value and $\mathbf{u}_i = -k_{p_i} \mathbf{I} \mathbf{q} - k_{d_i} \mathbf{I} \boldsymbol{\omega}$, $i = 1, 2$. The subscript i is used to switch between two sets of gains and should not be confused with control contributions on different axes. Note that, for the case of rest-to-rest maneuvers, the switching regions could have been determined based on the initial conditions instead of the conditions on $\|\mathbf{u}_2\|_\infty$, but the control law of (17) is proposed as a more general solution, not limited to the case of satellite rest-to-rest maneuvers. This switched inverse optimal approach is compared in Section III-F.3 to the continuous inverse optimal controller of [11], which is outlined in Section III-E.

E. Continuous Inverse Optimal Control Approach used for Comparison

In this subsection we describe a widely cited continuous inverse optimal control approach to compare it with the proposed gain-scheduled formulation of the switched inverse optimal minimum norm controller.

In [11], the authors also consider an affine control system and the assumption is made that a static state feedback control law of the form $\mathbf{R}^{-1}(\mathbf{x}) L_g V^T(\mathbf{x})$, where \mathbf{R} is a positive definite matrix (of dimension 7×7 for our attitude control model) stabilizes the affine system given in (5) with respect to a Lyapunov function V , which is assumed to be radially unbounded to establish global stability.

Under these assumptions, it is known (see for example [20]) that the control law $\mathbf{u} = \beta k(\mathbf{x})$, $\beta > 2$ is inverse optimal with respect to a meaningful cost function, of the form $\int_0^\infty \{l(\mathbf{x}(t)) + \mathbf{u}^T(t) \mathbf{R} \mathbf{u}(t)\} dt$, where $l(\mathbf{x})$ is a positive function of the states with a complicated expression given in [11, eq. (5)]. In [11], a backstepping design approach was used to construct a Lyapunov function and the following inverse optimal $L_g V$ controller was proposed:

$$\mathbf{u} = -\mathbf{I} \left((2k_2 + k_1) \mathbf{1}_{3 \times 3} + k_1 \mathbf{p} \mathbf{p}^T + \frac{4}{k_1} \mathbf{I}^{-1} \mathbf{S}^T(\boldsymbol{\omega})^2 \mathbf{S}(\boldsymbol{\omega}) \mathbf{I}^{-1} \right) \times (\boldsymbol{\omega} + k_1 \mathbf{p}) \quad (18)$$

where $k_1 > 0$, $k_2 > 0$ are constant gains, $\mathbf{p} = \mathbf{q}/q_4$ stands for the Rodrigues-Gibbs parameters, used to describe attitude kinematics, and $\mathbf{S}(\boldsymbol{\omega})$ is the inverse of the cross product matrix, defined by $\mathbf{S}(\boldsymbol{\omega}) \mathbf{y} = \boldsymbol{\omega} \times \mathbf{y}$, $\forall \mathbf{y} \in \mathbb{R}^3$.

Note that the controller in (18) is analogous to a PD law with state varying gains, which are high initially and lowest as the trajectory approaches the origin. This design may first appear to be practically justified since aggressive controllers such as bang-bang control are generally used away from the origin in conjunction with feedback control in a neighborhood of the origin. However, this inverse optimal controller has a counterproductive effect since gains start to continuously decrease as soon as a maximum velocity is reached. Also, the use of low gains is only justified very close to the origin under uncertainties (on inertias, controller time delays) and noise (refer to [22, p. 198]) and does not improve performance for the uncertainty and disturbance-free case considered here and in [11]. This control mode can easily be added in a close neighborhood of the origin with the proposed gain-scheduled min norm controller, where the high gain mode can be seen as analogous to the deceleration phase of a bang-off-bang controller. The proposed gain-scheduled controller is shown in Section III-F.3 to compare favorably with the controller in (18). It is not ruled out that the approach used in [11] may lead to more efficient controllers, but the state varying gains would then have to be designed differently and still prove stability.

Remark on the meaningful cost of [17]: Having described the controller of [11], which was nonstandard, it should be noted that a simpler continuous inverse optimal control law, consisting of a PD quaternion feedback was proposed in [17]. That controller was designed to solve an HJB problem. The Lyapunov function was similar to our expression in (15), implying that $L_g V^T = \boldsymbol{\omega} + \gamma^* \mathbf{q}$. In other words, any PD law of the form $\mathbf{u} = -k(\boldsymbol{\omega} + \gamma^* \mathbf{q})$, where k is an arbitrary positive gain and γ^* is sufficiently small to define a Lyapunov function, is inverse optimal in the sense of [17]. The inverse optimal controller in [17] is therefore a PD law with mild performance constraints on the gains, despite minimizing a mathematically meaningful cost function. This illustrates that inverse optimality does not always guarantee high performance. In this paper the use of minimum norm optimization introduces extra constraints and objectives that rule out such standard controllers. We also consider performance metrics (settling time, integrated torque, maximum torque limitations) for controller tuning.

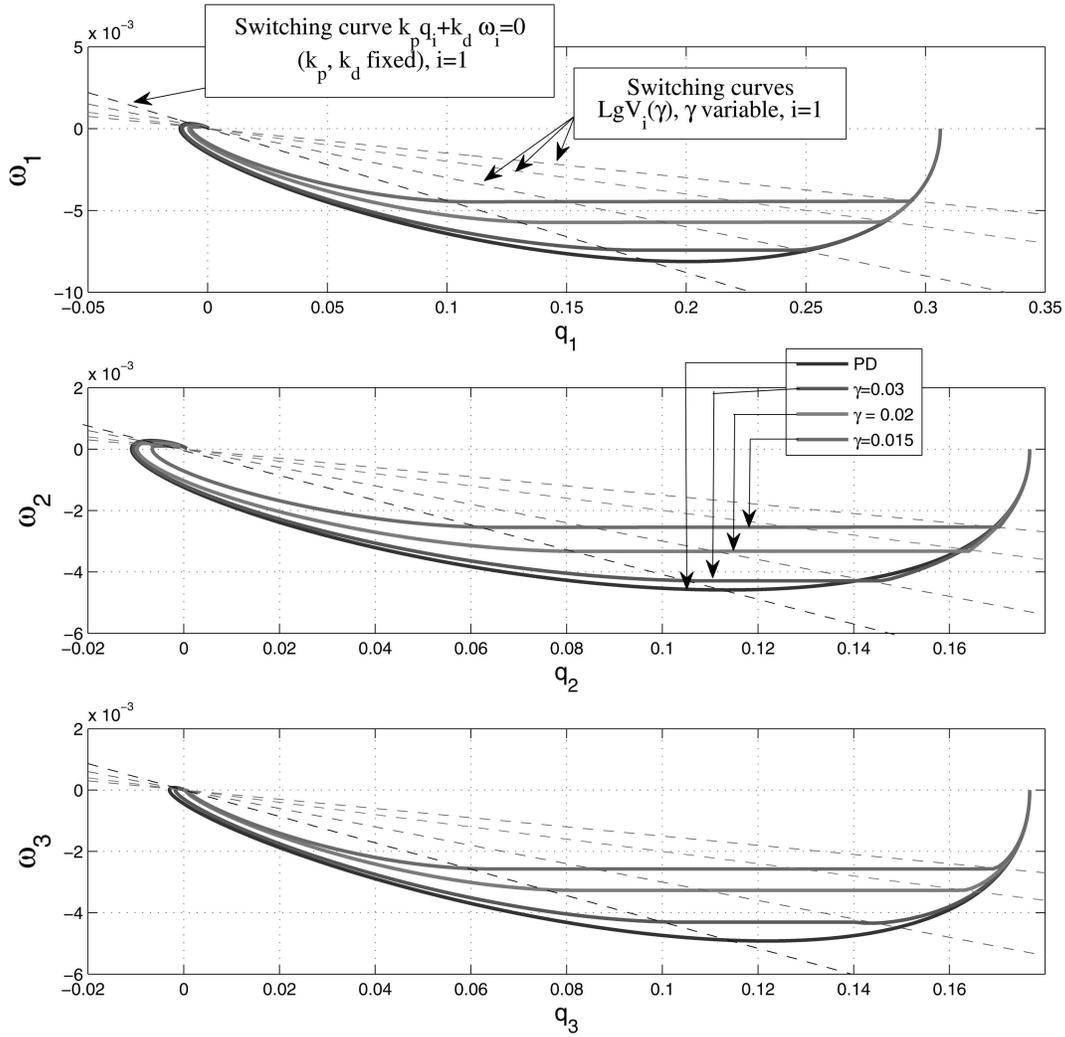


Fig. 1. Phase portrait of minimum norm controller for different values of γ .

F. Numerical Simulations of Nonlinear Satellite Attitude Control

The system parameters used for the simulations are

$$I_1 = 12 \text{ kg m}^2, \quad I_2 = 14 \text{ kg m}^2, \quad I_3 = 10 \text{ kg m}^2$$

which correspond to a slightly asymmetric microsatellite.

A realistic torque saturation level for this microsatellite is 0.01 Nm. Unless stated otherwise, the initial conditions are:

$$\begin{aligned} \omega_i(0) &= 0, & i &= 1, 3 \\ q_1(0) &= 0.3062, & q_2(0) &= 0.1768 \\ q_3(0) &= 0.1768, & q_4(0) &= 0.9186. \end{aligned}$$

This corresponds to initial attitude errors of 30 deg on all three axes for a 2-1-3 Euler rotation sequence.

The controller gains are: $k_p = k_{p1} = 0.002$, $k_d = k_{d1} = 0.05$, γ variable.

1) *Standard Minimum Norm Controller*: The phase portrait of the system described by (3) and (4), with the control law (14) is presented in Fig. 1 for

γ variable and k_p, k_d fixed. The Lyapunov function in (15) was used. We observe that the control torque is turned off between two straight lines in the phase space domain, determined by the Lyapunov function and the benchmark controller.

By decreasing the values of the parameter γ , the system proceeds more directly (on a shorter path) on all three axes towards the origin in the phase space domain and the torque expenditure is reduced as a result of a longer zero torque phase and lower velocities. The attitude response can, however, be shown to be slightly slower for the lower values of γ that generate lower velocities. The overall tradeoff is, however, still improved (see [8] for typical plots of the attitude response and torque profile).

Compared with the PD benchmark controller, the system proceeds more directly towards the origin in phase space with min-norm control and energy can be saved, while settling time is not significantly deteriorated. We chose to plot the phase space trajectory rather than the attitude and torque response because it better illustrates system behavior but

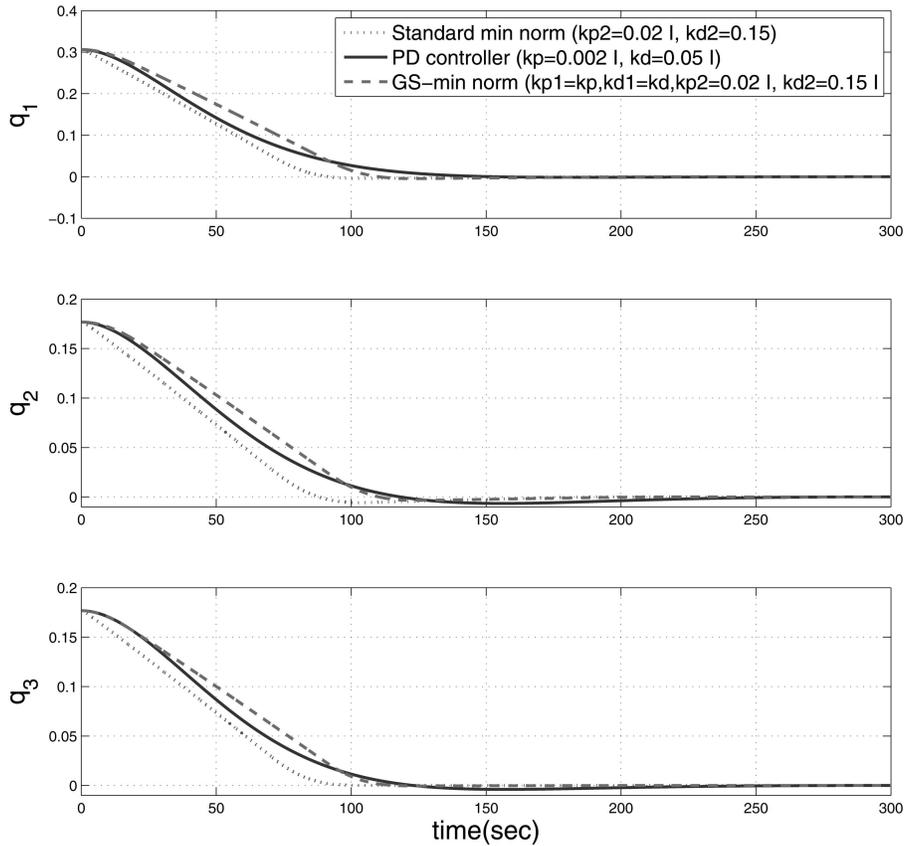


Fig. 2. Attitude response with high gain min norm controller (dotted line), gain-scheduled controller (dashed line), and with PD controller (line).

we found that, with $\gamma = 0.02$, the average torque expenditure is divided by 2 while the 2% settling time increases by less than 50 s for a 220 s maneuver. In [1] both settling time and torque were improved, but only because the gains k_p and k_d were inappropriately selected as the system was excessively underdamped. For a fairer and more realistic comparison, we have chosen the gains of the PD controller to achieve a 6% single overshoot. This is a slightly underdamped response, although similar comparison has also been made for a critically damped response with even better results. Minimum norm controllers would achieve an even higher tradeoff improvement if the PD gains happened to be excessively underdamped or overdamped. In Section III-F.2, we demonstrate how tradeoff enhancement becomes a reality by a gain-scheduled formulation of minimum-norm control.

2) *Tradeoff Improvement of a PD Benchmark Controller by Standard and Gain-Scheduled Minimum Norm Control:* In this subsection, we compare the performance of the control laws (14) and (17), for the same torque expenditure and the same attitude maneuver, with the benchmark PD law given by

$$\mathbf{u}_1 = -k_{p1} \mathbf{I} \mathbf{q} - k_{d1} \mathbf{I} \dot{\mathbf{q}}.$$

Practical issue with standard min-norm control: To conduct a fair comparison, the PD gains of (14) are taken to be larger than those of \mathbf{u}_1 : $k_p = 0.02 > k_{p1}$,

$k_d = 0.15 > k_{d1}$. These were tuned to consume the same overall integrated torque as the PD benchmark controller \mathbf{u}_1 . Settling times are then compared. This is done by increasing the gains of the standard min-norm controller until the overall torque reaches the same level as the PD law. Figure 2 shows a settling time enhancement by standard min-norm control. However, Fig. 3 shows that the standard minimum norm controller with higher gains causes a higher initial torque value, potentially causing torque saturation.

Performance enhancement by gain-scheduled min-norm control: The gain-scheduled controller in (17) used for comparison has the following gains: $k_{p1} = 0.002$, $k_{d1} = 0.05$, $k_{p2} = 0.02$, $k_{d2} = 0.15$, $\gamma = 0.02$. PD gains are multiplied by the moment of inertia \mathbf{I} .

Settling time enhancement: The gain-scheduled control law of (17) is shown to also reduce settling time compared with the PD benchmark (see Fig. 2), while avoiding this saturation issue (see Fig. 3). The settling time (to 0.5 deg accuracy on all 3 axes) of the gain-scheduled minimum norm controller is 123.8 s compared with 228.2 s with the PD benchmark. *Integrated torque comparison:* the overall torque expenditure of the gain-scheduled min-norm controller (0.214) was in fact tuned to be smaller than the PD benchmark (0.245).

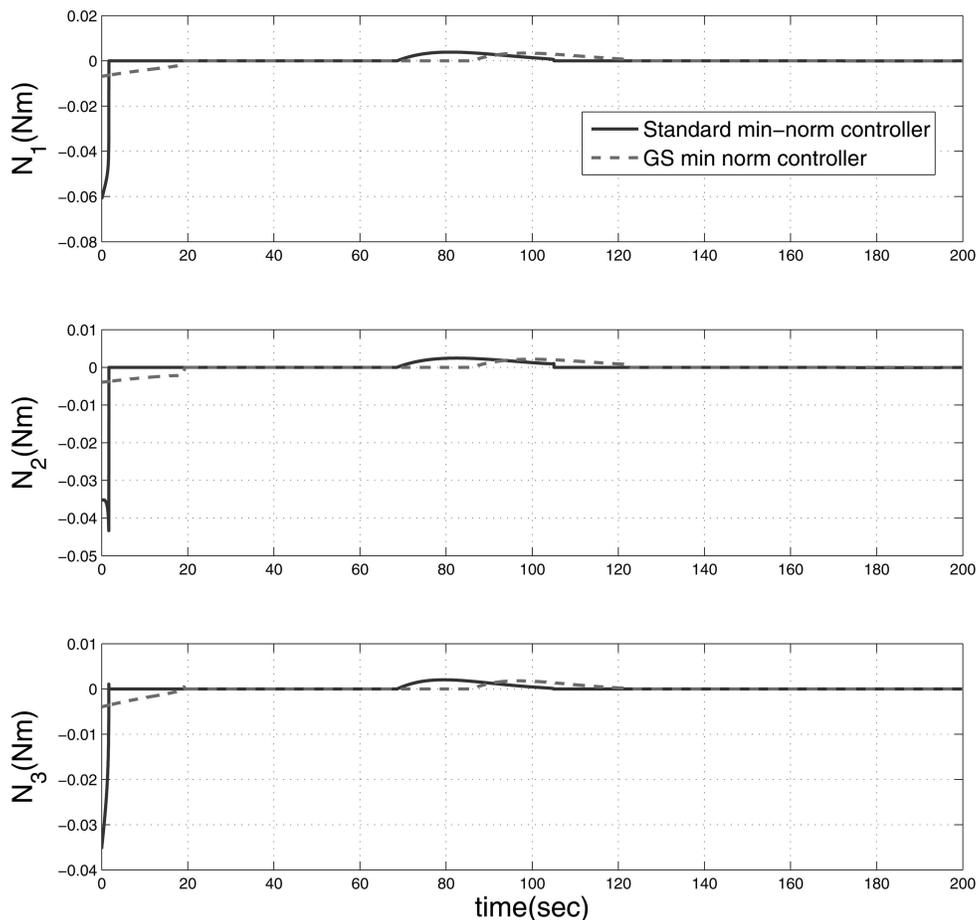


Fig. 3. Control torque of minimum norm controller, with and without gain scheduling.

The tradeoff between slew time and integrated torque is therefore enhanced by gain-scheduled min-norm control compared with the PD benchmark control law. Crucially, this enhancement is obtained with admissible instantaneous torque. Gain scheduling circumvents the torque saturation issues of standard min-norm control.

3) Comparison Between the Gain-Scheduled and Continuous Inverse Optimal Approach of Reference [11]: The continuous inverse optimal $L_g V$ controller of [11] and given in (18) of our paper is considered for comparison purposes. Note that for the purpose of this comparison only, we adopt exactly the same simulation parameters of [11], which are $I_1 = 10 \text{ kg m}^2$, $I_2 = 15 \text{ kg m}^2$, $I_3 = 20 \text{ kg m}^2$ which corresponds to an asymmetric microsatellite, $\omega_i(0) = 0$, $i = 1, 2, 3$, $q_1(0) = 0.4646$, $q_2(0) = 0.1928$, $q_3(0) = 0.8047$, $q_4(0) = 0.3153$, which is obtained by converting the initial conditions in [11] from Rodrigues parameters to quaternions. The controller gains in (18) are $k_1 = 0.5$, $k_2 = 0.1$ (taken from [11]). We have also checked that we obtained exactly the same simulation results as in [11]. (This might be visually difficult to see on the control torque figure, where a simulation time of only 1 s was used and the torque appeared stationary but was in fact still above

zero in [11], while we use a 10 s simulation time, which is more appropriate for the comparison.)

We assume a much higher torque saturation level (up to 140 Nm on any axis) than the 0.01 Nm under consideration elsewhere in this paper despite the fact that [11] also considers the parameters of a small satellite. This level is considered here to be oversized for small satellite actuators, but the conditions of [11] are used here for comparison purposes only.

The controller of [11] was expectedly found to produce nearly identical results, whether the control torques are external as in [11] or from reaction wheels. The results shown on the simulation figures are for reaction wheels for comparison purposes. Our gain-scheduled min norm controller has to be designed with internal torques from reaction wheels in order to maintain total momentum conservation during the zero torque mode and produce a horizontal path in the phase space during the zero torque mode of a rest-to-rest maneuver. The parameters of the gain-scheduled minimum norm controller used for comparison are $\mathbf{k}_{p1} = 5\mathbf{I}$, $\mathbf{k}_{d1} = 5\mathbf{I}$, $\mathbf{k}_{p2} = 30\mathbf{I}$, $\mathbf{k}_{d1} = 10\mathbf{I}$, $\gamma = 0.8$.

Figure 4 shows the control torque generated with both controllers. The gain-scheduled min-norm controller was tuned to deliver a lower maximum

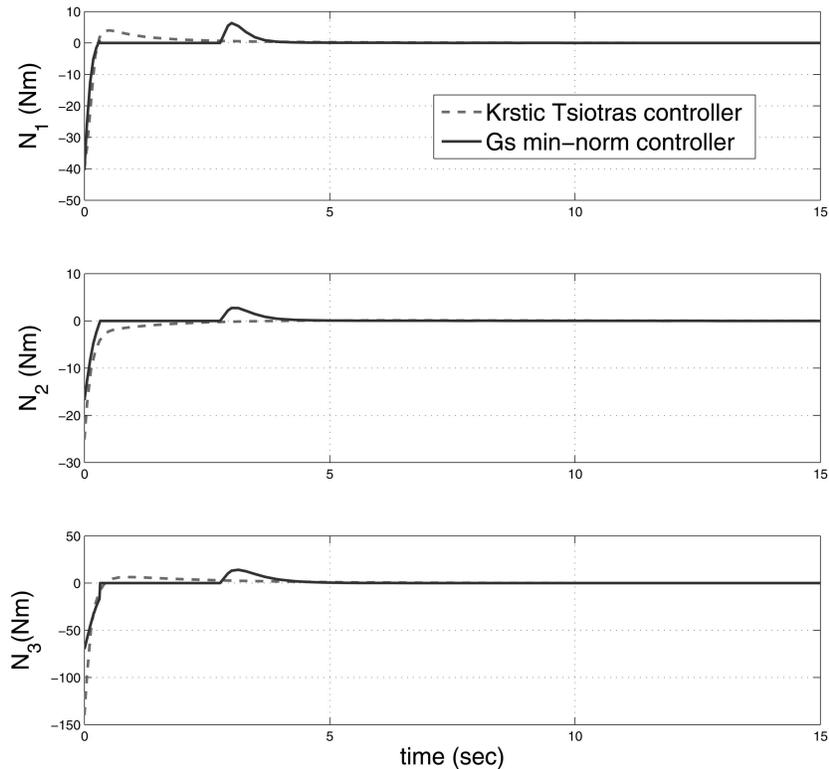


Fig. 4. Control torque comparison between gain-scheduled and the Krstic/Tsiotras controllers.

instantaneous torque than the controller of [11]. The gain-scheduled controller was tuned to consume less integrated torque (26.95 Nms) (and also less energy) than the Krstic/Tsiotras controller (35.66 Nms), which is feasible thanks to the zero torque mode. In other words, the proposed controller in (17) was allowed a smaller saturation torque level and a smaller total integrated torque. Despite that, Fig. 5 shows that the settling time to ± 1 deg and ± 0.5 deg (after conversion to a 2-1-3 sequence of Euler angles) is significantly enhanced on all three axes by the gain-scheduled min-norm control compared with the controller of [11]. The comparative results are summarized in Table I.

Why is performance enhanced? The performance enhancement is first explained by the fact that the gain-scheduled min-norm controller shows a greater similarity to the minimum time-fuel bang-off-bang solution. A more physical explanation can be obtained by comparing the phase space trajectories of the two controllers (see Fig. 6). The phase space trajectory of the continuous control law of [11] shows that angular velocity peaks relatively early, then the gains are decreased continuously by the controller of (18) from that point in time where both attitude and angular velocity decrease (in absolute value). More importantly, the controller of [11] consumes torque to slow down the response, as a linear state feedback law would do, while the proposed gain-scheduled controller proceeds towards the origin with maximum slew rate until a higher gain is activated when there

are no torque saturation issues. This phase space trajectory is therefore more effective at reducing settling time for a given level of the integrated torque. The values of the higher set of gains are shown in the next section to be limited by maximum torque constraints. Increasing the values of these gains is shown for a simplified model not to necessarily increase the total integrated torque. Indeed, increasing these gains, which define a switching curve, delays the reactivation of the controller and makes the zero torque duration increase.

The example shown in Figs. 4–6 is representative of the type of performance enhancement that is achievable and the phase space trajectory comparison in Fig. 6 is typical. For any rest-to-rest maneuver, a properly tuned Krstic/Tsiotras controller would generate an acceleration phase followed by a deceleration phase (qualitatively similar to Fig. 6) and the gain-scheduled controller can then be scaled to shape the phase space trajectory more effectively. Despite the fact that Figs. 4–6 can be seen as evidence of tradeoff enhancement compared fairly with the Krstic/Tsiotras law, similar comparisons have also been made at different levels of the integrated torque expenditure with both controllers and these are summarized in Fig. 7 for the same maneuver. For each level of the integrated torque, a procedure was applied to tune both controllers to approach their fastest response (by considering the overlap between the contours of equal torque and equal time).

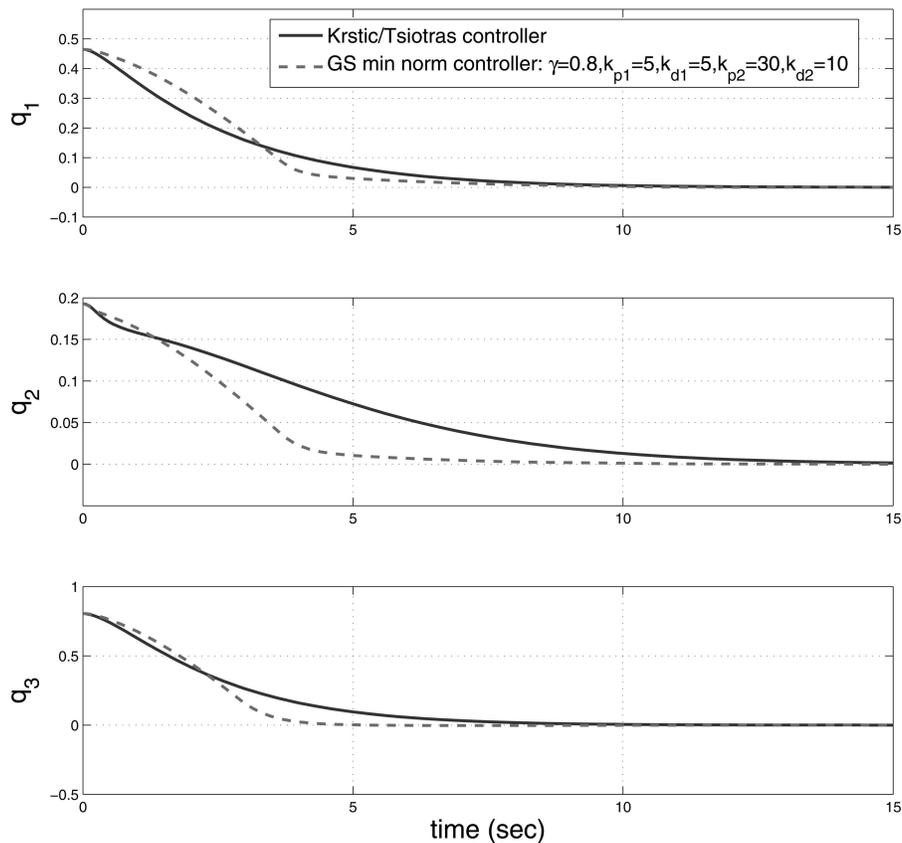


Fig. 5. Attitude response comparison between gain-scheduled and the Krstic/Tsiotras controllers.

The procedure was simplest in the case of the Krstic/Tsiotras controller that only has two gains to tune and the fastest response is approached. With the proposed gain-scheduled law, the data were partly generated by fixing two gains and adjusting the four others by trial and error and the best time was stored, which is not guaranteed to be close to the minimum, but the main objective was to show that it can be made smaller than the settling time of the Krstic/Tsiotras controller. For simplicity, maximum torque limitation was not directly imposed to the Krstic/Tsiotras controller, but the gain-scheduled controller was always tuned not to exceed the maximum torque generated during the maneuver by the Krstic/Tsiotras law. Figure 7 shows that settling time is enhanced for a given level of the integrated torque and that the percentage of the time saved increases with the integrated torque. Note that smaller rotations were also considered and significant settling time enhancement was also observed in those cases.

So far, we have shown the merits of the proposed inverse optimal controller in terms of settling time performance under integrated and saturation torque constraints by comparison with the continuous inverse optimal approach in [11]. We have also demonstrated in Section III-F.2 how the proposed gain-scheduling scheme achieves settling time

TABLE I
Controller Performance Comparison Between the Switched and Continuous Inverse Optimal Controllers

		Krstic/Tsiotras	GS min-norm
Total Integrated Torque (Nms)		35.66	26.95
Maximum torque during slew (Nm)	X axis	-40.4	-40.4
	Y axis	-25.13	-16.77
	Z axis	-139	-69.98
Settling time X axis (s)	± 1 deg accuracy	10.6	9.6
	± 0.5 deg accuracy	12.9	10.6
Settling time Y axis (s)	± 1 deg accuracy	12.6	7.1
	± 0.5 deg accuracy	15.8	10.1
Settling time Z axis (s)	± 1 deg accuracy	10.4	4.8
	± 0.5 deg accuracy	12.3	10.2

enhancement over a PD benchmark controller, for similar overall torque expenditure but without the risk of torque saturation. The possibility of tradeoff enhancement by gain-scheduled minimum norm optimization has therefore been shown for this nonlinear attitude control problem. However, the extent to which performance can be enhanced remains unknown. In the following section, we consider the problem of tuning this controller for simpler attitude

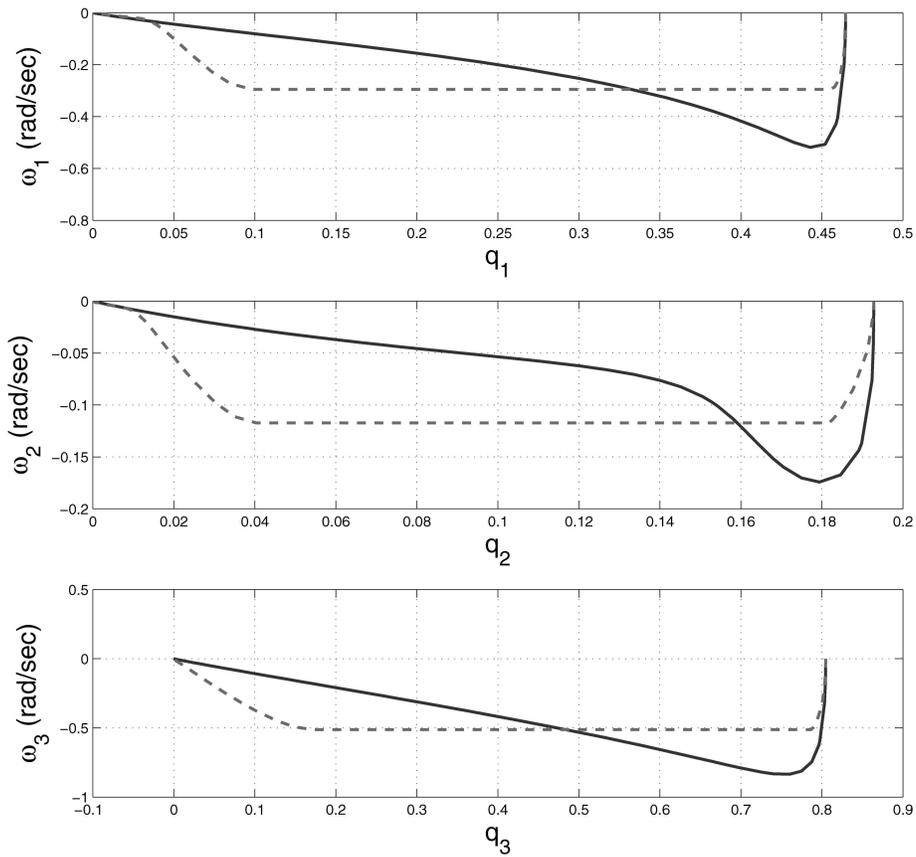


Fig. 6. Phase space trajectory with gain-scheduled and Krstic/Tsiotras controllers.

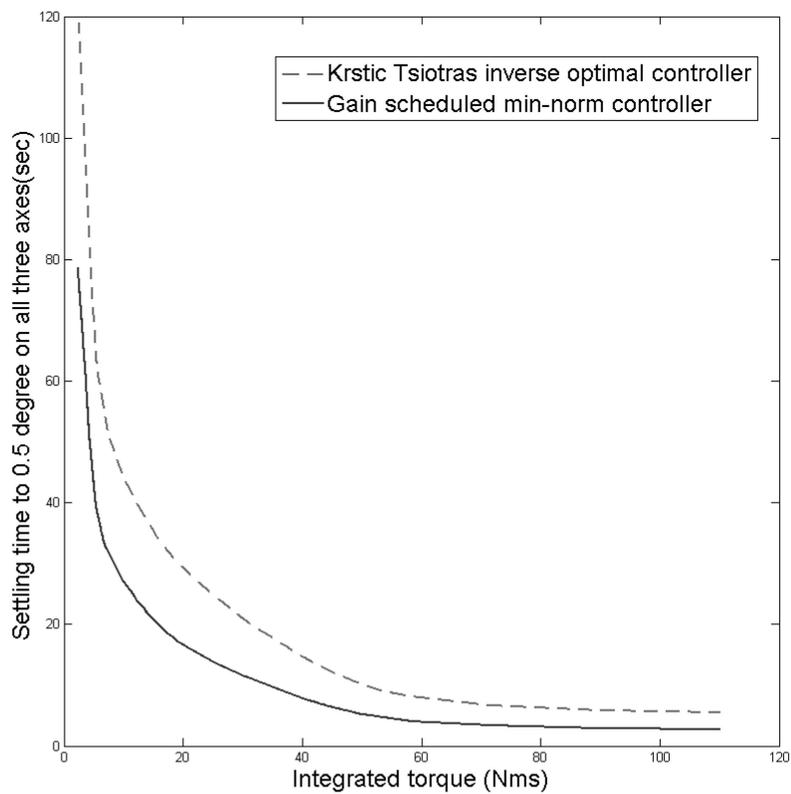


Fig. 7. Settling time for various levels of integrated torque.

dynamics based on the eigenaxis rotation case. The controller tuning analysis will determine the extent to which control performance can be enhanced given constraints on maximum and total integrated torque.

IV. PRACTICAL TUNING OF THE SWITCHED CONTROLLER

The problem of tuning the parameters of the proposed gain-scheduled switched controller of (17) is addressed here for a simplified attitude control model, based on the concept of eigenaxis rotations. We also analyze the extent to which settling time can be enhanced by varying controller parameters for a given level of the total integrated torque, and verify that the enhancement is achievable without exceeding torque saturation constraints. In other words, we relate controller parameters to settling time, within a set of constraints on the integrated torque and maximum instantaneous torque.

A. Eigenaxis Rotation

In [6] it was shown by Wie, et al., starting from the model of (3) and (4), that three axis attitude maneuvers, under the effect of PD-type feedback (with gyroscopic torque compensation):

$$\mathbf{u} = -\dot{\mathbf{h}} - \boldsymbol{\omega} \times \mathbf{h} = \boldsymbol{\omega} \times \mathbf{I}\boldsymbol{\omega} - \mathbf{k}_p \mathbf{q} - \mathbf{k}_d \boldsymbol{\omega} \quad (19)$$

where \mathbf{k}_p , \mathbf{k}_d are positive definite matrices, can also be modeled as a single rotation about an eigenaxis \mathbf{e} , such that $\mathbf{q} = \sin(\theta/2)\mathbf{e}$, $\boldsymbol{\omega} = \dot{\theta}\mathbf{e}$.

Bong Wie, et al. have shown that the satellite dynamics can then be modeled as a damped oscillator equation:

$$\mathbf{I}\ddot{\theta}\mathbf{e} = \mathbf{u} = \left(-\mathbf{k}_p \sin\left(\frac{\theta}{2}\right) - \mathbf{k}_d \dot{\theta} \right) \mathbf{e} \quad (20)$$

where \mathbf{k}_p , \mathbf{k}_d are generally diagonal gain matrices and \mathbf{I} is a diagonal moment of inertia matrix. By choosing $\mathbf{k}_p = K_p \mathbf{I}$, $\mathbf{k}_d = K_d \mathbf{I}$, the problem translates to a single nonlinear damped oscillator equation. In the following, we consider that $\theta/2$ is small enough (extreme limit at 15 deg) for a linear approximation of the sine term.

B. Switched Control for an Eigenaxis Rotation

To analyze our controller tuning problem, let's consider for simplicity that the eigenaxis is one of the principal axes of the satellite and that the moment of inertia of that axis is denoted J (making gyroscopic torque compensation unnecessary, which would also be the case for a nearly symmetric spacecraft or under small momentum conditions). For a diagonal moment of inertia matrix and a gain-scheduled minimum norm controller that switches between two PD control modes with the closed-loop dynamics of (20), the closed-loop attitude dynamics on a single axis are

given by

$$J\ddot{\theta} = u \quad (21)$$

where the moment of inertia J is a scalar and

$$u = \begin{cases} u_1 & \text{if } \dot{\theta}(\dot{\theta} + \gamma\theta) < 0 \\ 0 & \text{if } (\dot{\theta} + \gamma\theta) \left(\dot{\theta} + \frac{k_{p2}}{2k_{d2}}\theta \right) \leq 0 \\ u_2 & \text{elsewhere} \end{cases} \quad (22)$$

where $u_i = (-k_{pi}/2)\dot{\theta} - k_{di}\theta$, $i = 1, 2$. The index $i = 1$ is for the low gain mode and $i = 2$ is for the high gain mode. The switching curves define mutually exclusive areas in phase space. For rest-to-rest maneuvers, the controller is turned off at the intersection of the phase space trajectory with the line $\dot{\theta} + \gamma\theta = 0$ and reactivated after intersection with the line $\dot{\theta} + (k_{p2}/2k_{d2})\theta = 0$.

To simplify the controller tuning analysis, we assumed a choice of the parameter ϵ (determining transition from lower to higher gains in (17)) such that lower gains are used before turning the controller off and only higher gains are used afterwards. This is always possible for rest-to-rest maneuvers by choosing ϵ within an interval after selecting k_{p1}, k_{d1} . The parameter ϵ would generate a feasible torque whenever k_{p1}, k_{d1} are tuned to generate admissible torque. We still assume that the parameters k_{p1}, k_{d1} are those of the benchmark controller to be enhanced by gain-scheduled min-norm optimization.

The following analysis shows how the tuning of the controller parameters γ, k_{p2}, k_{d2} affects control performance. By applying the control law of (22) to (21), we have the following closed-loop dynamics:

$$\begin{cases} \ddot{\theta} + 2\zeta\omega_{n1}\dot{\theta} + \omega_{n1}^2\theta = 0 & \text{if } \dot{\theta}(\dot{\theta} + \gamma\theta) < 0 \\ \ddot{\theta} = 0 & \text{if } (\dot{\theta} + \gamma\theta) \left(\dot{\theta} + \frac{k_{p2}}{2k_{d2}}\theta \right) \leq 0 \\ \ddot{\theta} + 2\zeta\omega_{n2}\dot{\theta} + \omega_{n2}^2\theta = 0 & \text{elsewhere} \end{cases} \quad (23)$$

where $k_{pi}/2J = \omega_{ni}^2$, $k_{di}/J = 2\zeta\omega_{ni} = 2\zeta\sqrt{k_{pi}/2J}$, $i = 1, 2$.

In other words, the gains k_{p1}, k_{d1} are increased to $k_{p2} \geq k_{p1}$, $k_{d2} \geq k_{d1}$, while maintaining a constant damping ratio ζ , but with different natural frequencies ω_{n1}, ω_{n2} .

For the case $\lambda < 1$, which corresponds to an underdamped system, the analytical solution of the closed-loop dynamics is given by

$$\theta(t) = \begin{cases} A_1 \cos(\Omega_1 t - \phi_1) e^{-\lambda_1 t} & \text{if } \dot{\theta}(\dot{\theta} + \gamma\theta) < 0 \\ \theta(t_s) + \dot{\theta}(t_s)(t - t_s) & \text{if } (\dot{\theta} + \gamma\theta) \left(\dot{\theta} + \frac{k_{p2}}{2k_{d2}}\theta \right) \leq 0 \\ A_2 \cos(\Omega_2 t' - \phi_2) e^{-\lambda_2 t'} & \text{elsewhere} \end{cases} \quad (24)$$

with

$$\lambda_i = \zeta\omega_{ni}, \quad \Omega_i = \sqrt{\omega_{ni}^2 - \lambda_i^2} = \omega_{ni}\sqrt{1 - \zeta^2} = \sqrt{\frac{k_{pi}}{2J}(1 - \zeta^2)}$$

$$\phi_i = \arctan\left(\frac{\theta_{0i}\Omega_i}{\dot{\theta}_{0i} + \lambda_i\theta_{0i}}\right), \quad A_i^2 = \theta_{0i}^2 + \frac{\dot{\theta}_{0i} + \frac{k_{pi}}{2J}\theta_{0i}}{\Omega_i^2}, \quad i=1,2.$$

The switching time

$$t_s = \frac{\arctan\left(\frac{\gamma - \lambda_1}{\Omega_1}\right) + \phi_1}{\Omega_1}$$

is the time of intersection between the line $\dot{\theta} + \gamma\theta = 0$ and the low gain mode. Likewise,

$$t'_s = t_s + \frac{\theta'_s - \theta_s}{\dot{\theta}_s}$$

$$= t_s - \frac{2\left(k_{p2}\frac{\theta_s}{2} + k_{d2}\dot{\theta}_s\right)}{k_{p2}\dot{\theta}_s} > t_s \quad \text{for} \quad \frac{k_{p2}}{2k_{d2}} > \gamma$$

is the time when higher gains are activated and $t' = t - t'_s$.

Also note that the slope of the second switching curve where the high gain mode is activated increases by increasing the gains:

$$\frac{k_{p2}}{2k_{d2}} = \frac{\sqrt{\frac{k_{p2}}{2J}}}{4\zeta}. \quad (25)$$

C. Effect of Increasing k_{p2} , k_{d2} on the Settling Time (with Constant Damping Ratio)

The closed-loop dynamics are given by (23). During the first phase before the zero torque mode, the gains k_{p1} , k_{d1} (of the benchmark controller to be enhanced) are used. After reactivating the controller with the gains k_{p2} , k_{d2} , the condition is that $k_{p2} > k_{p1}$, $k_{d2} > k_{d1}$. By increasing the gains in this phase from $k_{p2} = k_{p1}$, $k_{d2} = k_{d1}$ (standard min-norm) to $k_{p2} > k_{p1}$, $k_{d2} > k_{d1}$ (gain-scheduled min-norm) the effect on settling time is:

1) The time spent in the zero torque phase increases by increasing the gains of the high gain phase from k_{p1} to k_{p2} (k_{d2} is uniquely determined from (25):

$$\Delta t_1 = t_{k_{p2}} - t_{k_{p1}} = \frac{\Delta\theta_2 - \Delta\theta_1}{\dot{\theta}_s} \quad (26)$$

with

$$\Delta\theta_1 = \frac{\dot{\theta}_s}{\gamma} - \frac{\dot{\theta}_s}{\frac{k_{p1}}{2k_{d1}}}, \quad \Delta\theta_2 = \frac{\dot{\theta}_s}{\gamma} - \frac{\dot{\theta}_s}{\frac{k_{p2}}{2k_{d2}}}$$

(where $\dot{\theta}_s$ is the constant attitude rate of the zero torque mode).

Equation (26) simplifies as follows (by applying the formula of (25) relating proportional and derivative gains):

$$\Delta t_1 = 8\zeta^2 J \left(\frac{1}{k_{d1}} - \frac{1}{k_{d2}} \right) > 0. \quad (27)$$

2) The time spent after reactivating the controller (until $\alpha\%$ error tolerance is reached) with higher gains decreases:

$$\Delta t_2 = \frac{-\ln(\alpha\%)}{\lambda_2 k_{d2}} + \frac{\ln(\alpha\%)}{\lambda_2 k_{d1}} \quad (28)$$

$$\Delta t_2 = 2J \ln(\alpha\%) \left(\frac{k_{d2} - k_{d1}}{k_{d1}k_{d2}} \right) < 0.$$

This is obtained by applying the $\alpha\%$ settling time law: $\alpha\% = e^{-\lambda_{set}}$, with $\ln(\alpha\%) < 0$, for $\alpha < 1$. Note that $\alpha\% = |\theta_0/\theta'_s|\epsilon\%$. It is indeed necessary to scale the error tolerance from a percentage of the total maneuver (usually $\epsilon\% = 2\%$) to a percentage of the maneuver from θ'_s , when the higher gains are activated.

The combined effect of increasing the gain on the settling time can be quantified as

$$\Delta t = \Delta t_1 + \Delta t_2 = (2J \ln(\alpha\%) + 8\zeta^2 J) \left(\frac{k_{d2} - k_{d1}}{k_{d1}k_{d2}} \right). \quad (29)$$

For α sufficiently small, the settling time decreases by increasing the gains. Therefore, for certain system parameters, the settling time keeps decreasing by increasing the gains k_{p2} , k_{d2} (which are respectively higher than k_{p1} , k_{d1}).

In particular, for an ‘‘optimal’’ damping ratio:

$$\zeta = \frac{\sqrt{2}}{2} \rightarrow \Delta t = (2 \ln(\alpha\%) + 4)J \left(\frac{k_{d2} - k_{d1}}{k_{d1}k_{d2}} \right). \quad (30)$$

In this case, with $\ln(\alpha\%) < -2 \Leftrightarrow \alpha\% < 13.6\%$ (of the attitude at time t'_s), the settling time would keep decreasing $\Delta t < 0$ by increasing the gains. Hence, settling times strictly decrease by increasing the gains for reasonable error tolerances. There is however a limit on the level to which the settling time can be reduced by increasing the gains k_{p2} , k_{d2} after reactivating the controller. As shown in Section IV-D, this limit is dependent on the parameter γ .

D. Tuning the Parameter γ using the Concept of Settling Time Limit

It is possible to uniquely select a value of γ to meet a specific settling time requirement. Indeed, when k_{p2} , k_{d2} go to infinity, there is a settling time limit for the controller that is only dependent on γ . This settling time limit is given by

$$t_{sl} = t_s + \left| \frac{\theta_s}{\dot{\theta}_s} \right| + 0. \quad (31)$$

The first term of the sum in (31) represents the time to reach the first switching curve. This term depends on the prescribed k_{p1}, k_{d1} only. The second term represents the time during which the controller is switched off. The expression of the second term stems from the fact that the second switching curve would have infinite slope with infinite gains and would correspond to the $\theta = 0$ axis. The third zero term represents the time spent with infinite gains. The effect of the gain tuning on the settling time and the meaning of the settling time limit are demonstrated in the numerical analysis Section IV-G.

E. Peak Torque and Total Integrated Torque

With the gain-scheduled minimum norm controller, a peak (or series of peaks of decreasing magnitude in the underdamped case) is produced when the controller is reactivated. The peak time t_M and peak torque u_{peak} are obtained from the condition $\dot{u}(t_M) = 0$ and after lengthy but straightforward calculations, we obtain

$$t_M = t'_{u=0} = \frac{\arctan\left(\frac{\zeta}{\sqrt{1-\zeta^2}}\right) + \phi_2}{\Omega_2} \quad (32)$$

$$u_{\text{peak}} = \frac{-\sqrt{1-\zeta^2}}{4\zeta^2} A_2 \frac{k_{d2}^2}{J} e^{-\lambda_2 t_M}. \quad (33)$$

The effect of changing gains on the peak torque is illustrated in Section IV-G for a fixed damping ratio and a range of values of the gains. The gains k_{p2}, k_{d2} are increased until the peak torque reaches the initial torque, which also represents the maximum torque value of the benchmark controller with gains k_{p1}, k_{d1} .

The total integrated torque is made up of two contributions: the torque consumed in the low gain mode (constant when γ, k_{p1}, k_{d1} are determined) and the torque consumed during the high gain mode. It has the following expression:

$$\int_0^t u(\tau) d\tau = \left[\left(\frac{D_1 \Omega_1}{\lambda_1^2 + \Omega_1^2} - \frac{D_2 \lambda_1}{\lambda_1^2 + \Omega_1^2} \right) \sin(\Omega_1 t + \phi_1) + \left(\frac{D_1 \lambda_1}{\lambda_1^2 + \Omega_1^2} - \frac{D_2 \Omega_1}{\lambda_1^2 + \Omega_1^2} \right) \sin(\Omega_1 t + \phi_1) \right] e^{-\lambda_1 t} \Big|_0^t \\ + \left[\left(\frac{E_1 \Omega_2}{\lambda_2^2 + \Omega_2^2} - \frac{E_2 \lambda_2}{\lambda_2^2 + \Omega_2^2} \right) \sin(\Omega_2 t + \phi_2) + \left(\frac{E_1 \lambda_2}{\lambda_2^2 + \Omega_2^2} - \frac{E_2 \Omega_2}{\lambda_2^2 + \Omega_2^2} \right) \sin(\Omega_2 t + \phi_2) \right] e^{-\lambda_2 t} \Big|_{t_s}^t \quad (34)$$

with $D_1 = k_{d1} \lambda_1 A_1 - k_{p1} A_1$, $D_2 = k_{d1} \Omega_1 A_1$, $E_1 = k_{d2} \lambda_2 A_2 - k_{p2} A_2$, $E_2 = k_{d2} \Omega_2 A_2$ and all other symbols were defined after (24).

The total integrated torque $\int_0^\infty |u(\tau)| d\tau$ does not constantly increase by increasing the gains. Indeed, increasing the gains k_{p2}, k_{d2} has two competing effects on the overall torque:

1) From (25), the slope of the switching curve $(k_{p2}/2)\mathbf{q} + k_{d2}\boldsymbol{\omega} = 0$ increases, meaning that the torque is turned off for a longer time (increases).

2) The control torque increases when the higher gains are activated on a shorter time interval.

The comparative effect of these two competing terms on the total integrated torque was found to be difficult to analyze mathematically for a PD benchmark. That is why a numerical analysis is used to show the effect of the gains (in a large range) on the integrated torque. In Section IV-G, the total integrated torque is shown to be more dependent on the parameter than the PD gains (k_{p2}, k_{d2}). Each value of γ therefore maps to a narrow band of total integrated torque values but also, as shown in the following, to a settling time limit given in (31).

F. Tuning Procedure for Single Axis Rest-to-Rest Maneuvers

We propose the following controller tuning procedure to approach the settling time limit with the controller of (22).

1) The gains k_{p1}, k_{d1} of the benchmark controller are typically chosen to deliver a maximum torque, which is lower by a certain margin (30% here), than the saturation torque. In practice, the gains k_{p1}, k_{d1} would be those already stored onboard the satellite for the normal operation mode based on PD control.

2) Tuning γ : The ‘‘settling time limit’’ defined in (31) can be used as a design parameter. A settling time limit objective is specified (with a margin from the desired settling time). The parameter is determined from this performance objective because of nondependency of (31) on other gains.

3) Tuning the set of higher gains (k_{p2}, k_{d2}): The gains (initialized at the values k_{p1}, k_{d1}) are then positively incremented to approach the desired settling time limit, obviously not to infinity, but until the ‘‘maximum design torque’’ value is reached. In practice, the settling time limit can be approached closely because these gains can be much higher than

the initial gains. The maximum design torque is set equal to the initial torque of the PD law with gains k_{p1}, k_{d1} . That ensures that the maximum torque during a maneuver with the benchmark controller is never exceeded by the controller in (22).

G. Numerical Simulation Analysis: Tuning the Gain-Scheduled Switched Controller

In this subsection, the problem of tuning of the gain-scheduled minimum norm controller is illustrated

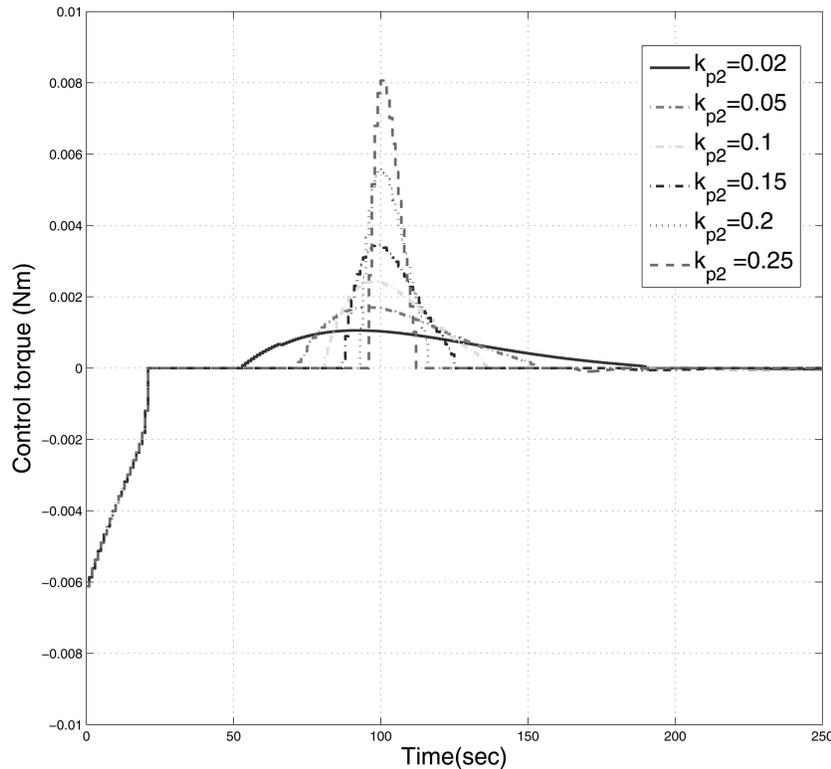


Fig. 8. Control torque of gain-scheduled minimum norm controller by varying gains of high gain mode.

by numerical simulations for the simple case of an eigenaxis rotation about a principal axis as modeled in equations (21), (22). The amplitude of the single axis rotation is 30 deg, which is considered here as the limit when half of this angle can be treated as sufficiently small in (20). The moment of inertia of the rotation axis is $J = 10$ (typical of a microsatellite having less than 100 kg mass). A realistic torque saturation limit is 10 milli Nm.

The effect of increasing the gains k_{p2} , k_{d2} in (22) is analyzed while maintaining the same damping ratio as the gains k_{p1} , k_{d1} , which is done by tuning the gains according to the rule $k_{di}/J = 2\zeta\sqrt{k_{pi}/2J}$, $i = 1, 2$.

1) Tuning the PD Gains (Assuming γ Fixed):

In Fig. 8, the effect of increasing the gains on the maximum instantaneous value of the torque is shown to correspond to a later and higher impulse. For γ fixed, a controller design condition is that the gains k_{p2} , k_{d2} in (22) are increased until the limit when the second peak of Fig. 8 reaches, in absolute value, the initial torque value, which is the maximum torque value during a maneuver by the PD benchmark used for comparison. The highest feasible gains (γ fixed) are then obtained and represent the minimum time for that value of γ .

Figure 8 also shows that the gains can increase significantly before the second peak (in absolute value) can be equal to the initial torque. In other words, settling time can be reduced by increasing the

gains when the high gain mode of the controller is reactivated, without causing torque saturation.

2) *Tuning the Parameter γ for the Gain-Scheduled Min-Norm Controller:* Figure 9 shows that, when the gains k_{p2} , k_{d2} go to infinity, the settling time converges to the settling time limit given by (31), which is dependent on γ . A settling time limit requirement can therefore be prescribed to determine.

Figure 10 illustrates the concept of settling time limit on the attitude response plot. If infinite gains were feasible, then the system would proceed on a straight line (when the torque is zero) then stop at the desired zero attitude. In practice, the gains are finite but can be significantly higher than those of the PD benchmark. The theoretical settling time limit can therefore be approached sufficiently to demonstrate significant settling time enhancement over the PD benchmark controller.

From the previous analysis, the controller parameters γ , k_{p2} , k_{d2} can be determined from performance requirements (settling time limit) and practical constraints (torque limitation). The gains k_{p1} , k_{d1} are chosen to be the same as the PD law used as a benchmark. In Figs. 8–11, $k_{p1} = 0.02$, $k_{d1} = 0.5$, and k_{p2} varies as shown on the figures' legends (with constant damping ratio).

3) *Total Integrated Torque of the Gain-Scheduled Min-Norm Controller:* Figure 11 shows that the overall torque expenditure is more dependent on the choice of than the value of k_{p2} . A large range

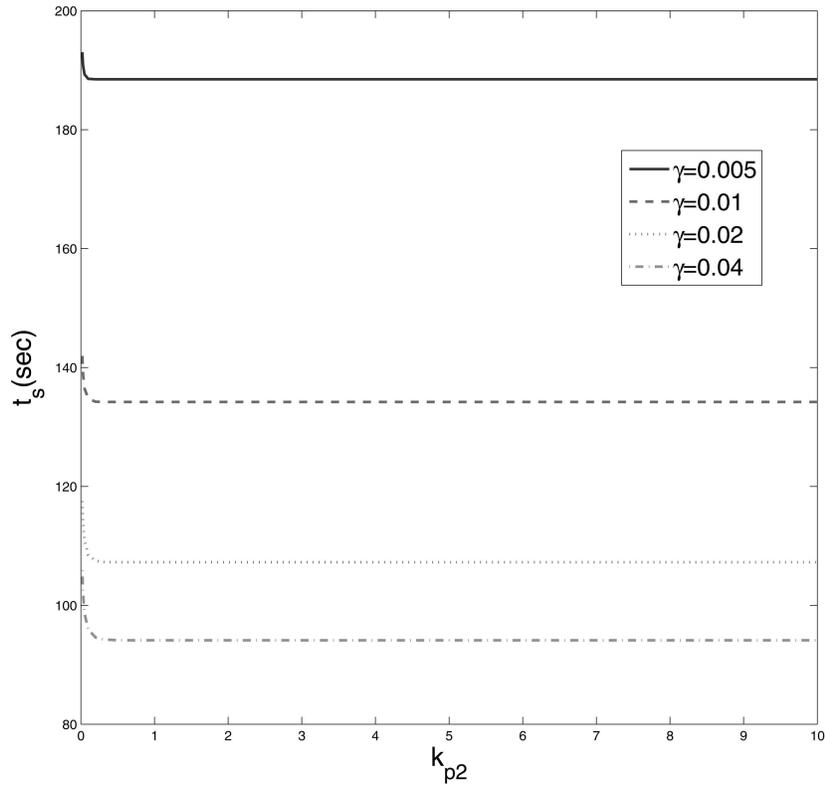


Fig. 9. Settling time as function of control parameters.

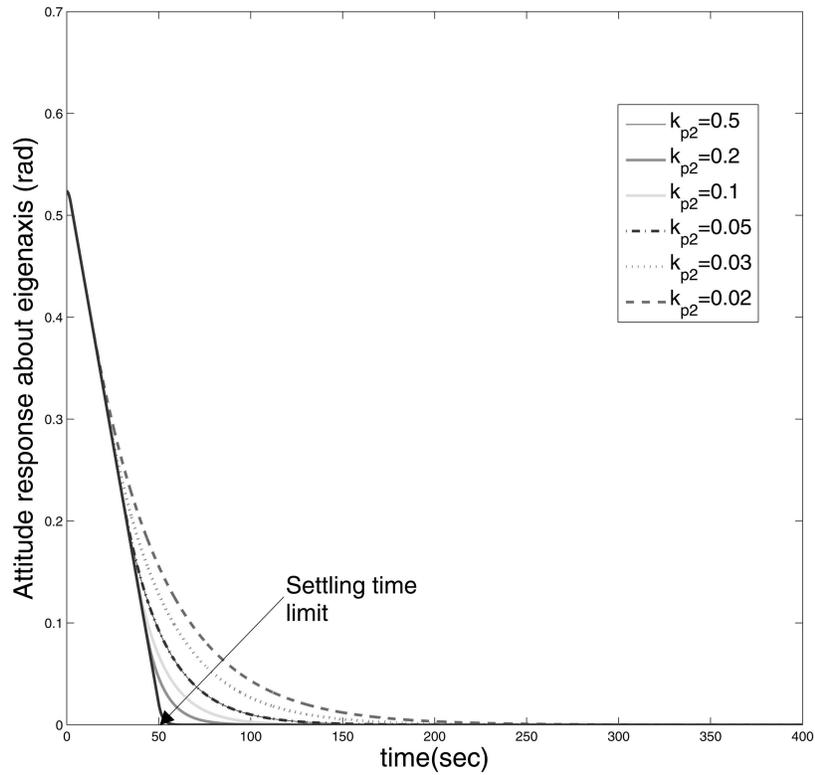


Fig. 10. Attitude response of gain-scheduled min-norm controller by varying gains of high gain mode.

of values of k_{p2} was used starting from k_{p1} until a multiple of five times the torque saturation limit was reached.

Increasing the gains does not only have the effect of increasing the overall torque because the slope of the straight line $(k_{p2}/2)\theta + k_{d2}\dot{\theta} = 0$ increases, meaning

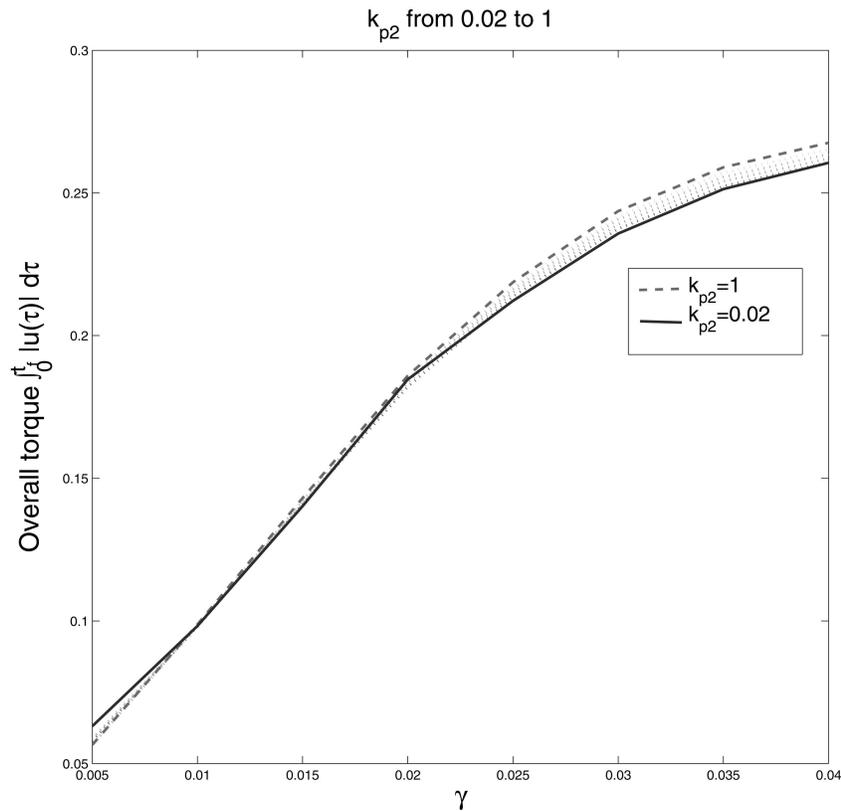


Fig. 11. Overall torque as function of control parameters for $0 < \gamma < kp/kd$.

that the zero torque mode duration also increases, as noted in Section IV-E.

H. Discussion of Implementation Considerations

The proposed gain-scheduled min-norm approach is particularly convenient for 3-axis rotations of up to 30 deg on all three axes. This range of rotations is particularly likely for small Earth observation satellites in their imaging mode, such as those developed by SSTL. Indeed, the high level attitude control requirements of SSTL satellites such as pointing accuracy and agility are generally specified for off-pointing maneuvers of up to 30 deg. Larger rotations are indeed not suitable for the imaging mode that this controller was initially designed for. The typical PD law on SSTL satellites is implemented with a constant set of gains, chosen to meet performance requirements for this range of rotations, despite the fact that performance is known not to be linear. Large rotations remain feasible, although less efficient. It is therefore proposed that the gain-scheduled min-norm controller is also implemented with fixed gains for all maneuvers of up to 30 deg on all axes. This controller was verified (for more initial conditions than shown in Section III-F) to significantly outperform the PD benchmark law with fixed gains, for this range of maneuvers (with settling time reductions between 35% and 52%), with the gains used in Section III-F.

Larger maneuvers should however be preplanned to optimize their performance. Otherwise, the controller with fixed gains would still outperform the benchmark law but the solution would move further away from the minimum time solution. Highly nonlinear maneuvers can be initially tuned by taking the gains of a benchmark law to be the low set of gains. The parameter γ can be determined numerically from a settling time limit on the scalar part q_4 of the quaternion, of which the deviation from 1 is a measure of pointing error. The settling time limit for a parameter γ can indeed be numerically approached by first applying sufficiently high gain values to the high gain mode to approach the solution with infinite gains. The parameter γ can then be adjusted to increase or decrease the settling time limit. The gains of the high gain mode can then be numerically determined by reducing them until maximum torque limitations are met on all axes. In the special case of nonlinear eigenaxis maneuvers beyond the 30 deg range assumed in Section IV, it is also conceivable to store gains in a look-up table.

While the proposed approach is not designed to determine the exact time-fuel optimal solution, it has advantages over the numerical methods of finding the optimal solution. The Euler Lagrange approach is more complex to implement and typically leads to locally optimal open-loop

control laws, with robustness and possibly stability issues. These problems are not fully solved by switching to a feedback law near the origin, in which case optimality is no longer ensured. The HJB approach can be implemented to determine optimal solutions off-line. The optimality of the approach proposed in [16] is however approached by high-order expansions of the control law, requiring the computation of a large number of terms. The controller based on this approach would have to be fully preplanned. The direct transcription methods, such as the pseudospectral technique, are also best implemented off-line and need to be preplanned. While the solution proposed here is generally suboptimal, the implementation complexity is as simple as that of the benchmark controller, compared with which performance is enhanced.

Compared with conventional bang-off-bang techniques (see [23] for example), the proposed gain-scheduled controller has the advantages of being a feedback law for a three-axis attitude control problem. The minimum norm approach ensures robust stability [2] as opposed to the techniques of determining switching times. These open-loop control techniques can also be formulated using switching curves, but this is generally the case for single axis or small rotations. Also, using a min-norm approach, a continuous state feedback is used near the origin, which is preferable under uncertainty and noise to applying a maximum torque. Certain references, such as [23], propose applying a dead zone near the origin, but the attitude error in this case follows a limit cycle and does not converge asymptotically to zero.

V. CONCLUSION

A switched control approach based on inverse optimality has been proposed for the attitude control of a small satellite. The controller was used to enhance the settling time of a benchmark controller, for a prescribed level of the integrated torque. The benchmark controller is a PD law without loss of generality. The settling time reduction was based on enhancing the convergence rate of a Lyapunov function. To incorporate torque saturation constraints, a gain-scheduled formulation of the switched inverse optimal controller was proposed. A controller tuning procedure was developed for this gain-scheduled controller to meet control performance requirements. With this tuning procedure, settling time can be reduced down to a prescribed limit, within the constraints on instantaneous and integrated torques. Without added implementation complexity, significant tradeoff enhancement was obtained compared with the benchmark controller, in the sense of a lower settling time for a given level of the integrated torque. The proposed approach only requires a switching logic between conventional state feedback controllers.

APPENDIX

We assume that \mathbf{I} is invertible. Let σ_{\max} , σ_{\min} , respectively, denote the minimum and maximum eigenvalues of the matrix \mathbf{I}^{-1} .

Positivity of V : The following inequality holds:

$$\mathbf{q}^T \boldsymbol{\omega} \geq -\|\boldsymbol{\omega}\| \|\mathbf{q}\|. \quad (35)$$

The positivity of V (given in (15)) can then be established from the inequality:

$$V \geq [\|\mathbf{q}\| \|\boldsymbol{\omega}\|] \mathbf{K} [\|\mathbf{q}\| \|\boldsymbol{\omega}\|]^T \quad (36)$$

$$\mathbf{K} = \begin{bmatrix} k_p + \gamma k_d & -\frac{\gamma}{2} \\ -\frac{\gamma}{2} & \frac{1}{2} \end{bmatrix}.$$

The parameter can be chosen small enough to make the matrix \mathbf{K} positive definite.

Negativity of \dot{V} : The time derivative of V is given by

$$\dot{V} = (k_p + \gamma k_d) \mathbf{q}^T \boldsymbol{\omega} + \boldsymbol{\omega}^T \dot{\boldsymbol{\omega}} + \gamma \dot{\mathbf{q}}^T \boldsymbol{\omega} + \gamma \mathbf{q}^T \dot{\boldsymbol{\omega}}. \quad (37)$$

By expanding the derivative terms of (37) under PD control $\mathbf{u} = -k_p \mathbf{I} \mathbf{q} - k_d \mathbf{I} \boldsymbol{\omega}$ (with \mathbf{I} invertible) and after simple but tedious calculations, similar to those that we derived in [8] or as in [17], [21], we have

$$\begin{aligned} \dot{V} &= -k_d \|\boldsymbol{\omega}\|^2 - \gamma k_p \|\mathbf{q}\|^2 \\ &+ \gamma [\frac{1}{2} g_4 \boldsymbol{\omega} - \frac{1}{2} \boldsymbol{\omega} \times \mathbf{q}]^T \boldsymbol{\omega} \\ &+ (\gamma \mathbf{q}^T + \boldsymbol{\omega}^T) \mathbf{I}^{-1} (-\boldsymbol{\omega} \times (\mathbf{I} \boldsymbol{\omega} + \mathbf{h})). \end{aligned} \quad (38)$$

The contribution of the term $\boldsymbol{\omega} \times \mathbf{q}$ vanishes because this vector is orthogonal to $\boldsymbol{\omega}$. In other references on quaternion feedback stabilization with external torques (see [21] for example), a bound on the angular velocity $\|\boldsymbol{\omega}\| \leq \beta$ was needed for the purpose of the stability proof. In our case, momentum exchange devices are used and it is sufficient and equivalent to use the momentum conservation constraint $\|\mathbf{I} \boldsymbol{\omega} + \mathbf{h}\| = \beta$ for the purpose of the proof.

We can then conclude that the time derivative of V satisfies the following inequality:

$$\dot{V} \leq -[\|\mathbf{q}\| \|\boldsymbol{\omega}\|] \mathbf{M} [\|\mathbf{q}\| \|\boldsymbol{\omega}\|]^T \quad (39)$$

with

$$\mathbf{M} = \begin{bmatrix} \gamma k_p & \frac{1}{2} \gamma \beta \sigma_{\max} \\ \frac{1}{2} \gamma \beta \sigma_{\max} & k_d - \frac{\gamma}{2} - \beta \sigma_{\max} \end{bmatrix}.$$

The parameter γ can again be chosen small enough (and positive) to make the matrix \mathbf{M} positive definite, provided that $k_d > \beta \sigma_{\max}$. In this case, the term $\gamma k_p (k_d - \beta \sigma_{\max})$ is positive and represents the most significant term in the expansion of $\det(\mathbf{M})$ as all other terms are proportional to γ^2 . σ_{\max} is typically small when moments of inertia are high. We have shown that $\dot{V}_{\text{PD}} \leq 0$, therefore by construction of the

controller in (14), $\dot{V}_{\mathbf{u}} \leq 0$. We also have $V(\mathbf{x}_{\text{eq}}) = 0$. Therefore, V is a Lyapunov function for the system of (3) and (4) and with γ sufficiently small, V can be made strictly negative for $\mathbf{x} \neq \mathbf{x}_{\text{eq}}$, therefore V is also a CLF function. From Barbalat's lemma [19], it follows that $\lim_{t \rightarrow \infty} \mathbf{x}(t) = [0_{1 \times 3}, 1, 0_{1 \times 3}]^T$. The choice of $q_{4\text{eq}} = 1$ rather than $q_{4\text{eq}} = -1$ was discussed following (15).

REFERENCES

- [1] Bharadwaj, S., et al.
Geometry and inverse optimality in global attitude stabilization.
AIAA Journal of Guidance, Control and Dynamics, **21**, 6 (1998), 930–939.
- [2] Freeman, F. A. and Kokotovic, P. V.
Inverse optimality in robust stabilization.
SIAM Journal on Control and Optimization, **34**, 4 (July 1996), 1365–1391.
- [3] Fragopoulos, D. and Innocenti, M.
Stability considerations in quaternion attitude control using discontinuous Lyapunov functions.
IEE Proceedings—Control Theory and Applications, **151**, 3 (May 2004), 253–258.
- [4] Yuqing, H. and Jianda, H.
Generalized point wise min-norm control based on control Lyapunov functions.
Proceedings of the Chinese 26th Control Conference, July 2007.
- [5] Wie, B. and Weiss, H.
Quaternion feedback for spacecraft large angle maneuvers.
AIAA Journal of Guidance, Control and Dynamics, **8**, 3 (1985), 360–365.
- [6] Wie, B. and Weiss, H.
Quaternion feedback regulator for spacecraft eigenaxis rotations.
AIAA Journal of Guidance, Control and Dynamics, **12**, 3 (1989), 375–380.
- [7] Joshi, S. M., Kelkar, K. G., and Wen, J. T.
Robust attitude stabilization using nonlinear quaternion feedback.
IEEE Transactions on Automatic Control, **40**, 10 (1995), 1800–1803.
- [8] Horri, N. M., Palmer, P., and Roberts, M.
Optimal satellite attitude control: A geometric approach.
Proceedings of the 2009 IEEE Aerospace Conference, Big Sky, MT, 2009.
- [9] Peterson, K. S., Grizzle, J. W., and Stefanopoulou, A. G.
Nonlinear control for magnetic levitation of automotive engine valves.
IEEE Transactions on Control Systems Technology, **14**, 2 (Mar. 2006), 346–354.
- [10] Ha, T., Lee, J., and Park, J. H.
Robust control by inverse optimal PID approach for a flexible joint robot manipulator.
IEEE International Conference on Robotics and Biomimetics, Sanya, China, 2007.
- [11] Krstic, M. and Tsiotras, P.
Inverse optimal stabilization of a rigid spacecraft.
IEEE Transactions on Automatic Control, **44**, 5 (May 1999), 1042–1049.
- [12] Freeman, R. A. and Primbs, J. A.
Control Lyapunov functions: New ideas from an old source.
Proceedings of the 35th Decision and Control Conference, Kobe, Japan, 1996, pp. 3926–3931.
- [13] Kanazawa, M., Nakaura, S., and Sampei, M.
Inverse optimal control problem for bilinear systems: Application to the inverted pendulum with horizontal and vertical movement.
Proceedings of the 48th IEEE Conference on Decision and Control, Shanghai, China, 2009.
- [14] Titus, H., Dee, S., and Cunningham, J.
Minimum fuel/time controls via pontryagin revisited.
Proceedings of IEEE International Conference on Control and Applications, 1989, pp. 140–143.
- [15] Luo, W., Chu, Y., and Ling, K.
Inverse optimal adaptive control for attitude tracking of spacecraft.
IEEE Transactions on Automatic Control, **50**, 11 (Nov. 2005), 1639–1654.
- [16] Lawton, J., Beard, R., and McLain, T.
Successive Galerkin approximation of nonlinear optimal attitude control.
Proceedings of the American Control Conference, San Diego, CA, June 1999.
- [17] Luo, W., Chu, Y., Ling, K.
H1 inverse optimal attitude tracking control of rigid spacecraft.
AIAA Journal of Guidance, Control and Dynamics, **28**, 3 (2005), 481–493.
- [18] Primbs, J. A., Nevistic, V., and Doyle, J. C.
Nonlinear optimal control: A control Lyapunov function and receding horizon perspective.
Asian Journal of Control, **1**, 1 (1999), 14–24.
- [19] Khalil, H. K.
Nonlinear Systems (2nd ed.).
Upper Saddle River, NJ: Prentice-Hall, 1996.
- [20] Freeman, R. A. and Kokotovic, P.
Robust Nonlinear Control Design: State Space and Lyapunov Techniques.
Boston: Birkhauser, 1996.
- [21] Wen, J. and Kreutz-Delgado, K.
The attitude control problem.
IEEE Transactions on Automatic Control, **36**, 10 (Oct 1991), pp. 1148–1162.
- [22] Sidi, M. J.
Spacecraft Dynamics and Control: A Practical Engineering Approach.
New York: Cambridge University Press, 1997.
- [23] Chobotov, V. A.
Spacecraft Attitude Dynamics and Control.
Melbourne, FL: Krieger Publishing Co., 2008.

Nadjim M. Horri was awarded a Ph.D. from the University of Surrey.

Dr. Horri is a research fellow at the Surrey Space Centre, University of Surrey. His research interests include optimal satellite attitude control, the control of underactuated satellites and the coordinated attitude and orbit control for satellite formations.



Phil Palmer is a reader at the University of Surrey and an academic member of the Surrey Space Centre, where he leads the astrodynamics research group. His research interests include formation flying, attitude and orbit estimation and control, optical navigation, and spacecraft autonomy.

Mark Roberts is a Professor of Mathematics at the University of Surrey. His research interests include applied mathematics, dynamical systems and astrodynamics.