Model Predictive Control of Low Earth Orbiting Spacecraft with Magneto-torquers

Mark Wood, Wen-Hua Chen, Denis Fertin

Abstract— The problem of attitude control using magnetic torque rods is addressed, in order to demonstrate predictive control as a suitable and effective technique of magnetic attitude control. The study addresses the key issues of magnetic field modeling, controller stability and implementation. Two controller designs are implemented, the first adopting an MPC approach with a constant magnetic field assumption, while the second method includes the true variation of the magnetic field within the control law. Both methods demonstrate significantly improved performance over PD control with the inclusion of the true magnetic field variation leading to the best results. Controller stability is considered with and without terminal penalty within the cost function. Floquet analysis demonstrates both methods to be stable, however the terminal penalty based method leads to a more stable controller.

Index Terms—Model Predictive Control, Magnetic Attitude Control, Time-varying systems

1. INTRODUCTION

The attitude control system (ACS) of a spacecraft or satellite plays a very important role in the success of a given mission. Although differing satellites have varying control requirements, the fundamental reliance on a successful control strategy makes the design and implementation of the ACS very important.

Typically, satellites can be controlled by a number of various actuation methods, including thrusters, reaction wheels, magnetic torque rods, or a combination of the above. Each actuation method has its own advantages and disadvantages and their suitability depends very much on the type of satellite and its overall mission goal.

Magnetic torque rods interact with the Earth's own magnetic field in order to generate a control torque acting on the spacecraft. Due to the nature of the interaction of the magnetic fields, this torque is perpendicular to the local Earth magnetic field direction, and hence the satellite is instantaneously under-actuated. Use of an inclined orbit allows controllability to be ensured in the long term, however this does present a significant challenge from a control perspective.

Use of magnetic torque rods has been studied for both direct attitude control, and for momentum dumping in reaction wheel based control systems. Psiaki (2000) presents a paper considering periodic LQ regulation as a possible approach to fully magnetic attitude control. Psiaki presents a stable control system even in the case of actuator saturation

Wisniewski (1996) presents a detailed PhD study on the application of various control techniques for use in magnetic attitude control. These include linear and non-linear controllers adopting LQR and sliding mode control as well as considering an energy based method. These techniques consider both the attitude regulation problem and the rate detumbling presenting a thorough analysis of the use of magnetic torque rods for full attitude control.

The use of model predictive control (MPC) for magnetic attitude control is less well documented, with very few studies available to the current knowledge of the authors. Krogstad, Gravdahl and Tondel (2005) present an MPC based method for approaching the satellite attitude problem. The design is carried out for a small magnetically actuated satellite and demonstrates significant onboard energy savings in comparison to other control methods. A very significant paper also presented recently is that of Silani and Lovera (2005). Silani and Lovera not only present a thorough review of the area of magnetic attitude control, a derivation of a closed-form solution for the MPC based controller is carried out, allowing realistic implementation without the need for online optimization.

This paper further investigates MPC for attitude control using magnetic torque rods. The main contributions of this paper are threefold. Firstly the effect of magnetic field modeling is considered. Within the MPC scheme, information about the local magnetic field is required. In the simplest case the magnetic field can be measured and assumed constant over any prediction horizon, or alternatively complex prediction models can be used to try

Mark Wood is a PhD research student in the department of Aeronautical and Automotive Engineering at Loughborough University, UK.
(Email M.Wood@lboro.ac.uk)

Dr. Wen-Hua Chen is Lecturer in flight control systems in the Department of Aeronautical and Automotive Engineering at Loughborough University, UK. (Email: W.Chen@lboro.ac.uk)

Denis Fertin is Drag Free and Attitude Control Engineer, Lisapathfinder project, ESA (Email: Denis.Fertin@esa.int)

and improve accuracy. This paper looks at two different approaches to magnetic field modeling and how they affect performance. Secondly the issue of stability is considered. Of the recent publications for MPC for magnetic attitude control, the issue of stability is often omitted, and this paper intends to introduce some initial stability results through use of Floquet analysis. Finally, the focus is also on implementation. To accurately reflect the true performance of the controller, realistic disturbances are included within the simulations, including aerodynamic drag torques, residual magnetic dipole moments, thruster disturbance torques and gravity gradient moments.

In Section 2 the dynamics of magnetically controlled spacecraft are introduced, with a linear model sufficient in a first approach for use in this application, where depointing is limited. Section 3 introduces the model predictive control approach, showing the prediction equations and cost function. Section 4 shows the performance of the predictive control system under simulation including full disturbance model. Section 5 considers stability of the controller with and without terminal penalty, while Section 6 discusses the results found from the various simulations.

2. SPACECRAFT DYNAMICS UNDER MAGNETIC CONTROL

2.1 Linear dynamics model

If linearization is carried out about the equilibrium nadir-pointing attitude, assuming a circular orbit, small Euler angles and deviation of body rates from nominal value, the following linearized model can be produced (Psiaki (2000)). Once a satellite has acquired an earth pointing attitude on orbit (ie one the initial high pointing angle and angular rates have been removed), the satellite dynamics are approximated well by a linear model

$$\dot{x} = Ax + \begin{bmatrix} 0_{3,3} \\ I^{-1} \end{bmatrix} (T_c + T_d) \qquad (1)$$

where,

$$A = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ -4\omega_0^2\sigma_1 & 0 & 0 & 0 & 0 & \omega_0(1-\sigma_1) \\ 0 & 3\omega_0^2\sigma_2 & 0 & 0 & 0 & 0 \\ 0 & 0 & \omega_0^2\sigma_3 & -\omega_0(1+\sigma_3) & 0 & 0 \end{bmatrix}$$

$$x = \begin{bmatrix} \varphi & \theta & \psi & \omega_x & \omega_y & \omega_z \end{bmatrix}^T$$

ϕ, θ, ψ are the spacecraft pointing angles about roll, pitch and yaw axes respectively

$\omega_x, \omega_y, \omega_z$ are the spacecraft angular rates about roll, pitch and yaw axes respectively

ω_0 is the orbital rate, and $\sigma_i = (I_j - I_k)/I_i$ for the index sets $(1,2,3),(3,2,1)$ and $(3,1,2)$

Note that the co-ordinate system used throughout this paper defines the spacecraft orientation relative to a local-level co-ordinate system. The local-level system has the +z axis pointing towards the nadir, the y axis perpendicular to the orbital plane (defined by position and velocity vector), and the x axis defined by the right hand rule.

2.2 Actuator Dynamics

Magnetic torque-rods generate control torques through interaction with the Earth's own magnetic field. This torque is perpendicular to the Earth magnetic field vector (B) and is described below in equation (2).

$$T_c = M \times B \qquad (2)$$

2.3 Disturbance Torque

Due to the low Earth orbit assumed in these investigations, the disturbance torques acting the on the spacecraft can be quite significant. Aerodynamic forces from the upper atmosphere lead to the need for a thruster assembly to maintain near constant velocity and maintain orbit. This leads to disturbance torques not only from the drag forces, but from the corrective thrust applied to the satellite because of thruster misalignment and off positioning. Gravity gradient torques lead to further disturbances, while residual dipole moments further disturb the satellite attitude.

3. MODEL PREDICTIVE CONTROL

3.1 The MPC Approach

MPC is an advanced control strategy based largely around an internal model of the plant under control. This model is used to predict the future expected behavior of the plant, with the resulting information being used to determine the 'optimal' control input to be applied. A typical MPC strategy adopts the following basic procedure.

- An internal model is used to predict the future satellite behavior over a future time period known as the prediction horizon.
- This predicted behavior can be compared to a desired reference trajectory, with an error vector being generated.
- The optimum control sequence is defined by minimizing the predicted error over the prediction

2909

horizon.
- The first input in the 'optimal' control sequence is applied to the plant, with the remaining sequence being discarded and the whole process then repeated at the next sampling interval.

MPC also has the significant advantage of being able to include hard constraints within the problem formulations. With the constraint defined in (2) this makes MPC particularly beneficial for this application.

3.2 Control problem formulation

Consider a linear system described by the discrete-time state space model (3)

$$x(k+1) = Ax(k) + Bu(k) \qquad (3)$$

Iterating equation (3) over a finite future period allows formulation of the expected state vector to be derived in terms of the current state and any future control signal.

$$\mathbf{X} = \begin{bmatrix} A \\ A^2 \\ A^3 \\ \vdots \\ A^P \end{bmatrix} x(k) + \begin{bmatrix} B & 0 & 0 & 0 & 0 \\ AB & B & 0 & 0 & 0 \\ A^2B & AB & B & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ A^{P-1} & A^{P-2}B & A^{P-3}B & \cdots & B \end{bmatrix} \mathbf{U}$$

where $X = \begin{bmatrix} \hat{x}(k+1) & \hat{x}(k+2) & \cdots & \hat{x}(k+P) \end{bmatrix}^T$
and $U = \begin{bmatrix} U(k) & U(k+1) & \cdots & U(k+P-1) \end{bmatrix}^T$

Where X is the predicted future state vector, U is the future control sequence and P is the prediction horizon. In more compact notation this can be illustrated as:

$$\mathbf{X} = \mathbf{Y}x(k) + \mathbf{\Gamma}\mathbf{U} \qquad (4)$$

It should be noted that this analysis assumes that control inputs are calculated at each step over the prediction horizon. This does not need to be the case, and to reduce computational burden the control can be assumed constant after a certain number of time steps (the control horizon).

In order to determine the optimum control sequence defined by U, some optimization process must be carried out. This takes the form of minimization of a cost function illustrated below:

$$\sum_{i=1}^{P} x(k+i)Qx(k+i) + u(k+i-1)Ru(k+i-1) \quad (5)$$

where Q is the state penalty weighting matrix and R is the control weighting matrix.

This optimization process may be carried out however it is very important to note that the control torque U is restricted by equation (2). This means that the available set of control inputs must be within the set of those control torques perpendicular to the local magnetic field vector B. Silani and Lovera present a useful method of including this orthogonality constraint while still achieving a closed-form control law, removing the need for the usual online quadratic programming for constrained systems. Equation (6) presents the solution of (5) in closed-form (Silani and Lovera (2005)).

$$U_{opt} = -\Lambda \left(I - \left(\Psi\Lambda\Psi' \right)^{-1} \Psi\Lambda \right) \Gamma' QYx \qquad (6)$$

where $\Lambda = \left(\Gamma'Q\Gamma + R \right)^{-1}$ and Ψ defines the condition $B \cdot T = 0$ over the prediction horizon.

3.3 Terminal Penalty Weighting

In order to improve controller stability a terminal penalty can be added within the cost function. This effectively modifies equation (5) to the following

$$\sum_{i=1}^{P-1} x(k+i)Qx(k+i) + u(k+i-1)Ru(k+i-1) + \\ x(k+P)Q_T x(k+P) \qquad (7)$$

where Q_T is the terminal penalty weighting

By more heavily penalizing the final output value, the optimization process is encouraged to drive the state outputs to zero, hence stabilizing the control system. This formulation can easily be adopted into the closed-form solution shown in equation (6) by simply modifying the last element of the Q weighting matrix to that of the terminal penalty.

3.4 Magnetic Field Modeling

With the controller carrying out prediction over a finite future period, a suitable model for the variation of the magnetic field vector is required in order to calculate the torque created by a given dipole moment. There are 3 main methods of modeling the magnetic field:

1. Use the measured magnetic field and assume constant
2. Use an onboard prediction model
3. Use a series of magnetic field look-up tables

Method 1 offers a very simple way of modeling the field. Onboard magnetometers measure the current field and this is assumed constant over the prediction horizon. From an implementation point of view, using an onboard prediction model increases the computational burden and hence the

2910

computing power required for the control system. As a result this study considers the assumption of a constant magnetic field, and also the use of a series of magnetic field reference tables, which should offer the same benefits as a prediction model but without the computational burden

4. SIMULATION

4.1 GOCE Satellite

Set for launch within the next year, the GOCE satellite is part of the ESA living planet program. The main aim of the satellite is to measure the Earth's gravity field gradient and hence the control aim is to cancel all non-gravitational angular accelerations acting upon GOCE. The selected orbit is a sun-synchronous, near-polar orbit of 96.5° inclination, at an altitude of 270km. The satellite is controlled entirely by magnetic torque rods. The GOCE mission has been used primarily due to the challenge of high disturbances acting on the satellite, and also as a simulator was available from ESA. Further information regarding the GOCE configuration can be found in Sechi et al (2006).

4.2 Controller Design

The typical MPC structure consists of inputs of current state information, which is then fed into an optimizer in order to determine the control sequence to apply. In this case, the controller is also combined with a feed-forward disturbance observer due to the various expected external disturbances acting on the satellite. As information is available on the current expected controller input and the actual measured acceleration, an estimate of the disturbance can easily be formed for use within a simple feed-forward controller. The structure is summarized below, while the MPC tuning parameters are listed in Table 1.
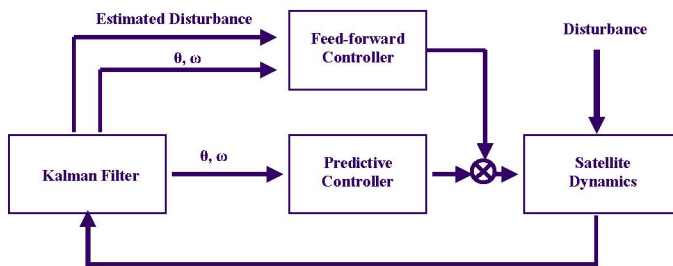

Figure 1 - Controller Structure

Prediction horizon (P)	30
Control horizon	30
Sampling interval	10s
Angular weighting (φ,θ,ψ)	5e-6, 2e-3, 1e-7
Angular rate weighting ($\omega_x,\omega_y,\omega_z$)	106, 0.83, 106
Control weighting (R)	6.6 0.4 1.1

Table 1 – Controller design parameters

4.3 Simulation Results

The following simulations demonstrate the controller performance under a full disturbance model as illustrated in Section 2. The control system is initialized at satellite pointing angles of 1° about each axis and angular rates of 0.0005 deg/s about the roll and yaw axes. Results are presented for both the constant magnetic field assumption MPC and MPC with look-up tables. The results are compared to those initially achieved with PD control in Table 2. Figure 2 shows that the pointing angle time history for both controllers is quite similar. Both comfortably meet the design requirements.

Figure 3 now illustrates the angular rate performance for the two controllers. The pitch and roll response are very similar however there is a marked improvement in the yaw rate performance. Figure 4 illustrate the angular acceleration performance under both schemes and it is here that the main benefits of the look-up table approach can be seen. Acceleration is reduced quite significantly about all three axes and, with this being the most important design parameter for GOCE, this is quite an important finding.
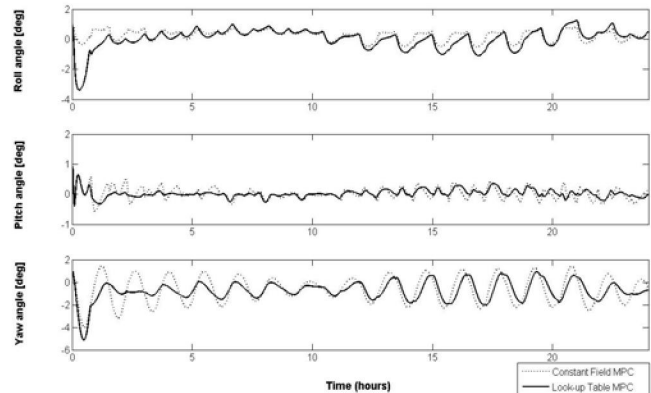

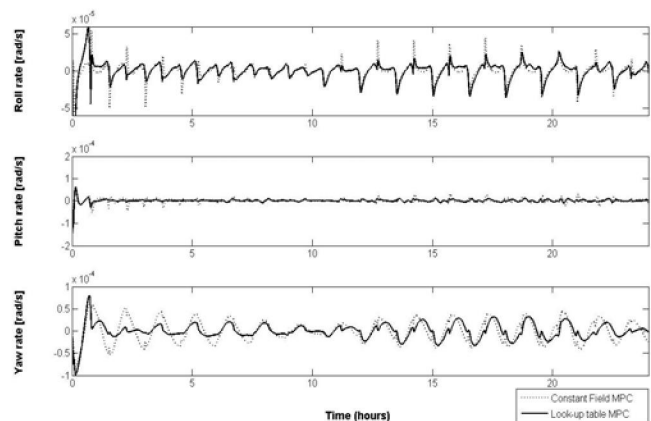Figure 2 – Pointing angle time history for constant field and look-up table MPC

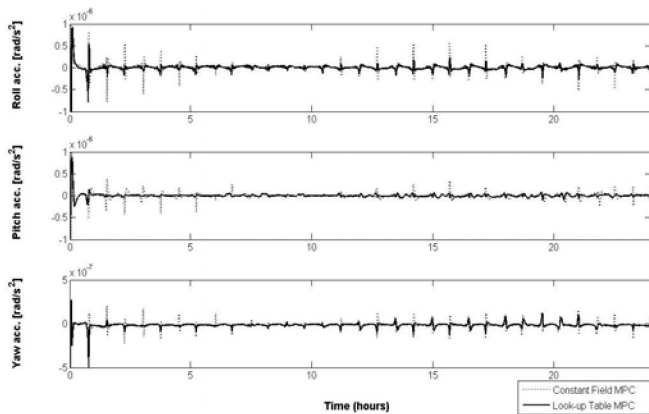Figure 3 – Angular rate time history for constant field and look-up table MPC

2911

Figure 4 – Angular acceleration time history for constant field and look-up table MPC

	Pointing Angle [deg]	Angular Rate [rad/s]	Angular Acc.[rad/s²]
PD Control	2	2×10^{-4}	1×10^{-6}
	1	2×10^{-3}	1×10^{-6}
	5	2×10^{-4}	1×10^{-7}
MPC	1	2×10^{-5}	7×10^{-7}
	<0.1	2×10^{-5}	4×10^{-7}
	2.5	5×10^{-5}	2×10^{-7}
MPC Lookup	1	2×10^{-5}	2×10^{-7}
	0.5	1×10^{-5}	1×10^{-7}
	2	3.5×10^{-5}	1×10^{-7}

Table 2 – Comparison of maximum angle, angular rate, and angular acceleration values for various controllers

Table 2 compares the results achieved from a typical simple PD controller with those from the two MPC designs. Using PD control alone it was very difficult to achieve good performance on all three axes, with the roll axis particularly difficult to tune. The results show that the MPC approaches lead to significantly better performance.

5. STABILITY ANALYSIS

For time invariant systems, the eigenvalues of the closed-loop transfer function allow stability of the system to be defined. In the case of time-varying systems however, the process is more complex. Consider again the discrete-time system described by equation (3). If the closed-loop form of the optimal control input U is substituted into (3), the closed loop dynamics can be described by:

$$x(k+1) = \left\{A - \left[B \quad 0_{6,3(P-1)}\right]\Lambda \left(I - \left(\Psi\Lambda\Psi'\right)^{-1}\Psi\Lambda\right)\Gamma QY \right\}x(k)$$

where the zero matrix is introduced to select just the first input of the control sequence U_{opt}.

For convenience write, $x(k+1) = A_{cl}(k)\hat{x}$ (8)

Consider iteration of equation (8) over one period (T) of the system, the resulting equation can be written as

$$x(k+T) = A_{cl}(k+T-1)A_{cl}(k+T-2)...(k)A_{cl}(k)\hat{x}$$

or

$$x(k+T) = \Phi\hat{x} \tag{9}$$

where Φ is defined as the discrete transition matrix of the closed-loop pseudo-periodic system.

If the eigenvalues of the transition matrix all lie within the unit disc the system can be said to be stable.

5.1 Terminal penalty for stability

The use of a terminal penalty is a well known technique of stabilizing model predictive control schemes. The addition of a penalty within the cost function to penalize the state at the end of the prediction horizon has a tendency to drive the state outputs towards zero and hence stabilize the system.

Figures 5 and 6 show the evolution of the eigenvalues of the transition matrix for the magnetic field look-up table MPC (although both MPC methods have very similar stability characteristics). Although it is the value of the transition matrix eigenvalues after one day that truly indicate stability, the overall evolution of the eigenvalues gives a good indication as to the level of stability of the system (where clearly if the eigenvalues reduce more quickly the system is more stable).

Figure 5 shows the eigenvalues without terminal penalty while figure 6 shows the eigenvalue evolution with a terminal penalty of 500 included. It should be noted that the tuning for the two cases is different due to the effect of the terminal penalty on the overall cost function.
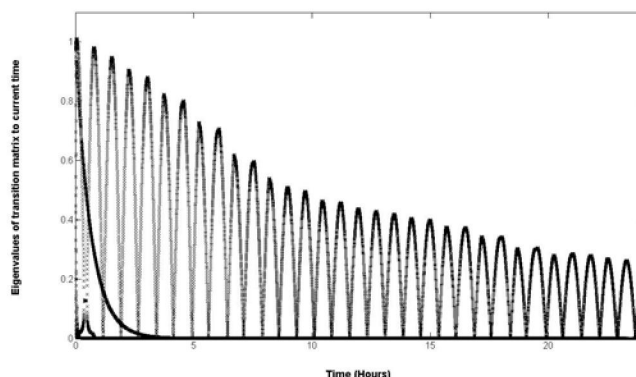

Figure 5 – Evolution of eigenvalues of transition matrix with no terminal penalty
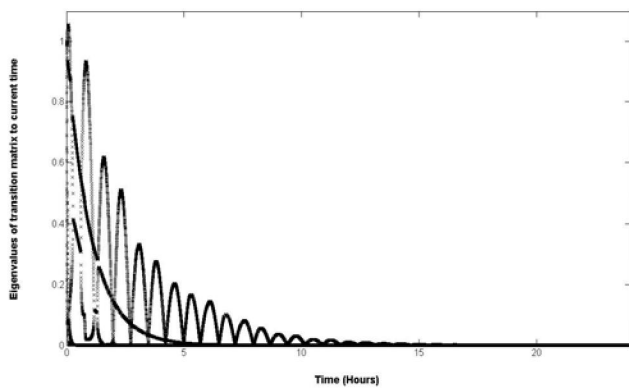
2912

Figure 6 – Evolution of eigenvalues of transition matrix with terminal penalty of 500

As can be seen from the two figures above, the terminal penalty has a significant stabilizing effect on the control system. Although both designs are stable (as the transition matrix eigenvalues are less than 1 after a period of 24 hours), the inclusion of the terminal penalty leads to a much more stable controller, with any initial conditions being removed much more rapidly. This is also reinforced by simulation results, with any initial condition being removed much more quickly by the terminal penalty MPC, with the control system demonstrating significantly higher damping.

6. CONCLUSIONS

From the results shown, the use of MPC for magnetic attitude control can lead to impressive performance in comparison to traditional PD control techniques. The use of a constant magnetic field assumption provides a simple but very useful method of modeling the magnetic field and leads to good results. The inclusion of the true variation in the magnetic field leads to better performance but would rely on accurate magnetic field information.

Primarily the improvements can be seen in the angular rate and angular acceleration performance. When considering the angular accelerations, the fact that use of a more accurate magnetic field model gives improved results is very intuitive when considering the predictive control approach. Considering the points in each orbit where there is weak controllability on one of the spacecraft axes, the control system has to decide what action to take to try and improve the performance on that axis.

In the case of the constant magnetic field assumption, the controller sees weak controllability across the entire prediction horizon, and applies a large dipole moment to try and get even a small influence on the weakly controlled axis. This has the effect of causing large accelerations on the other axes under control.

When using magnetic field look-up tables however, the

controller is aware that controllability will improve in the future and hence the requirement to implement an immediate aggressive control action is no longer necessary. As a result, the control system issues a series of less aggressive control commands and the acceleration performance is seen to be much improved. These smaller accelerations in turn lead to smaller angular rates, which is again reflected in the results shown.

From the issue of stability, the inclusion of a terminal penalty within the problem formulation leads to significant improvement in the overall stability of the control system. The inclusion of the terminal penalty encourages the control system to achieve close to zero state predictions by the end of the prediction horizon. This has the effect of stabilizing the controller and as can be seen from the eigenvalues of the transition matrix, the stabilizing effect is quite significant

The work illustrated here has provided two main developments to the current MPC and magnetic attitude control literature. Use of a more accurate magnetic field model is seen to improve the performance of the control system, with a very noticeable effect on the acceleration performance seen. Secondly, the addition of a terminal penalty within the predictive control scheme is shown to have a very stabilizing effect on the control law. It should be noted that although the transition matrix analysis provides a stability check, it is essentially an "after the fact" check and no guidance can be given for choosing tuning parameters to achieve stability. Research is currently being carried out to develop stability guaranteed MPC.

7. REFERENCES

Krogstad, T., Gravdahl, J., Tondel, P. (2005) Explicit Model Predictive Control of a Satellite with Magnetic Torquers. *In proceedings of the 2005 IEEE international symposium on, Mediterranean conference on, control and automation.*

Psiaki, M. (2000). Magnetic torquer attitude control via asymptotic periodic linear quadratic regulation. *In AIAA guidance, navigation, and control conference*, Denver, Colorado, USA.

Sechi, G., Andre, G., Andreis, D., Saponara, M. Magnetic Attitude of the GOCE Satellite *In Proceedings of the 6ᵗʰ International ESA Conference on Guidance Navigation and Control.* Loutraki, Greece.

Silani, E., Lovera, M. (2005). Magnetic spacecraft attitude control: a survey and some new results. *Control Engineering Practice*, 13, 357-371.

Wisniewski, R. (1996). Satellite attitude control using only electromagnetic actuation. PhD thesis, Aalborg University, Denmark.