ℓ_{asso} MPC: Smart Regulation of over-actuated Systems

Marco Gallieri and Jan M. Maciejowski

***Abstract*— In this paper, a novel MPC strategy is proposed, and referred to as "ℓ_{asso} MPC". The new paradigm features an ℓ_1-regularised least squares loss function, in which the control error variance competes with the sum of input channels magnitude (or slew rate) over the whole horizon length. This cost choice is motivated by the successful development of LASSO theory in signal processing and machine learning. In the latter fields, "sum-of-norms regularisation" have shown a strong capability to provide robust and sparse solutions for system identification and feature selection. In this paper, a discrete-time dual-mode ℓ_{asso} MPC is formulated, and its stability is proven by application of standard MPC arguments. The controller is then tested for the problem of ship course keeping and roll reduction with rudder and fins, in a directional stochastic sea. Simulations show the ℓ_{asso} MPC to inherit positive features from its corresponding regressor: extreme reduction of decision variables' magnitude, namely, actuators' magnitude (or variations), with a finite energy error, being particularly promising for over-actuated systems.**

I. INTRODUCTION

Nowadays, the most common MPC implementations are based on a quadratic-input-quadratic-state cost, because of its simplicity and due to the possibility of inheriting proprieties from the Linear Quadratic Regulator (LQR). Recently, 1 and ∞-norm costs have also been becoming popular [1].

The theory of ℓ_1-regularised Least Squares (LS), or LASSO, [2], [3], [4], [5], [6], [7], has been widely developed in the fields of machine learning and signal processing. The approach, used to overcome over fitting and to reduce the effects of measurement noise and outliers, penalises the input ℓ_1-norm, providing a sparse solution. ℓ_1-regularised LS is a convex problem that forces most of the decision variables to be equal to zero. This can also be obtained by other ℓ_q-regularisations, with $q \in [0,\ 1)$, resulting in non-convex problems. On the other hand, the solution of a convex ℓ_q-regularisation with $q \geq 2$ will generally feature many small nonzero values. Because of this characteristic, the use of LASSO appears to be promising for MPC, in terms of reduction of actuator activity. LASSO has been used for trajectory optimisation in [8], where its use in predictive control was also suggested. The latter contribution motivates the development of the theory of sum-of-norms regularisation in MPC, in this paper.

Deparment of Engineering, University of Cambridge, UK. {mg574, jmm1}@cam.ac.uk

Research supported by the EPSRC grant "Control for Energy and Sustainability", EP/G066477/1.

II. ℓ_{asso} MPC

Consider an observation vector Y, a design matrix Λ and a vector of decision variables $\mathbf{X}$. The problem

$$\mathbf{X}^\star = \arg\min_{\mathbf{X}} \|\Lambda \mathbf{X} - Y\|_2^2 + \lambda \|\mathbf{X}\|_1 \tag{1}$$

is referred to as an ℓ_1*-regularised least squares* problem, [2], [8], [9], [10], [3], or LASSO regression, in its *unconstrained form*. Equation (1) is a convex problems, non-differentiable at the origin. Non-differentiability causes LASSO to feature, differently from a standard least-squares problem, a piece-wise affine solution (as a function of Y) with $\|X^\star\|_1 \leq t$, for some $t \geq 0$, [10].

The ℓ_1-norm condition forces most of the solution's elements to be null. As discussed before, this feature marks a substantial difference between LASSO and other convex regularised least-squares approaches, for instance the *sum-of-squares* or *Tikhonov regularisation* [4], in which most of the decision variables will always be non-zero. Interestingly, the cost of a Tikhonov regularisation has the same form as for a finite-horizon LQR, which suggests using quadratic MPC as a benchmark for ℓ_{asso} MPC. This parallel allows the authors to investigate whether results achieved in signal processing can be expected in control. The focus of this paper will be on over-actuated systems, where LASSO is expected to combine a regulation task with control allocation.

A. Dual-mode ℓ_{asso} MPC

This paper addresses the control of discrete-time Linear Time-Invariant (LTI) systems[1]

$$x^+ = Ax + Bu. \tag{2}$$

The considered system (2) is subject to constraints of the form

$$u \in \mathbb{U} \subset \mathbb{R}^m, \quad x \in \mathbb{X} \subset \mathbb{R}^n. \tag{3}$$

Consider the following constrained optimal control problem

$$V_N^o = \min_{\underline{\mathbf{u}}} \left\{ V_N(x, \underline{\mathbf{u}}) \hat{=} F(x_N) + \sum_{j=0}^{N-1} \ell(x_j, u_j) \right\} \tag{4}$$

$$\begin{aligned} s.t.: \quad & x_{j+1} = Ax_j + Bu_j, \\ & u_j \in \mathbb{U}, \quad x_j \in \mathbb{X}, \ \forall j \geq 0, \\ & x_0 = x, \ x_N \in \mathbb{X}_f, \end{aligned} \tag{5}$$

[1]Equation (2) expresses the system dynamics, $x(k+1) = Ax(k) + Bu(k)$, in a compact form. Predictions used by the MPC are, on the other hand, denoted as $x_{j+1} = Ax_j + Bu_j$.

with stage cost

$$\ell(x_j,u_j) = \begin{cases} x_j^T Q x_j + u_j^T R u_j + \lambda \|u_j\|_1 & \text{for } x \notin \mathbb{X}_f \\ x_j^T Q x_j + u_j^T R u_j & \text{for } x \in \mathbb{X}_f \end{cases} \tag{6}$$

where $\mathbb{X}_f$ will be a positively invariant set and the terminal cost $F(x)$ a Control Lyapunov Function (CLF) in $\mathbb{X}_f$, to be defined later. Throughout the paper, it is assumed that

Assumption 1.

- **(H0)** (A,B) stabilizable, $(Q^{1/2},A)$ detectable.
- **(H1)** $Q \succeq 0$, $R \succ 0$, $\lambda > 0$,
- **(H2)** $\mathbb{X}$ and $\mathbb{U}$ are closed, bounded and convex,
- **(H3)** $0 \in \text{int}\{\mathbb{X}\}$, $0 \in \text{int}\{\mathbb{U}\}$,
- **(H4)** $x_N \in \mathbb{X}_N = \{x \in \mathbb{R}^n \mid \exists\ \underline{u} \in \mathbb{U} \mid x \in \mathbb{X},\ V_N(x_0,\underline{u}) < \infty\}$.

Define the following

$$H = 2(\mathbf{R} + \Theta^T \mathbf{Q} \Theta), \quad G = 2\Theta^T \mathbf{Q} \Psi x, \tag{7}$$

with

$$\Psi = \begin{bmatrix} A \\ A^2 \\ \vdots \\ A^N \end{bmatrix}, \quad \Theta = \begin{bmatrix} B & 0 & \cdots & 0 \\ AB & B & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A^{N-1}B & A^{N-2}B & \cdots & B \end{bmatrix}, \tag{8}$$

$$\mathbf{Q} = \text{BlkDiag}\{I_{N\times N} \otimes Q, P\}, \quad \mathbf{R} = I_{N\times N} \otimes R, \tag{9}$$

where BlkDiag$\{A,B\}$ denotes a block-diagonal matrix, and $\otimes$ denotes the Kronecker product.

Theorem II.1. Problem (4–5) is an ℓ_1-regularized LS problem, subject to convex constraints.

Proof. Using a standard condensed QP form for the quadratic cost [11], problem (4) can be recast as

$$\min_{\underline{u}} \left\| \frac{1}{\sqrt{2}} H^{1/2} \underline{u} - \left(-\frac{1}{\sqrt{2}} H^{-1/2} G \right) \right\|_2^2 + \lambda \|\underline{u}\|_1. \tag{10}$$

Substituting

$$\mathbf{X} = \underline{u}, \quad \Lambda = \tfrac{1}{\sqrt{2}} H^{1/2}, \quad Y = -\tfrac{1}{\sqrt{2}} H^{-1/2} G \tag{11}$$

(10) has the same form as (1). Constraints are in (5). ■

In the proposed approach, the solution, $\underline{u}^\star$, of (4) is computed online and only the first command is applied to the plant. This is then repeated in a receding horizon fashion. The resulting controller is defined as ℓ_{asso} MPC with (input) magnitude regularisation. Similarly, replacing u_j by its one step time difference, Δu_j, problem (4) will define a ℓ_{asso} MPC with rate regularisation. In contrast to quadratic MPC, non-differentiability of the 1-norm penalty prevents the computation of a closed form solution for the unconstrained problem.

1) Unconstrained ℓ_{asso} MPC: Define $\mathbf{r} = \Lambda \mathbf{X} - Y$ and $\mathbf{X} = \underline{u}$. The Karush-Kuhn-Tucker (KKT) conditions for a sum-of-norms-constrained least-squares problem are [10]

$$(2\Lambda^T(\Lambda \underline{u}^\star - Y))_i \in \begin{cases} \{+\lambda\} & \text{if } \underline{u}_i > 0 \\ \{-\lambda\} & \text{if } \underline{u}_i < 0 \\ [-\lambda,\ \lambda] & \text{if } \underline{u}_i = 0, \end{cases} \tag{12}$$

that is [3]

$$\lambda \geq 2\|\Lambda^T \mathbf{r}^\star\|_\infty. \tag{13}$$

From (13), a null control solution is possible for a finite value of λ, given x

$$\lambda^{\max}(x) = 2\|\Lambda^T Y\|_\infty \equiv 2\|\Theta^T \mathbf{Q} \Psi x\|_\infty. \tag{14}$$

Consider an unconstrained finite horizon ℓ_{asso} MPC problem. Substituting (7) and (11) into (13) gives

$$\lambda^\star(x, \underline{u}^\star) = 2\|\Theta^T \mathbf{Q}(\Psi x + \Theta\ \underline{u}^\star) + \mathbf{R}\ \underline{u}^\star\|_\infty. \tag{15}$$

Theorem II.2. Given $x = x(k)$, a null optimal solution for the unconstrained ℓ_{asso} MPC, $\underline{u}^\star = 0$, is obtained for

$$\lambda \geq \lambda^\star(x, 0) \mathrel{\hat{=}} \lambda^{\max}(x). \tag{16}$$

Proof. The KKT conditions (12) give the first inequality. Evaluate (15) at $\underline{u}^\star = 0$ to verify the second equality. ■

Furthermore, an implicit thresholding capability can be deduced, from (12), for the unconstrained ℓ_{asso} MPC.

2) Proposed approach: In this paper, stabilisation to zero is achieved by means of a *dual-mode* approach [12]. Since in (6) it is assumed that $\lambda = 0$ for $x \in \mathbb{X}_f$, the terminal set and controller can be defined as for a "constrained LQR" [13]. Define the following, as respectively, *maximal admissible set*, *input admissible set*, *terminal set*, *terminal cost* and *terminal controller* [14]

$$\mathcal{O}_\infty = \{x \mid (A - BK_\infty)^k x \in \bar{X},\ \forall\ k \geq 0\} \tag{17}$$

$$\bar{X} = \{x \in \mathbb{X} \mid -K_\infty x \in \mathbb{U}\} \tag{18}$$

$$\mathbb{X}_f = \{x \mid x^T P_\infty x \leq c\} \tag{19}$$

$$F(x) = x^T P_\infty x \tag{20}$$

$$K_f(x) = -K_\infty x \tag{21}$$

where

$$0 < c < c_m \mathrel{\hat{=}} \inf_{x \notin \bar{X}} \{x^T P_\infty x\} \tag{22}$$

Consequently, the *stabilisable set* is given by

$$\mathcal{S}_N = \mathcal{K}_N(\mathbb{X}, O_\infty) \subseteq \mathcal{S}_\infty \tag{23}$$

where $\mathcal{K}_N$ is the N-step *controllable set* and $\mathcal{S}_\infty$ is the *maximal stabilisable set* [15], [14]. The ℓ_{asso} MPC control law will be referred to as $K_N(x)$.

Theorem II.3. (Zero-regulation) Assume (17–22). The constrained dual-mode ℓ_{asso} MPC solving problem (4–5) is exponentially stabilising $\forall\ x \in \mathbb{X}_N$. Hence, the state origin is locally exponentially stable (LES) with domain of attraction $\mathbb{X}_N \subseteq \mathcal{S}_\infty$.

Proof. The proof relies on the direct method in [16], [17]. The optimal cost function, V_N^o, is a positive definite function of the state. Moreover, it is radially unbounded. Given, at time k, the optimal sequence, $\underline{\mathrm{u}}^\star = \{u_0,\dots,u_N\}$, the sequence, $\tilde{\underline{\mathrm{u}}} = \{u_1,\dots,u_N,K_f\}$, is also admissible at time $k+1$. Hence

$$\begin{aligned} V_N^o(x^+) \le V_N(x,\tilde{\underline{\mathrm{u}}}) = V_N^o(x) - \ell(x, K_N(x)) \\ -F(x_N) + \ell_f(x_N, K_f) + F(x_N^+). \end{aligned} \tag{24}$$

Since P_∞ solves a matrix Lyapunov equation, $F(x)$ satisfies $\ell_f(x,K_f)+F(x^+)-F(x)=0$, $\forall x \in \mathbb{X}_f$. Hence, $\forall x \in \mathbb{X}_N$, it follows that

$$V_N^o(x^+) - V_N^o(x) \le -\ell(x, K_N(x)). \tag{25}$$

Therefore, the optimal value function, V_N^o, is a Lyapunov function for the closed-loop system, $\forall x \in \mathbb{X}_N$. Hence, the origin is asymptotically stable. Finally, since $\exists\ a_1, a_2$ such that $a_1\|x\|_2^2 \le \ell(x,\underline{\mathrm{u}})$, $\forall x \in \mathbb{X}_N$ and $F(x) \le a_2\|x\|_2^2$, $\forall x \in \mathbb{X}_f$, it follows [17] that the origin is locally exponentially stable, with domain of attraction $\mathbb{X}_N \subseteq \mathcal{S}_\infty$. ■

Corollary II.4. **(Zone-regulation)** If the terminal controller is not applied, the system state is ultimately bounded by

$$\mathcal{B}_{N,d^\star} = \{x \mid V_N^o(x) \le d^\star\}, \tag{26}$$

with

$$\begin{aligned} d^\star = \min d \\ \text{s.t.: } \mathcal{B}_{N,d} \supseteq \{Ax \mid x \in \mathbb{X}_f\}. \end{aligned} \tag{27}$$

B. Implementation

A 1-norm cost can be represented by introducing appropriate slack constraints, and penalising the new slack variables, σ_k. Problem (4–5) can be now formulated as a set of N constrained semi-definite QPs, for example in the condensed formulation

$$X_k^\star = \arg\min_{X_k} \frac{1}{2}\ X_k^T \tilde{H} X_k + X_k^T \tilde{G} \tag{28}$$

where:

$$X_k = \begin{bmatrix} \underline{\mathrm{u}} \\ \sigma_k \end{bmatrix},\quad \tilde{H} = \begin{bmatrix} H & 0 \\ 0 & 0 \end{bmatrix},\quad \tilde{G} = \begin{bmatrix} G \\ \lambda\ S_{(mN\times 1)} \end{bmatrix} \tag{29}$$

where S is a selection matrices, consisting of an identity matrix and zero block, in order to distinguish which input are regularised and which are not. The problem is subject to the following constraints, similarly as in [11],

$$\Omega X_k \le b \tag{30}$$

$$\Omega = \begin{bmatrix} \Gamma_g\Theta & 0 \\ W & 0 \\ \multicolumn{2}{c}{\Pi} \end{bmatrix}, \tag{31}$$

$$b = \begin{bmatrix} -\Gamma_g \Psi X_k - \bar{g} \\ -\bar{w} \\ 0_{(3mN\times 1)} \end{bmatrix}, \tag{32}$$

$$\Pi = \begin{bmatrix} I_{mN} & -I_{mN} \\ -I_{mN} & -I_{mN} \\ 0 & -I_{mN} \end{bmatrix}. \tag{33}$$

In eq.(31), Γ_g and $\bar{g}$ represent the state constraints, as well as W and $\bar{w}$ provide input constraints. These variables are all in standard form, as in [11]. Finally, ϖ provides the 1-norm slack constraints. Because of this slack constraints, for LTI models the ℓ_{asso} MPC is expected to be a piecewise affine state-feedback controller, even in the absence of state or input constraints.

To impose an attainable terminal constraint, $x_N \in \mathbb{X}_f$, the following steps are performed offline

Algorithm 1. Compute:

1) An admissible polyhedron, $\mathcal{X}_\infty \subseteq \mathcal{O}_\infty$ [15],
2) The level set c_m in (22) outside $\mathcal{X}_\infty$, by solving a set of QPs [13]. Choose $c < c_m$ to determine $\mathbb{X}_f$,
3) A (non-invariant) polyhedral approximation, $\mathcal{X}_f \subseteq \mathbb{X}_f$, using *Algorithm 2*,
4) The domain of attraction $\mathbb{X}_N \subseteq \mathcal{S}_N$, as in [14].

Since the terminal set is an ellipsoid defined by $x^T(\frac{1}{c}P_\infty)x \le 1$, its polyhedral approximation can be determined using elliptical geometry and intersection.

Algorithm 2:

1) Compute the eigenvalues λ_i and eigenvectors v_i, for $i = 1,\dots,n$ of the matrix $\mathcal{E} = cP_\infty^{-1}$.
2) For each pair $\{\lambda_i, \lambda_j\}$ with $i \ne j$, compute:
   a) The semi-axes length, the focus and semi-latus rectum, $a = \sqrt{\max\{\lambda_i,\lambda_j\}}$, $b = \sqrt{\min\{\lambda_i,\lambda_j\}}$, $f = \sqrt{a^2-b^2}$, $l = b^2/a$,
   b) The n-dimensional cone, $\mathcal{P}_{ij}$, defined by $(\pm)v_a^T x \le f$, together with $(\pm)v_b^T x \le l$.
3) Compute $\mathcal{X}_f = \bigcap_{i\ne j} \mathcal{P}_{ij}$.

In Algorithm 2, v_a is the eigenvector associated with $\max\{\lambda_i,\lambda_j\}$, the direction of the major semi-axis. Similarly, v_b is for the minor semi-axis.

Online, $x_N \in \mathcal{X}_f$ is imposed, for ease of computation, providing $x_N \in \mathbb{X}_f$. The terminal controller K_f is then applied iff $F(x_N) \le c$, ensuring positive invariance for $\mathbb{X}_f$. As an example, the sets computed for a second order LTI system are shown in Figure 1.

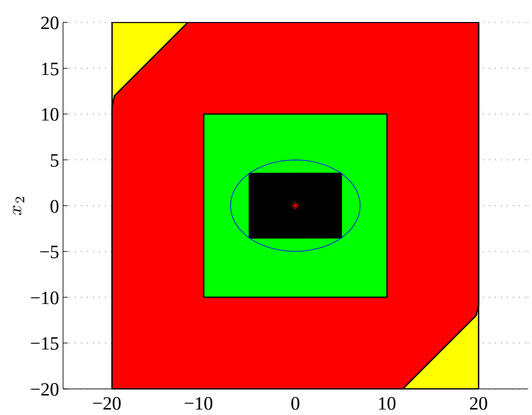

Fig. 1. Sets for a 2nd order LTI system, using Matlab Invariant set toolbox and Multi-parametric toolbox: $\mathbb{X}$ (yellow), $\mathbb{X}_N$ (red), $\mathcal{X}_\infty$ (green), $\mathbb{X}_f$ (ellipse interior), $\mathcal{X}_f$ (black).

III. SHIP ROLL DAMPING

Roll motion in ships is a major cause of problems to passengers and human operators. The issue concernes most marine activities, from luxury yachting to offshore operations. The reduction of vessel roll, or roll damping, has been widely addressed in the past, via several techniques, [18]. Recent results have shown the use of combined rudder and fins control to provide several benefits, including less interference with the vessel course keeping. In the latter case, the system to be controlled can be considered to be *fully-actuated*.

Whenever fins come into play, constraints on the effective angle of attack are vital, in order to avoid non-linear fluid separation phenomena, a major cause of performance degradation and instability. The necessity of handling constraints motivated the recent development of MPC-based fin stabilisers [19]. Unfortunately, quadratic MPC causes all actuators to be in use for most of the time, whereas it would be preferable for only fins to act to reduce roll most of the time, with the rudder applying additional torque only when the fins have reached the limit of their authority. Less roll-induced rudder action is recommendable, in order to have less yaw interference, the main motivation for the presence of fins, which are expensive devices. This motivates our consideration for the problem of roll damping, as a 'demonstrator' for ℓ_{asso} MPC. (Of course the rudder should also act unimpeded in its primary role of steering the ship.)

A. Vessel dynamics and MPC design

Vessel dynamics are formulated according to marine craft maneuvring theory [20]. The equations of motion include rigid-body dynamics and linear hydrodynamic effects, given by the so-called hydrodynamic derivatives, generally obtained by system identification [19]. Waves are treated as a Gaussian output disturbance whose distribution is based on a power spectrum, the choice of which depends on several factors [20]. In this paper, a JONSWAP spectrum [20] is chosen, as an example. The vessel moves in 6 degrees of freedom (DOF), and its motion is generally expressed in standard SNAME coordinates and reference frames[2] [20], [18], summarised in Table I. Full state measurement is assumed.

Variable name	Description
$n,\ e,\ d$	North, east, down positions, n-frame
$\phi,\ \theta,\ \psi$	Roll, pitch, yaw (Euler) angles, $n \to b$
$u,\ v,\ w$	Surge, sway, heave velocities, b-frame
$p,\ q,\ r$	Roll, pitch, yaw rate, b-frame
ν	Vector of gen. velocities and rates, b-frame
η	Vector of positions and attitude, n-frame

TABLE I

SNAME COORDINATES

The controller is required to maintain a desired average yaw angle, while reducing the roll variance, despite the action of sea waves. For control design, a discrete-time LTI model is obtained from the original nonlinear system. For comparison, the model and controller in [18] will be considered as a benchmark. In particular, a quadratic MPC is implemented, as in [18], with matrices of the form:

$$Q = \mathrm{diag}\{0,\ 0,\ q_r,\ q_\phi,\ q_\psi\}, \quad R = \mathrm{diag}\{q_R,\ q_F,\ q_F\}. \tag{34}$$

Constraints are clearly described in [18].

Our ℓ_{asso} MPC will have the same Q matrix as for the above "constrained LQR", while the regularisation parameter will be tuned by trial and improvement. Figure 2 shows the projection on the roll and yaw coordinates of several approximations of $\mathcal{O}_\infty$, computed with the Matlab Invariant-set toolbox and Multi-parametric toolbox.

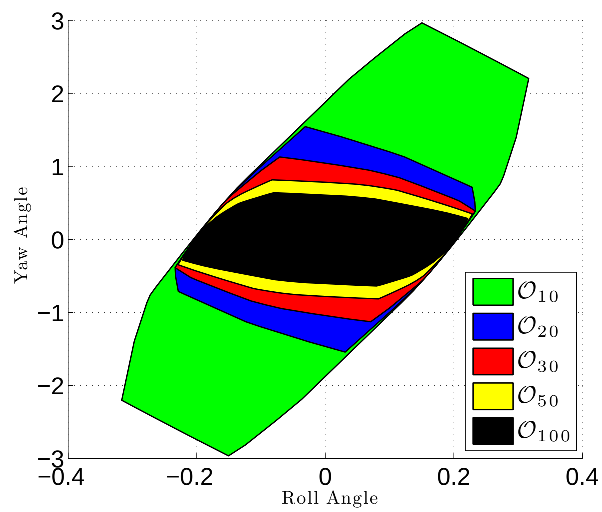

Fig. 2. Approximations of $\mathcal{O}_\infty$, using Matlab Invariant set toolbox and Multi-parametric toolbox (roll–yaw projection).

B. Simulations

Results are shown, first for an input-magnitude-penalised formulation, then for the case of input rate penalisation. The horizon length is chosen to be of 15 steps, with a sampling period of 0.1 sec. The terminal controller is never applied, in a zone-regulation fashion. The chosen target for the system is $\mathcal{O}_{100}$, in Figure 2. Performances are evaluated by computing the mean and standard deviation (STD) of the roll angle, and of the low-frequency yaw, obtained by a low-pass filter with a cut-off frequency of $1.5\ \omega_0$, where ω_0 is the mean wave frequency. The ℓ_1-norm of the overall input signal is also computed. Simulations are performed for an irregular sea, namely, a Sea State 5 (SS5) [18], and results are averaged over 40 simulations.

For the first case, $q_\phi = 20$, $q_r = 10$, $q_\psi = 100$, $q_R = 20$, $q_F = 10$, for the quadratic MPC while, for the ℓ_{asso} MPC, $q_\phi = 20$, $q_r = 10$, $q_\psi = 100$, $q_R = 2$, $q_F = 1$, for $\lambda = 0.9$ and $\lambda = 1.8$.

Table II shows that, for the dual mode ℓ_{asso} MPC, the tradeoff between control error statistics and input amplitude can be regulated by choosing an appropriate λ. This parameter, together with the quadratic input cost, can be tuned in order to provide results which are close to the ones given by quadratic MPC. To challenge the new approach, the R

[2] SNAME stands for "Society of Naval Architects and Marine Engineers".

TABLE II
COMPARISON IN SEA STATE 5, FOR VARIOUS λ

λ	Roll std	Yaw std	Roll mean	Yaw mean	E($\|\underline{u}\|_1$)
-	0.2396	0.6367	-0.0117	0.0167	6999.7
0.9	0.1625	0.4587	-0.0023	0.1513	7231.8
1.8	0.2108	0.6621	0.0045	0.1306	4386.7

matrix for quadratic MPC (indicated by "$\lambda = -$" in Table II) is chosen to be 10 times the one used in ℓ_{asso} MPC. Despite this, the results obtained with ℓ_{asso} MPC are better than the ones achieved with quadratic MPC, in terms of the tradeoff between error standard deviation and input magnitude. Note that benefits on the ℓ_1-norm of the input signals, expected in the nominal LTI case, are still achieved for the nonlinear stochastic system.

Figures 3, 4 show the performances and control signals obtained with quadratic MPC (dashed blue) and by the ℓ_{asso} MPC (continous black). As expected, in contrast to quadratic MPC, the ℓ_{asso} MPC is capable of holding the rudder angle exactly at zero for most of the time. Similarly for the fins — namely, the solution is 'sparse in time' as well as 'sparse in actuators'. On the other hand, quadratic MPC seems to generate rather unrealistic commands, with magnitudes lower than 1 degree. These are not likely to be achievable by ether the rudder or the fin machinery.

Wave filters are generally used to reduce control signals' sensitivity to high frequency yaw motion, and to reduce rudder's activity [18]. From Figure 4, it can be noticed how ℓ_{asso} MPC implements an *input thresholding*, without the use of such wave filters. The thresholding resembles what obtained in LASSO-based feature detection [21], [22].

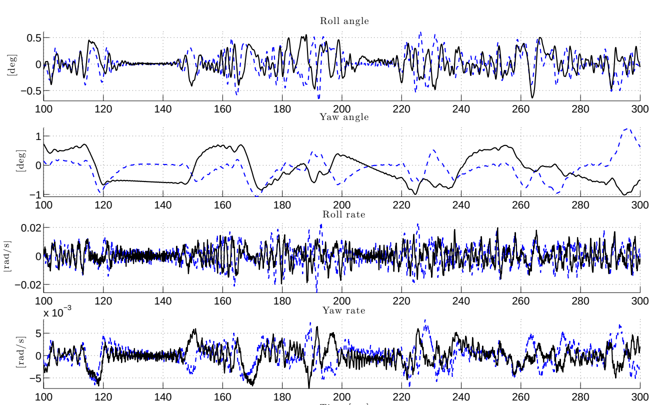

Fig. 3. Comparison in sea state 5: Roll and yaw angles and rates, ℓ_{asso} MPC (solid) $\lambda = 1.8$, quadratic MPC (dashed).

Results are now shown for an input rate penalty, with $q_\phi = 20$, $q_r = 10$, $q_\psi = 100$, $q_R = 20$, $q_F = 10$, for the quadratic MPC and, $q_\phi = 20$, $q_r = 10$, $q_\psi = 100$, $q_R = 2$, $q_F = 1$, or the ℓ_{asso} MPC with $\lambda = 2.5$ and $\lambda = 5$.

From Table III, it can be seen that, for $\lambda = 1.5$, performances can be similar to the ones obtained with a quadratic MPC (this time with the same quadratic cost), in addiction to a significant reduction (approx. 40%) of the ℓ_1-norm of the input variations, that is, with more constant input signals.

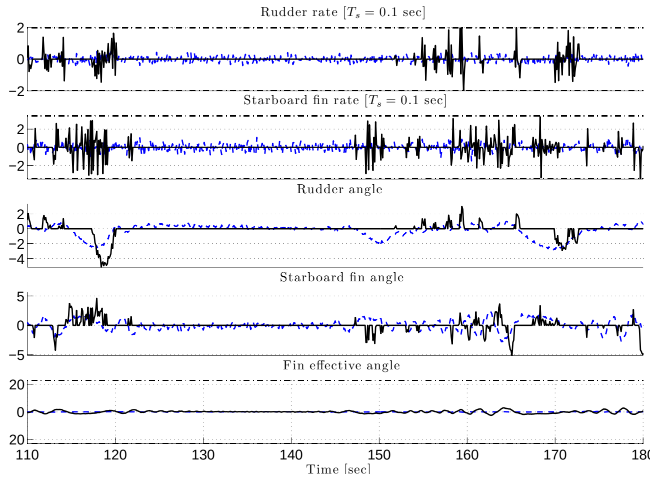

Fig. 4. Comparison in sea state 5: Actuators. ℓ_{asso} MPC (solid) $\lambda = 1.8$, quadratic MPC (dashed).

TABLE III
COMPARISON IN SEA STATE 5, FOR VARIOUS λ (Δu FORMULATION)

λ	Roll std	Yaw std	Roll mean	Yaw mean	E($\|\Delta\underline{u}\|_1$)
-	0.2267	0.4061	-0.0021	0.0552	2.76e+03
1.5	0.2576	0.4512	-0.0074	0.0283	1.9e+03
2.5	0.2552	0.4437	-0.0261	0.1222	1.54e+03

In this case, it is interesting to see how performances change for different λs: in particular, a smaller value will provide the reduction of the mean values, while a bigger one reduces the error variances. Figures 5, 6 show the performances and control signals obtained with quadratic MPC (dashed blue) and by the ℓ_{asso} MPC (continuous black).

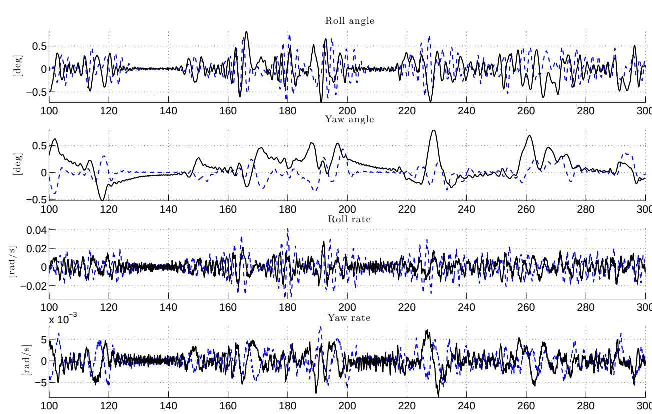

Fig. 5. Comparison (Δu) in sea state 5: Roll and yaw angles and rates, ℓ_{asso} MPC (solid) $\lambda = 1.5$, quadratic MPC (dashed).

Simulations seem to confirm what is experienced in compressed sensing, where LASSO can outperform quadratic regularisation, in terms of accuracy of reconstruction, and robustness to noise [10].

IV. CONCLUSIONS

In this paper, a novel stabilising dual-mode MPC has been formulated, and referred to as ℓ_{asso} MPC. The approach is based on sum-of-norms regularisation, and is inspired by the successful development of LASSO regression in signal

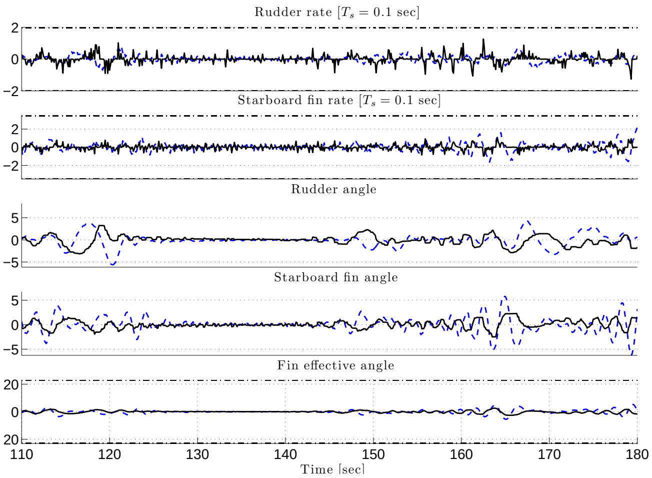

Fig. 6. Comparison in sea state 5 (Δu): Actuators. ℓ_{asso} MPC (solid) $\lambda = 1.5$, quadratic MPC (dashed).

processing. Nominal stability has been proven using standard MPC arguments, and an algorithm has been proposed to implement the dual-mode controller. The problem of vessel course keeping and roll reduction with rudder and fins has been addressed, to demonstrate the capabilities of the new controller. A quadratic MPC has been chosen for comparison.

Simulations have shown that ℓ_{asso} MPC can provide similar control performance to that achieved by quadratic MPC, while providing a very different input behaviour. In particular, penalising the input magnitude causes the actuators to be set at zero for most of the time, and used only when necessary. Small magnitudes seem not to be allowed for the input commands. Similarly, penalising the input time-differences causes the actuators to have a more constant or *piecewise-smooth* behaviour.

ℓ_{asso} MPC seems to be particularly suitable for fully or over-actuated systems, where it is desirable to use the least number of actuators per task, namely to have a smart control allocation policy. Reduced actuators usage could provide several benefits, such as fuel efficiency and increased life-time. Potentially, it could be possible to reduce some feedback-induced negative effects, such as excitation of unmodelled dynamics, or sensitivity to noise, in a way which differs from existing MPC approaches. Finally, the possibility of thresholding the input signals, although still conjectural, is of significant interest. Further investigation is needed, to understand the potential and the limits of sum-of-norms regularisation in the field of feedback control systems.

V. ACKNOWLEDGMENTS

The authors are grateful to Dr Edward Hartley for interesting and useful discussions on the achieved results.

REFERENCES

[1] C. Rao and J. Rawlings, "Linear programming and model predictive contr," *Journal Process Control*, vol. 10, pp. 283–289, 2000.

[2] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. R. Statist. Soc.*, vol. 58, no. 1, p. 267:288, 1996.

[3] M. Osborne and B. Presnell, "A new approach to variable selection in least squares problems," *IMA J. Num. Analysis*, vol. 20, no. 3, p. 389:403, 2000.

[4] S. Kim, K. Koh, M. Lusting, and S. Boyd, "An efficient method for compressed sensing," in *IEEE Int. Conf. Image Proc.*, 2007.

[5] H. Ohlsson, "Regularization for sparseness and smoothness," Ph.D. dissertation, Department of Electrical Engineering, Linkping University, Sweden, 2010.

[6] M. Schmidt, "Graphical model structure learning with l1-regularization, phd thesis," Ph.D. dissertation, Univ. of British Columbia, 2010.

[7] S. Schuet, "Wiring diagnostics via ℓ_1-regularized least sqaures," in *IEEE Sensor J.*, vol. 10, no. 7, 2010.

[8] H. Ohlsson, F. Gustafsson, L. Ljung, and S. Boyd, "Trajectory generation using sum-of-norms regularization," in *Conf. Decision & Control (CDC)*, 2010.

[9] H. Ohlsson, L. Ljung, and S. Boyd, "Segmentation of arx-models using sum-of-norms regularization," *Automatica*, vol. 46, no. 6, p. 1107:1111, 2010.

[10] S. Kim, M. Lusting, S. Boyd, and D. Gorinevsky, "An interior-point method for large-scale ℓ_1-regularised least squares," in *IEEE Journ. of selec. topics in sign. proc.*, vol. 1, no. 4, 2007, p. 606:617.

[11] J. Maciejowski, *Predictive Control with Constraints*, P. Hall, Ed. Prentice Hall, 2002.

[12] P. Scokaert, D. Mayne, and J. Rawlings, "Suboptimal model predictive control (feasibility implies stability)," *IEEE Trans. Automatic Control*, vol. 44, no. 3, pp. 648 –654, mar 1999.

[13] D. Chmielewski and V. Manousiouthakis, "On constrained infinite-time linear quadratic optimal control," *System & Control Letters*, pp. 121–129, 1996.

[14] E. Kerrigan and J. Maciejowski, "Invariant sets for constrained nonlinear discrete-time systems with application to feasibility in model predictive control," in *Proceedings of the 39th IEEE Conference on Decision and Control.*, vol. 5, 2000, p. 4951:4956.

[15] E. Gilbert and K. Tan, "Linear systems with state and control constraints: The theory and application of maximal output admissible sets," *IEEE Trans. Automatic Control*, vol. 36, pp. 1008–1020, 1991.

[16] D. Q. Mayne, J. Rawlings, C. Rao, and P. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, p. 789, 2000.

[17] J. B. Rawlings and D. Q. Mayne, *Model Predictive Control Theory and Design*, L. Nob Hill Pub, Ed. Nob Hill Pub, Llc, 2010.

[18] T. Perez, *Ship Motion Control. Course Keeping and Roll Stabilization Using Rudder and Fins*. Springer, 2005.

[19] T. Perez and G. Goodwin, "Constrained predictive control of ship fin stabilisers to prevent dynamic stall," *Control Engineering Practice*, vol. 16, p. 482:494, 2008.

[20] T. I. Fossen, *Marine Control Systems: Guidance, Navigation and Control of Ships, Rigs and Underwater Vehicles.*, M. C. AS, Ed. Trondheim: Marine Cybernetics AS, 2002.

[21] V. Roth, "The generalized lasso," in *IEEE Trans. Neural Networks*, vol. 15, no. 1, January 2004, p. 16:28.

[22] T. Zhang, "Some sharp performance bounds for least squares regression with l_1 regularization," *The Annals of Statistics*, vol. 37, pp. 2109–2144, 2009.