

Article

Terminal Spacecraft Rendezvous and Capture with LASSO Model Predictive Control

Edward N. Hartley, Marco Gallieri, and Jan M. Maciejowski*

Department of Engineering, University of Cambridge, United Kingdom

(v1.0 January 2013, v1.1 March 2013)

This is an Author's Accepted Manuscript of an article published in *International Journal of Control*, first published online on 20 Aug 2013 [copyright Taylor & Francis], available online at:
<http://www.tandfonline.com/10.1080/00207179.2013.789608>

Changes may have been made to this work since it was submitted for publication, and may not be reflected in this document. The version of record can be found at:

Hartley, E. N. Gallieri, M. and Maciejowski, J. M. (2013), 'Terminal Spacecraft Rendezvous and Capture with LASSO Model Predictive Control', *International Journal of Control*, 86(11), 2104--2113.

doi: 10.1080/00207179.2013.789608

The recently investigated ℓ_{asso} model predictive control (MPC) is applied to the terminal phase of a spacecraft rendezvous and capture mission. The interaction between the cost function and the treatment of minimum impulse bit (MIB) is also investigated. The propellant consumption with ℓ_{asso} MPC for the considered scenario is noticeably less than with a conventional quadratic cost and control actions are sparser in time. Propellant consumption and sparsity are competitive with those achieved using a zone-based ℓ_1 cost function, whilst requiring fewer decision variables in the optimisation problem than the latter. The ℓ_{asso} MPC is demonstrated to meet tighter specifications on control precision, and also avoids the risk of undesirable behaviours often associated with pure ℓ_1 stage costs.

Keywords: Predictive control, LASSO, Spacecraft control, Sparse control, Tuning, Non-smooth characteristics

1 Introduction

Recently, applications of model predictive control (MPC) to a variety of manoeuvres occurring in spacecraft rendezvous have been investigated (Richards and How 2003, Breger and How 2005, 2006, 2007, 2008, Larsson et al. 2006, Saponara et al. 2011, Hartley et al. 2012, Di-Cairano et al. 2012, Gavilan et al. 2012, Kolmanovsky et al. 2012, Sauter and Palmer 2012). MPC offers the ability to explicitly handle physical and operational constraints whilst (approximately) optimising a given performance metric through repeated solution of a fixed, but receding horizon

*Corresponding author. Email: jmm@eng.cam.ac.uk

constrained optimal control problem (e.g. Maciejowski (2002), Camacho and Bordons (2004), Rawlings and Mayne (2009)). Whilst readily acknowledged as more computationally demanding than classical control methods, ever improving computational capabilities and algorithmic developments (e.g. to name a few, Bemporad et al. (2002), Wang and Boyd (2008), Ferreau et al. (2008), Jerez et al. (2011), Mattingley and Boyd (2012), Richter et al. (2012), Domahidi et al. (2012)) increase the range of feasible applications.

As well as efficiently using propellant to satisfy core mission objectives, it can be prudent to consider the number of times the thrusters are switched on/off to avoid excessive mechanical wear on the valves. This paper explores how the recently investigated l_1 -regularised MPC (Ohlsson et al. 2010, Nagahara and Quevedo 2011, Annergren et al. 2012, Gallieri and Maciejowski 2012, 2013) —which amalgamates the classical quadratic cost function with an additional 1-norm penalty on the control input to induce temporal and spatial control sparsity—might perform for this application. This class of cost function is informally dubbed ℓ_{asso} -MPC in Gallieri and Maciejowski (2012, 2013). Also, in the present paper, ℓ_{asso} -MPC is compared with classical linear quadratic MPC, ℓ_1 MPC and MPC with a zone-based (Larsson et al. 2006) cost function.

For simplicity, in this paper only translational control is considered, with the assumption that attitude control is handled by an independent system. A single positive and negative thruster is assumed on each of 3 axes, each providing a commanded force, and zero torque.

The main body of this paper is structured thus: Section 2 summarises the background motivating the investigation; Section 3 defines the control scenario in which the controllers will be evaluated; Section 4 details the tuning methodology used to meet design requirements with each class of cost function; Section 5 presents numerical results; and Section 6 concludes.

2 Background

2.1 Predictive control

Consider a linear state space model of the plant of form $\mathbf{x}(k+1) = \mathbf{A}\mathbf{x}(k) + \mathbf{B}\mathbf{u}(k)$, where $\mathbf{u}(k) \in \mathbb{R}^{n_u}$ and $\mathbf{x}(k) \in \mathbb{R}^{n_x}$ are the input and state at discrete time step k , and \mathbf{A} and \mathbf{B} are appropriately sized matrices with the pair (\mathbf{A}, \mathbf{B}) stabilisable. Let \mathbf{x}_i denote a prediction of $\mathbf{x}(k+i)$, $F_N: \mathbb{R}^{n_x} \rightarrow \mathbb{R}_+$ be a terminal cost, $\ell: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}_+$ be a stage cost, and \mathbf{x}_r and \mathbf{u}_r be state and input reference setpoints. An archetypal (linear) predictive controller solves:

$$J^* = \min_{\mathbf{u}_i, \mathbf{x}_i} F_N(\mathbf{x}_N - \mathbf{x}_r) + \sum_{i=0}^{N-1} \ell(\mathbf{x}_i - \mathbf{x}_r, \mathbf{u}_i - \mathbf{u}_r) \quad (1a)$$

$$\text{s.t. } \mathbf{x}_0 = \mathbf{x}(k) \quad (1b)$$

$$\mathbf{x}_{i+1} = \mathbf{A}\mathbf{x}_i + \mathbf{B}\mathbf{u}_i \quad (1c)$$

$$\mathbf{x}_i \in \mathbb{X}, \forall i \in \{1, \dots, N-1\} \quad (1d)$$

$$\mathbf{u}_i \in \mathbb{U}, \forall i \in \{0, \dots, N-1\} \quad (1e)$$

$$\mathbf{x}_N \in \mathbb{T}. \quad (1f)$$

The first optimal input, denoted \mathbf{u}_0^* , is applied to the plant, the remainder of the solution is discarded (or used for a subsequent warm-start), and the process is repeated at the next time step with $k \leftarrow k+1$. The input \mathbf{u} can be a commanded value for a manipulated variable, or an increment to the manipulated variable (often denoted as $\Delta\mathbf{u}$ (Maciejowski 2002)), in which case the accumulated value is included in the state \mathbf{x} . In the scenario considered here a zero-valued input is preferable to it being at any other steady state, so the former is used.

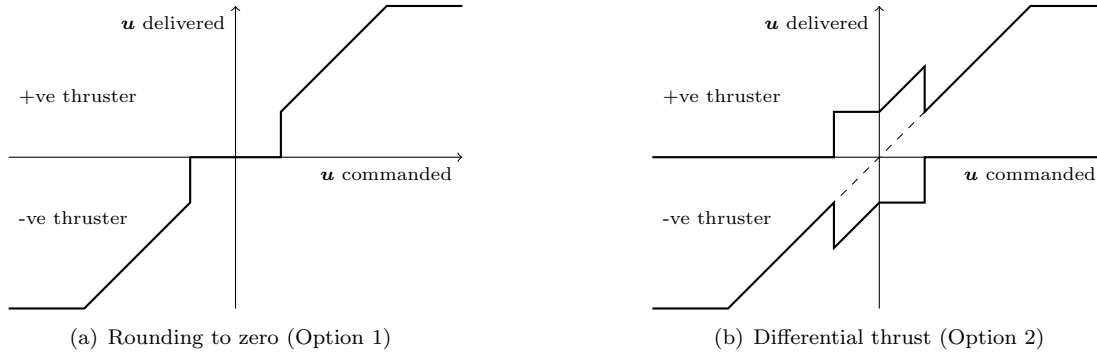


Figure 1. Treatment of the minimum impulse bit

2.2 The minimum impulse bit

When using thrusters for spacecraft trajectory control, not only are there constraints on the maximum force that can be applied at any given instant, and mission specific constraints on the trajectory (e.g. visibility cone constraints, maximum approach velocity), but there is also the physical constraint of a thrust “dead-zone” between the thruster being fully off, and delivering its minimum non-zero thrust, often referred to as the “Minimum Impulse Bit” (MIB). This is neglected or implicitly treated as a disturbance for purposes of control design in the majority of the above-cited references. Unlike certain classes of non-smooth characteristics (del Carmen Rodríguez Liñán and Heath 2012), this non-convex constraint cannot simply be inverted. Possible remedies are:

- (0) using continuous thrusters (evading the issue);
- (1) rounding low thrust requests down to zero (implicitly modelling the MIB as an input disturbance);
- (2) using differential thrust with two opposing forces above the MIB with a net difference of the control command;
- (3) buffering the accumulation of low thrusts and delivering the command only once the deficit of undelivered thrusts exceeds the MIB.

Option 0 presents a conceptual benchmark for comparison. Option 1 precludes small control actions (thus increasing sparsity), but can limit the achievable control accuracy and for low thrusts the plant trajectory will not match the predictions even in otherwise ideal conditions. Option 2 recovers linearity of dynamics (under otherwise ideal conditions) and the ability for fine-grained control at the cost of considerable fuel-burn for low thrusts. Option 3 introduces some unpredictability and can cause the “same” correction to be requested repeatedly.

When inserted between a linear or piecewise-affine feedback control law and the actuators as a “thrust management algorithm”, each of these options can require a different control tuning, due to different assumptions about the relationship between the “real-world” plant dynamics and the prediction model. In Larsson et al. (2006) it is noted that mixed-integer programming could theoretically be used to encode the MIB in an open loop optimal control problem (to be applied in a receding horizon manner), but the complexity grows combinatorially with the number of binary variables.

2.3 Cost functions for predictive control

The most frequently used stage cost in MPC is a weighted quadratic function of the error between the predicted state and input and their respective steady-state targets at current and future sampling instants (letting $\|\bullet\|_{\mathbf{M}}^2 \triangleq \bullet^T \mathbf{M} \bullet$):

$$\ell(\mathbf{x}, \mathbf{u}) = \|\mathbf{x}\|_{\mathbf{Q}}^2 + \|\mathbf{u}\|_{\mathbf{R}}^2, \quad F_N(\mathbf{x}) = \|\mathbf{x}\|_{\mathbf{P}}^2. \quad (2)$$

The weights $\mathbf{Q} \geq 0$, $\mathbf{R} > 0$ and $\mathbf{P} \geq 0$ are tuned to suit the high-level objectives and levels of uncertainty of a given application. The weight \mathbf{P} is often chosen as the solution of the Discrete Algebraic Riccati Equation (DARE) so that the MPC is equivalent to a linear quadratic regulator when constraints are inactive (Chmielewski and Manousiouthakis 1996, Scokaert and Rawlings 1998). For many applications such a design works well. However in the particular spacecraft rendezvous scenario considered here, it has a tendency to favour continuous levels of low-thrust, which gives poor fuel economy if Option 2 is used to compensate for the MIB.

Alternative cost functions can be used to discourage long periods of low-level thrust (i.e. to encourage short and sparsely located pulses of relatively large magnitude). In Rao and Rawlings (2000), the properties of 1-norm and ∞ -norm cost functions are explored, and it is noted that very careful tuning is necessary to avoid dead-beat (minimum time) or idle control. As an alternative to the quadratic cost, Larsson et al. (2006) proposes a zone-based cost, penalising the 1-norm of the thrust (since the magnitude of the force delivered is directly proportional to the propellant consumption) and the 1-norm of deviations outside a given zone around a setpoint. A way to represent this is with an additive slack variable on a constraint:

$$\ell(\mathbf{x}, \mathbf{u}) = \|\mathbf{Q}\mathbf{s}\|_1 + \|\mathbf{R}_\lambda \mathbf{u}\|_1 \quad (3a)$$

$$\text{s.t.} \quad \mathbf{x} \leq \mathbf{b} + \mathbf{s}, \quad -\mathbf{x} \leq \mathbf{b} + \mathbf{s} \quad (3b)$$

$$\mathbf{s} \geq 0. \quad (3c)$$

If the (whole) predicted trajectory is maintained within the target zone then no control activity is necessary. This reduces unnecessary control actions, for example due to sensor noise and is economical in terms of propellant consumption.

The recently considered l_1 regularised (ℓ_{asso}) cost function for MPC (Ohlsson et al. 2010, Nagahara and Quevedo 2011, Annergren et al. 2012, Gallieri and Maciejowski 2012, 2013) could also be an appropriate tool for this application. This cost augments a classical quadratic stage cost with a 1-norm regularisation term:

$$\ell(\mathbf{x}, \mathbf{u}) = \|\mathbf{x}\|_{\mathbf{Q}}^2 + \|\mathbf{u}\|_{\mathbf{R}}^2 + \|\mathbf{R}_\lambda \mathbf{u}\|_1. \quad (4)$$

The term $\|\mathbf{R}_\lambda \mathbf{u}\|_1$ induces a level of temporal (and spatial in the case of multiple inputs) sparsity in the control action. One way to perform the ℓ_1 norm minimisation is through introduction of a slack variable \mathbf{s} :

$$\ell(\mathbf{x}, \mathbf{u}) = \|\mathbf{x}\|_{\mathbf{Q}}^2 + \|\mathbf{u}\|_{\mathbf{R}}^2 + \mathbf{1}^T \mathbf{s} \quad (5a)$$

$$\text{s.t.} \quad \mathbf{s} \geq 0, \quad \mathbf{s} \geq \mathbf{R}_\lambda \mathbf{u}, \quad \mathbf{s} \geq -\mathbf{R}_\lambda \mathbf{u}. \quad (5b)$$

When $\mathbf{R} > 0$ and $\mathbf{R}_\lambda > 0$ are diagonal, the problem can also be formed as:

$$\ell(\mathbf{x}, \mathbf{u}) = \|\mathbf{x}\|_{\mathbf{Q}}^2 + \|\mathbf{u}_-\|_{\mathbf{R}}^2 + \|\mathbf{u}_+\|_{\mathbf{R}}^2 + \mathbf{1}^T \mathbf{R}_\lambda \mathbf{u}_- + \mathbf{1}^T \mathbf{R}_\lambda \mathbf{u}_+ \quad (6a)$$

$$\text{s.t.} \quad \mathbf{u} = \mathbf{u}_+ - \mathbf{u}_-, \quad \mathbf{u}_- \geq 0, \quad \mathbf{u}_+ \geq 0. \quad (6b)$$

If the decision variables are \mathbf{u}_+ and \mathbf{u}_- , with net input $\mathbf{u} = \mathbf{u}_+ - \mathbf{u}_-$ formed as a postprocessing step then, unlike with form (5), the cost function in a condensed MPC will have a strictly positive definite Hessian matrix if $\mathbf{R} > 0$ and $\mathbf{Q} \geq 0$. Matrix \mathbf{R}_λ only affects the linear term in the quadratic program.

Table 1. Prediction model parameters

Name	Symbol	Value	Unit
Gravitational constant	G	6.673×10^{-11}	Nm/kg ²
Mass of mars	M	6.4191×10^{23}	kg
Semimajor axis	a_0	3893.4×10^3	m
Eccentricity	e_0	0	–
Gravitational parameter	μ	4.2835×10^{13}	m ³ s ⁻¹
Orbit angular velocity	$\dot{\nu}$	8.5193×10^{-4}	rad s ⁻¹
Chaser mass	$m_{chs,0}$	1575	kg




Figure 2. Visualisation of relative co-ordinate system in target-centred reference frame




Figure 3. Nominal reference trajectory (not to scale)

3 Control scenario

A relatively simple control scenario based on the final capture phase of the Mars Sample Return (MSR) mission (Beatty et al. 2008, Mura 2007, Saponara et al. 2011, Hartley et al. 2012) is investigated. The linearised relative dynamics of a controller chaser spacecraft with respect to a passive target in a circular orbit are described by the Hill-Clohessey-Wiltshire equations (e.g. Fehse (2003), Sidi (1997)). The parameters used are given in Table 1.

The chaser's position is described in a relative reference frame centred on the target with the x axis aligned with the target velocity, z axis pointing towards the central body, and y axis completing the right-handed set and normal to the orbital plane (Figure 2).

The chaser spacecraft starts at zero velocity, 200 m behind a passive target in a circular orbit around Mars and 7.7 cm below the x -axis (positive z). The objective is to track a step increase in velocity in the x direction to 0.2 ms^{-1} until the separation has reduced to 100 m and then reduce to 0.1 ms^{-1} whilst keeping the position component in the y -axis near zero and in the z -axis near 7.7 cm. Once a separation of $\leq 3 \text{ m}$ is achieved, the remainder of the manoeuvre must be carried out passively. Hence, the value 7.7 cm comes from the distance the chaser will travel in the z direction over 30 s when its velocity in the x direction is 0.1 ms^{-1} and no control is applied. The manoeuvre is summarised in Figure 3.

The prediction model in the controller omits the x position, since only the velocity \dot{x} is to be regulated. A sampling time $T_s = 1 \text{ s}$ is used, and the continuous-time model is discretised using a zero-order-hold, with inputs corresponding to thruster forces. For simplicity, no state or terminal constraint is imposed ($\mathbb{X} = \mathbb{T} \triangleq \mathbb{R}^5$), and box constraints are imposed on the inputs: $\mathbb{U} \triangleq \{\mathbf{u}: \|\mathbf{u}\|_\infty \leq u_{\max}\}$, where $u_{\max} = 8 \text{ N}$. A fixed (receding) prediction horizon of $N = 20$

is used. Navigation error due to sensor noise is modelled as additive, Gaussian and white, with 3σ values of $[0.0247 \text{ m}, 0.0247 \text{ m}, 0.009 \text{ ms}^{-1}, 0.007 \text{ ms}^{-1}, 0.007 \text{ ms}^{-1}]$ on $[y, z, \dot{x}, \dot{y}, \dot{z}]$ respectively. As well as the MIB (selected at 0.544 N), thrust on each axis is subject to a multiplicative uncertainty with $\sigma = 0.01$.

4 Controller tuning

To use the same numerical control weightings for all formulations would be insufficient to make a valid comparison. Each of the cost functions and scenarios should be tuned to achieve some metric of “best” performance in terms of high-level mission objectives, in their own way. Our definition of “best” performance here is:

- minimising propellant consumption to reach target;
- the Euclidean length of the tracking error vector in y - z direction must be within a prespecified tolerance (the specification used by Hartley et al. (2012) is 20 cm , but tighter tolerances leave a margin for further uncertainty);
- the velocity tracking error must be $\leq 0.01 \text{ ms}^{-1}$ after 300 s from the start of the scenario, and remain as such until the step change in velocity set-point;
- the velocity tracking error must be 0.01 ms^{-1} at 3 m from the target; and
- the 3 m point from the target must be passed by the end of a 2000 s simulation.

Following from ideas presented in Kawai et al. (2007), Vega et al. (2008), Joelianto and Hernawan (2009), where stochastic optimisation is used to tune the MPC cost weightings to achieve higher-level objectives (since these are not necessarily convex with respect to the tuning parameters), a Simulated Annealing algorithm (ASA) (Ingber 2012) is used via the MATLAB MEX interface ASAMIN to obtain control parameters satisfying these specifications (with violations of requirements penalised linearly in the cost function, with a weight of 10^3). Other than limiting the number of trials to 3500 (4500) and the number of acceptances to 350 (450), for non-zone (zone) based controller cost functions other configuration parameters retain their default values. The matrices \mathbf{Q} , \mathbf{R} and \mathbf{R}_λ are structured as $\mathbf{Q} = \text{diag}(Q_{yz}, Q_{yz}, Q_{\dot{x}}, Q_{\dot{y}}, Q_{\dot{z}})$, $\mathbf{R} = \text{diag}(R_x, R_{yz}, R_{yz})$, $\mathbf{R}_\lambda = \text{diag}(R_{\lambda,x}, R_{\lambda,yz}, R_{\lambda,yz})$, and the parameters Q_{yz} , $Q_{\dot{x}}$, $Q_{\dot{y}}$, $Q_{\dot{z}}$, R_x , R_{yz} , $R_{\lambda,x}$ and $R_{\lambda,yz}$ are chosen using ASA. The base-10 logarithms of these parameter are used as the decision variable, to reflect that each element may be of a different order of magnitude. This is constrained between -4 and 4 , and initialised as 3 for Q_\bullet and -3 for R_\bullet and $R_{\lambda,\bullet}$. For zone-based costs, the elements (actual, not logarithm) of the vector $\mathbf{b} \triangleq [b_{yz}, b_{yz}, b_{\dot{x}}, b_{\dot{y}}, b_{\dot{z}}]^T$ are also decision variables, constrained between 10^{-4} and 0.2 . For LQ and ℓ_{asso} controllers, $Q_{\dot{y}} = 0$ and $F_N(\mathbf{x}_N) = \mathbf{x}_N^T \mathbf{P} \mathbf{x}_N$ where \mathbf{P} solves the DARE. For ℓ_1 -based controllers $F_N(\mathbf{x}_N) = \ell(\mathbf{x}_N, 0)$.

Each cost evaluation within the Simulated Annealing algorithm is taken as the maximum over four 2000 s simulations where the simulation model takes on each of the four extreme points of $a = a_0 \pm 50 \text{ km}$ and $m_{chs} = (1 \pm 0.15)m_{chs,0}$ as a measure to prevent “over-tuning” (the prediction model retains nominal values) and to provide some practical, if not formal, robustness. In these scenarios, the actual capture is not performed and the tracking continues until the end of the simulation. The tracking accuracy over the full trajectory is used as a related measure. This metric is considered over the full simulation rather than performing a single capture per simulation (i.e. it is considered that the capture might happen at any point in the simulation). In addition, a 30 s open-loop propagation is performed at every point at which the chaser is closer than 75 m from the target. Any lateral ($y - z$) error of the propagated point from the origin, greater than the desired capture tolerance (tolerances of 10 cm , 12.5 cm , 15 cm , 17.5 cm and 20 cm were used) is then penalised with a weighting of 10^3 in the Simulated Annealing cost function. If this measure were not taken, it would be necessary to dramatically increase the number of simulations carried out within the Simulated Annealing process to prevent over-tuning to the specific disturbance sequence used. This would make it computationally intractable. Fuel

Table 2. Number of slacks and decision variables per time step

Cost	States	Inputs	Slacks
Linear quadratic (LQ)	5	3	0
ℓ_{asso} (LASSO) ^a	5	6	0
ℓ_1 (L1)	5	6	5
ℓ_1 with zone (L1Z)	5	6	5

^a Assuming inputs are split into positive and negative components where needed.

Table 3. ΔV usage

Tolerance	10 cm	12.5 cm	15 cm	17.5 cm	20 cm
LQR MIB0	0.91	0.88	0.87	0.86	0.86
LQR MIB1	0.75	0.73	0.72	0.71	0.71
LQR MIB2	3.23	3.23	3.23	3.23	3.24
LQR MIB3	0.91	0.88	0.87	0.87	0.86
LASSO MIB0	0.75	0.73	0.70	0.70	0.70
LASSO MIB1	0.74	0.73	0.72	0.70	0.70
LASSO MIB2	0.91	0.85	0.71	0.74	0.72
LASSO MIB3	0.76	0.72	0.71	0.74	0.71
LP MIB0	2.31	4.90	9.82	9.47	8.49
LP MIB1	8.72	4.69	8.72	1.77	8.58
LP MIB2	7.43	4.82	5.05	4.79	8.70
LP MIB3	4.62	8.52	4.63	8.53	8.52
LPZ MIB0	0.83	0.79	0.73	0.70	0.70
LPZ MIB1	0.87	0.79	0.70	0.70	0.70
LPZ MIB2	1.71	0.80	0.76	0.70	0.70
LPZ MIB3	0.81	0.78	0.72	0.70	0.70

usage is only counted until the point -3 m from the target is reached.

Due to the total number of simulations run, this method takes a substantial time (days) to run, and was carried out on a multi-node compute stack. The primary focus here is neither to advocate nor investigate the efficiency of tuning to meet high level objectives using this particular algorithm over any other, rather simply to explore the degrees of freedom of the different cost functions in an automated way based on the premise that such methods should be more effective than manual tuning, due to the increased number of possible trials.

Since $N = 20$, and the large number of simulations means that fast QP solutions are desirable, an uncondensed predictive control form (Rao et al. 1998) is used and a custom interior point QP solver for the predictive controller within the simulations is generated using CVXGEN (Mattingley and Boyd 2012). The numbers of states, slack variables and inputs required for each class of cost function considered are presented in Table 2.

5 Results

Some metrics pertaining to the worst-case closed loop results from the trials used in the Simulated annealing algorithm are presented in Tables 3, 4 and 5. ΔV (Table 3) represents the sum of the absolute commanded velocity changes, the control density (Table 4) is the fraction of non-zero elements in $[\mathbf{u}_+^T, \mathbf{u}_-^T]$. The density should usually be less than 0.5, except when differential thrust is used to deliver commands from the MPC that are lower than the MIB. The capture accuracy in Table 5 represents the worst case Euclidean deviation from the origin in the lateral (z and y) directions, following a 30 s open-loop propagation from each point along the trajectory where the x position is closer than 75 m from the target. Results in brackets indicate that the tuning objective could not be achieved. From the tabulated results, the following observations can be made:

- The pure ℓ_1 (LP) cost is uncompetitive in terms of ΔV usage in all scenarios examined. The ℓ_{asso} and ℓ_1 -zone based MPC controllers use similar or less ΔV to the basic constrained LQR

Table 4. Control density

Tolerance	10 cm	12.5 cm	15 cm	17.5 cm	20 cm
LQR MIB0	0.48	0.48	0.48	0.48	0.48
LQR MIB1	0.12	0.11	0.11	0.11	0.11
LQR MIB2	0.83	0.83	0.85	0.85	0.83
LQR MIB3	0.48	0.48	0.48	0.48	0.48
LASSO MIB0	0.10	0.11	0.07	0.04	0.08
LASSO MIB1	0.11	0.10	0.11	0.05	0.03
LASSO MIB2	0.11	0.08	0.04	0.04	0.03
LASSO MIB3	0.10	0.09	0.08	0.14	0.10
LP MIB0	0.13	0.21	0.46	0.41	0.37
LP MIB1	0.35	0.17	0.35	0.05	0.34
LP MIB2	0.31	0.21	0.23	0.20	0.40
LP MIB3	0.17	0.37	0.18	0.37	0.37
LPZ MIB0	0.04	0.04	0.04	0.04	0.03
LPZ MIB1	0.04	0.04	0.03	0.03	0.03
LPZ MIB2	0.10	0.04	0.04	0.03	0.03
LPZ MIB3	0.04	0.04	0.04	0.03	0.03

Table 5. Capture accuracy

Tolerance	10 cm	12.5 cm	15 cm	17.5 cm	20 cm
LQR MIB0	0.100	0.124	0.137	0.148	0.162
LQR MIB1	0.100	0.125	0.145	0.173	0.170
LQR MIB2	0.064	0.062	0.064	0.066	0.062
LQR MIB3	0.100	0.123	0.137	0.149	0.165
LASSO MIB0	0.100	0.125	0.146	0.175	0.188
LASSO MIB1	0.100	0.121	0.150	0.149	0.175
LASSO MIB2	0.097	0.124	0.148	0.158	0.188
LASSO MIB3	0.100	0.125	0.145	0.139	0.198
LP MIB0	(0.184)	(0.146)	(0.184)	(0.186)	0.184
LP MIB1	(0.182)	(0.162)	(0.182)	0.169	0.184
LP MIB2	(0.133)	(0.167)	(0.163)	0.146	0.185
LP MIB3	(0.182)	(0.183)	0.134	(0.182)	0.183
LPZ MIB0	(0.118)	0.121	0.150	0.164	0.161
LPZ MIB1	(0.119)	0.125	0.146	0.168	0.168
LPZ MIB2	(0.118)	0.125	0.147	0.168	0.175
LPZ MIB3	(0.120)	0.125	0.144	0.154	0.186

cost function. In the MIB2 case (differential thrust for low values), ℓ_{asso} and ℓ_1 -zone MPC have substantially lower ΔV usage. The ΔV usage for ℓ_{asso} MPC and ℓ_1 -zone MPC is similar except for the tightest tolerance.

- In terms of control density, ℓ_{asso} MPC and ℓ_1 -zone MPC exhibits substantially lower control density, even when differential thrust is used for low thrust values. In contrast, the LQ MPC uses a low, but non-zero level of thrust for a substantial proportion of the trajectory. ℓ_1 -zone gives a slightly sparser solution than ℓ_{asso} when the capture tolerance is tighter. With MIB1, where thresholding is used to forcibly induce sparsity, this gives the best sparsity for LQR-based MPC. It has little effect on ℓ_1 -zone MPC but has an adverse effect on sparsity for ℓ_{asso} . For the less tight capture tolerances, the ℓ_{asso} MPC gives sparser solutions than the LQ MPC even when the former uses MIB2 (differential thrust) and the latter uses MIB1 (thresholding to zero).
- The optimisation-based tuning was unable to find a suitable tuning for pure ℓ_1 -based MPC to meet capture tolerances for 10 cm, and 12.5 cm capture accuracy. Similarly, it was unable to find parameters for the ℓ_1 -zone based MPC to achieve the tightest 10 cm tolerance. ℓ_{asso} and constrained LQR MPC were able to meet these tolerances. This does not guarantee that suitable parameters do not exist, but demonstrates that finding them is sufficiently difficult that the Simulated Annealing algorithm was unable to do so.

Figures 4 and 5 depict the trajectories for LQ, LASSO and L1Z control with the tunings obtained for MIB 0 with 12.5 cm capture tolerance (as an ideal baseline) and the tunings obtained for MIB 2 with the same tolerance, performed using the nominal scenario where plant and model match (but sensor noise and thrust uncertainty are still present). As with ℓ_1 zone control, the

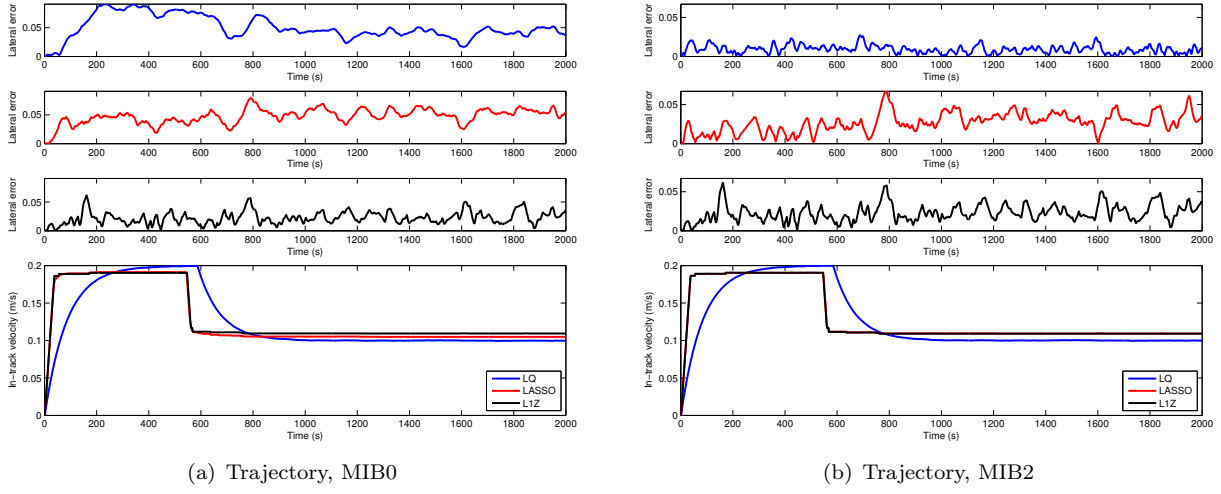


Figure 4. Example closed-loop state trajectories for nominal scenario

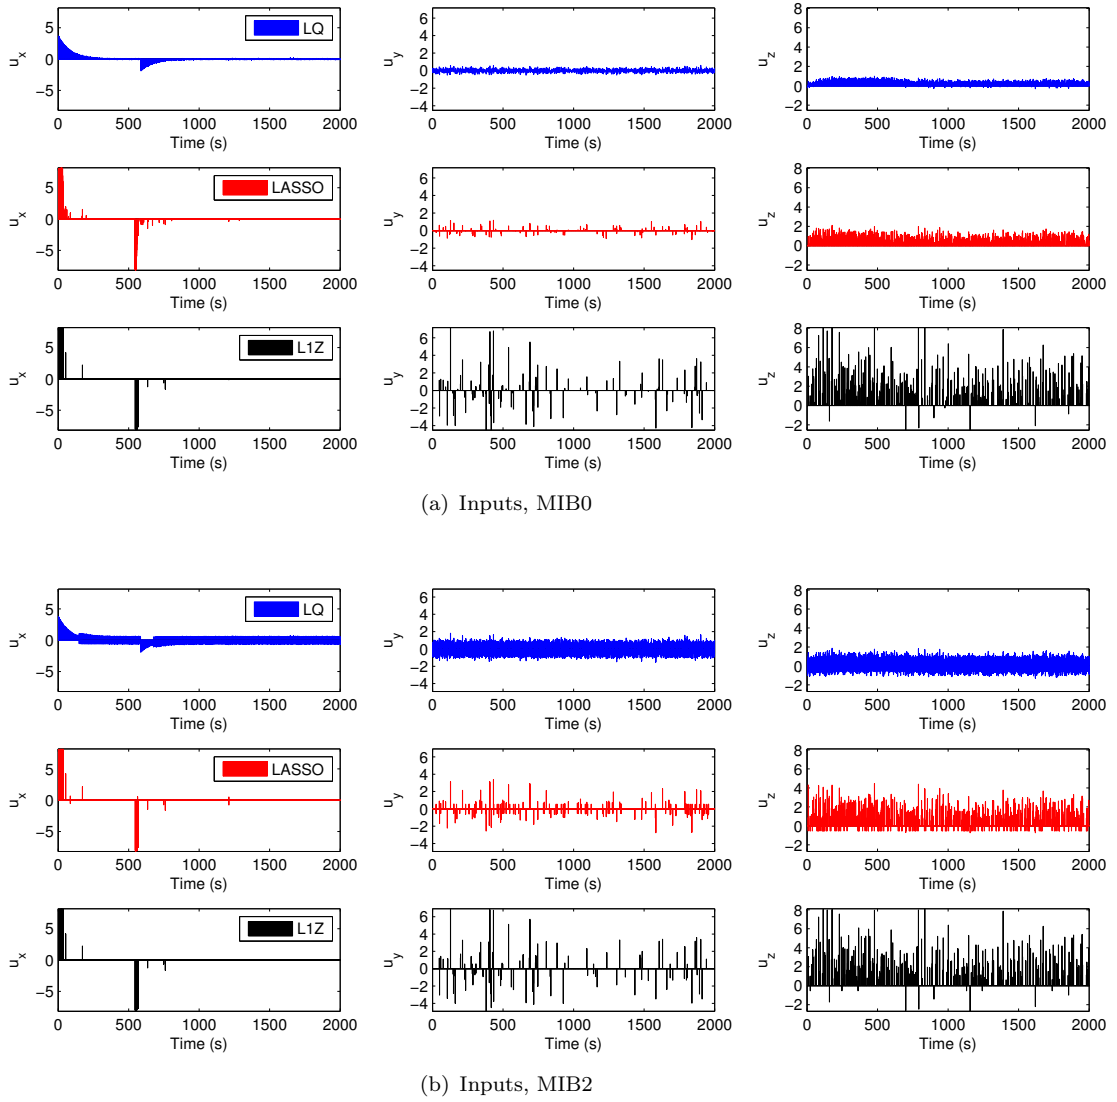


Figure 5. Example closed-loop input trajectories for nominal scenario


Figure 6. Frequency spectrum of error in $y - z$ direction

Table 6. Tuning parameters for 12.5 cm capture tolerance

	Q_{yz}	$Q_{\dot{x}}$	$Q_{\dot{y}\dot{z}}$	R_x	R_{yz}	$R_{\lambda,x}$	$R_{\lambda,yz}$	b_{yz}	$b_{\dot{x}}$	$b_{\dot{y}\dot{z}}$
LQR MIB0	1.983	299.078	—	0.5942	0.4101	—	—	—	—	—
LQR MIB1	22.576	206.721	—	0.0728	1.6408	—	—	—	—	—
LQR MIB2	17.399	128.859	—	0.2445	0.0924	—	—	—	—	—
LQR MIB3	2.005	304.686	—	0.6076	0.4150	—	—	—	—	—
LASSO MIB0	26.860	1096.465	—	0.0230	0.3100	0.4734	2.4157	—	—	—
LASSO MIB1	855.990	6298.239	—	1.3363	48.0799	0.7679	5.3826	—	—	—
LASSO MIB2	0.474	331.488	—	0.0001	0.0005	0.1530	0.0236	—	—	—
LASSO MIB3	2.549	1264.017	—	0.0005	0.0159	0.5822	0.2026	—	—	—
LPZ MIB0	207.273	45.040	2.4167	—	—	0.0370	1.0661	0.1035	0.0180	0.1970
LPZ MIB1	0.147	0.165	0.0072	—	—	0.0007	0.0076	0.0900	0.0101	0.1976
LPZ MIB2	2681.102	3824.480	6.1223	—	—	16.3400	7.1787	0.1027	0.0176	0.1775
LPZ MIB3	120.935	286.031	0.0155	—	—	1.2174	5.8839	0.0900	0.0148	0.1316

ℓ_{asso} controls the position in the z direction with sparse pulses rather than a continuous low-level thrust. Similarly, velocity control in the x direction is sparser. For both options shown, the velocity trajectory more closely resembles that of the ℓ_1 zone control than the LQ, with acceleration performed with a few large pulses. However the lateral error trajectory is slightly smoother as evidenced by Figure 6 which shows the fast fourier transform (Figure 6) of the lateral error amplitude, evaluated as the average of a moving 64 element block multiplied by a Kaiser window. The response drops off slightly faster with ℓ_{asso} than with the ℓ_1 -zone (L1box) control. (Note that the pure LQ control has unfavourable sparsity and/or ΔV use for these scenarios.)

Finally, a Monte-Carlo simulation is carried out to evaluate the variability in performance for the obtained tunings for a capture tolerance of 12.5 cm for LQ MPC, ℓ_{asso} MPC and ℓ_1 zone MPC with MIB 2. (This is the tightest tolerance for which a compliant tuning was found for all of these three classes of cost function.) The random number seed for navigation and thrust uncertainty is varied, and offsets in orbital radius (equal to semimajor axis for circular orbits), chaser mass and additive perturbations to initial conditions are sampled from independent uniform distributions with bounds shown in Table 7.

The numerical values of the tuning parameters for a tolerance of 12.5 cm are presented in Table 6 and the plots of the final captures are shown in Figure 7. Key observations are that

Table 7. Simulated plant model parameter error distributions

Variable	Description	Minimum value	Maximum value
Δx	Initial rel. position err.	-1 m	1 m
Δy	Initial rel. position err.	-1 m	1 m
Δz	Initial rel. position err.	-1 m	1 m
$\Delta \dot{x}$	Initial rel. velocity err.	-5 cm s ⁻¹	5 cm s ⁻¹
$\Delta \dot{y}$	Initial rel. velocity err.	-5 cm s ⁻¹	5 cm s ⁻¹
$\Delta \dot{z}$	Initial rel. velocity err.	-5 cm s ⁻¹	5 cm s ⁻¹
Δa	Orbital semimajor axis err.	-50 km	50 km
Δm_{chs}	Chaser mass err.	-236 kg	236 kg

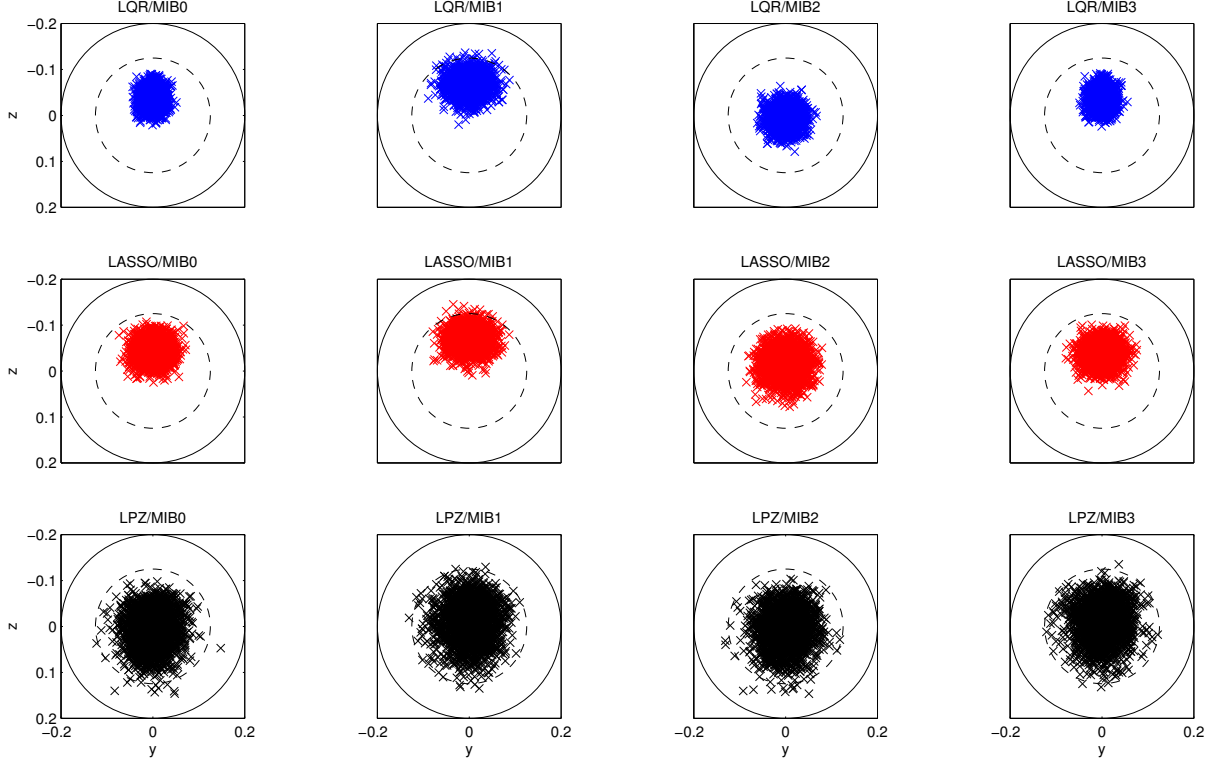


Figure 7. Capture accuracy for tuning with 12.5 cm capture tolerance

MIB1 (thresholding inputs) appears to lead to more bias in the LQ and ℓ_{asso} MPC than the other options, but has little effect on the ℓ_1 zone-based MPC. Also, with MIB0, MIB2 and MIB3, both the LQ and ℓ_{asso} MPC satisfy capture tolerance for each of the 2000 simulations. However, the ℓ_1 zone-based MPC implementation behaves less reliably. It is not clear whether this is an inherent feature of the cost function, or whether the method chosen for tuning is simply less well suited to this structure of cost function.

6 Conclusions

The key observation of this paper is that when differential thrust is used to circumvent the non-convex minimum impulse bit, unlike with a standard LQ cost function, it is possible to find a set of tuning parameters for this application such that a ℓ_{asso} stage cost function can offer similar performance in terms of propellant consumption, sparsity of control action and control accuracy to the zone-control based ℓ_1 cost function advocated by Larsson et al. (2006). With ℓ_{asso} , there is no need to explicitly define the limits of the target zone, thus avoiding an additional set of tuning parameters, and an additional set of slack variables which would increase the computational burden of the online optimisation. In other words, the ℓ_{asso} cost requires fewer

decision variables for tuning and online solution. Whilst for the given tunings, the ℓ_{asso} cost gives a slightly higher “steady state” tracking error, this is offset by a slightly quicker roll off of the spectral content of the error with increasing frequency. As tolerances become tighter, the tuning of the ℓ_{asso} cost causes a slight increase in ΔV usage, but is able to reliably achieve 10 cm capture accuracy. In contrast, a suitable tuning was not achieved for ℓ_1 or ℓ_1 -zone MPC. As a final potential advantage, as with standard linear quadratic MPC, the MPC with ℓ_{asso} cost can be formed in a strictly positive definite way so that the results of Richter et al. (2012) could be invoked during the certification process of a controller. Further investigation would be needed to determine whether the ℓ_{asso} cost function could offer any further advantages when combining attitude and translational control with the same set of actuators and how the controllers respond to larger initial errors.

Acknowledgements

This work was supported by EPSRC Grants EP/G030308/1 and EP/G066477/1.

References

- Annergren, M., Hansson, A., and Wahlberg, B. (2012), “An ADMM Algorithm for solving L_1 regularized MPC,” in *Proceedings of the 51st IEEE Conference on Decision and Control*, Maui, HI, December 10–13, pp. 4486–4491.
- Beaty, D., Grady, M., May, L., and Gardini, B., “Preliminary Planning for an International Mars Sample Return mission,” Report of the International Mars Architecture for the Return of Samples (iMARS) Working Group (2008).
- Bemporad, A., Morari, M., Dua, V., and Pistikopoulos, E.N. (2002), “The explicit linear quadratic regulator for constrained systems,” *Automatica*, 38, 3–20.
- Breger, J.S., and How, J.P. (2007), “Powered Safe Abort for Autonomous Rendezvous of Spacecraft,” in *Proceedings of the AIAA Guidance, Navigation and Control Conference*, Hilton Head, SC, August 20–23.
- Breger, L.S., and How, J.P. (2006), “Safe Trajectories for Autonomous Rendezvous of Spacecraft,” in *Proceedings of the AIAA Guidance, Navigation and Control Conference and Exhibit*, Keystone, CO, Aug. 21–24.
- Breger, L., and How, J.P. (2005), “J2-Modified GVE-based MPC for formation flying spacecraft,” in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, San Francisco, CA, Aug. 15–18, Vol. 1, pp. 158–169.
- Breger, L., and How, J.P. (2008), “Safe Trajectories for Autonomous Rendezvous of Spacecraft,” *Journal of Guidance, Control, and Dynamics*, 31, 1478–1489.
- Camacho, E.F., and Bordons, C., *Model predictive control*, London: Springer-Verlag (2004).
- Chmielewski, D., and Manousiouthakis, V. (1996), “On constrained infinite-time linear quadratic optimal control,” *Systems & Control Letters*, 29, 121–129.
- del Carmen Rodríguez Liñán, M., and Heath, W.P. (2012), “MPC for Plants Subject to Saturation and Deadzone, Backlash or Stiction,” in *Proceedings of the IFAC Conference Nonlinear Model Predictive Control*, Noordwijkerhout, August 23–27, pp. 418–423.
- Di-Cairano, S., Park, H., and Kolmanovsky, I. (2012), “Model predictive control approach for guidance of spacecraft rendezvous and proximity maneuvering,” *International Journal of Robust and Nonlinear Control*, 22, 1398–1427.
- Domahidi, A., Zraggen, A., Zeilinger, M.N., and Jones, C.N. (2012), “Efficient interior point methods for multistage problems arising in receding horizon control,” in *Proceedings of the IEEE Conference on Decision and Control*, Maui, HI, December, pp. 668–674.
- Fehse, W., *Introduction to Automated Rendezvous and Docking of Spacecraft*, Cambridge University Press (2003).

- Ferreau, H.J., Bock, H.G., and Diehl, M. (2008), “An online active set strategy to overcome the limitations of explicit MPC,” *International Journal of Robust and Nonlinear Control*, 18, 816–830.
- Gallieri, M., and Maciejowski, J.M. (2012), “LASSO MPC: Smart regulation of over-actuated systems,” in *Proceedings of the American Control Conference*, Montreal, June, pp. 1217–1222.
- Gallieri, M., and Maciejowski, J.M. (2013), “Stabilising terminal cost and terminal controller for ℓ_{asso} -MPC: enhanced optimality and region of attraction,” in *(To appear in to European Control Conference)*, Zurich, Switzerland, July 17–19.
- Gavilan, F., Vazquez, R., and Camacho, E.F. (2012), “Chance-constrained model predictive control for spacecraft rendezvous with disturbance estimation,” *Control Engineering Practice*, 20, 111–122.
- Hartley, E.N., Trodden, P.A., Richards, A.G., and Maciejowski, J.M. (2012), “Model Predictive Control System Design and Implementation for Spacecraft Rendezvous,” *Control Engineering Practice*, 20, 695–713.
- Ingber, L., *Adaptive simulated annealing (ASA) README* (2012).
- Jerez, J.L., Ling, K.V., Constantinides, G.A., and Kerrigan, E.C. (2011), “MPC for deeply pipelined FPGA implementation: Algorithms and Circuitry,” *IET Control Theory & Applications*, (Under review).
- Joelianto, E., and Hernawan, F.M. (2009), “Multiplexed Model Predictive Control Weighting Selection using Genetic Algorithm,” in *Proc. Int. Conf. Instrumentation, Communication Information Technology, and Biomedical Engineering*, Badung, 23–25 Nov, pp. 435–439.
- Kawai, F., Ito, H., Nakazawa, C., Matsui, T., Fukuyama, Y., Suzuki, R., and Aiyoshi, E. (2007), “Automatic Tuning for Model Predictive Control: Can Particle Swarm Optimization find a better parameter?,” in *IEEE 22nd Int. Symp. Intelligent Control*, Singapore, Oct. 1–3, pp. 646–651.
- Kolmanovsky, I., Baldwin, M., and Erwin, R.S. (2012), “Model Predictive Control of three dimensional spacecraft relative motion,” in *Proceedings of the American Control Conference*, Montreal, June 27–29, pp. 173–178.
- Larsson, R., Berge, S., Bodin, P., and Jönsson, U. (2006), “Fuel Efficient Relative Orbit Control Strategies for Formation Flying and Rendezvous within PRISMA,” in *Proceedings of the 29th Annual AAS Guidance and Control Conference*, Breckenridge, CO, Feb 4–8.
- Maciejowski, J.M., *Predictive Control with Constraints*, Pearson Education (2002).
- Mattingley, J., and Boyd, S. (2012), “CVXGEN: A code generator for embedded convex optimization,” *Optimization and Engineering*, 13, 1–27.
- Mura, F. (2007), “Mars Sample Return: ENG-02 Mission Architecture Definition,” Technical report, European Space Agency.
- Nagahara, M., and Quevedo, D.E. (2011), “Sparse Representations for Packetized Predictive Networked Control,” in *Proceedings of the 18th IFAC World Congress*, Milano, August 28 – September 2, pp. 84–89.
- Ohlsson, H., Gustafsson, F., Ljung, L., and Boyd, S. (2010), “Trajectory generation using sum-of-norms regularization,” in *Proceedings of the 49th IEEE Conference on Decision and Control*, Atlanta, GA, December 15–17, pp. 540–545.
- Rao, C.V., and Rawlings, J.B. (2000), “Linear programming and model predictive control,” *Journal of Process Control*, 10, 283–289.
- Rao, C.V., Wright, S.J., and Rawlings, J.B. (1998), “Application of interior-point methods to model predictive control,” *Journal of Optimization Theory and Applications*, 99, 723–757.
- Rawlings, J.B., and Mayne, D.Q., *Model predictive control: Theory and design*, Nob Hill Publishing (2009).
- Richards, A., and How, J. (2003), “Performance Evaluation of Rendezvous Using Model Predictive Control,” in *AIAA Guidance, Navigation and Control Conference and Exhibit*, Austin, TX, Aug. 11–14.
- Richter, S., Jones, C.N., and Morari, M. (2012), “Computational complexity certification for real-

- time MPC with input constrained based on the fast gradient method,” *IEEE Transactions on Automatic Control*, 57, 1391–1403.
- Saponara, M., Barrena, V., Bemporad, A., Hartley, E.N., Maciejowski, J., Richards, A., Tramutola, A., and Trodden, P. (2011), “Model predictive control application to spacecraft rendezvous in Mars Sample Return scenario,” in *Proceedings of the 4th European Conference for Aerospace Sciences (EUCASS)*, Saint Petersburg, July 4–8.
- Sauter, L., and Palmer, P. (2012), “Analytic Model Predictive Controller for Collision-Free Relative Motion Reconfiguration,” *Journal of Guidance, Control, and Dynamics*, 35, 1069–1079.
- Scokaert, P.O.M., and Rawlings, J.B. (1998), “Constrained linear quadratic regulation,” *IEEE Transactions on Automatic Control*, 43, 1163–1169.
- Sidi, M.J., *Spacecraft dynamics and control: A practical engineering approach*, Cambridge University Press (1997).
- Vega, P., Francisco, M., and Tadeo, F. (2008), “Multiobjective optimization for automatic tuning of Model Predictive Controllers,” in *Proceedings of 17th IFAC World Congress*, Seoul, July, pp. 6980–6985.
- Wang, Y., and Boyd, S. (2008), “Fast Model Predictive Control Using Online Optimization,” in *IFAC World Congress*, Seoul, Korea, July, pp. 6974–6997.