Simple adaptive control system design trades

Mooij, E.

DOI
10.2514/6.2017-1502

Publication date
2017

Document Version
Accepted author manuscript

Published in
AIAA Guidance, Navigation, and Control Conference, 2017

AIAA-17-1502

Simple Adaptive Control System Design Trades

Mooij, E.
Delft University of Technology, Delft, The Netherlands

AIAA Guidance, Navigation, and Control Conference
January 9-13, Grapevine, TX

Simple Adaptive Control System Design Trades

E. Mooij*

Delft University of Technology, Faculty of Aerospace Engineering,

Kluyverweg 1, 2629 HS Delft, The Netherlands

In the design of a Model Reference Adaptive Control system, a reference model serves as the (well-known) basis through which system and user requirements can find their way into the design. By tuning the design parameters, the response of the actual vehicle should track the response of the reference model as well as possible. In the current study, the actual vehicle is an unpowered, winged re-entry vehicle similar to the Space Shuttle. To understand the influence of each of the design parameters on the controller performance, a sensitivity analysis is conducted on sample frequency, weighting-matrix coefficients and compensator parameters. Also the influence of varying flight conditions is studied. Normalisation of the reference signals has led to a more unified definition of the design parameters, and has introduced more independence from the flight conditions. Even though the influence of the design parameters is well understood, the large number makes tuning them a tedious process. An automated optimisation procedure should be used, since the interaction between the related design parameters makes an all-in-one-go approach mandatory. Experience of the designer, though, can easily reduce the design time, if a sub-optimal yet satisfactory response is strived for.

Nomenclature

$\mathbf{A}$	system matrix
$\mathbf{B}$	control matrix
$\mathbf{C}$	output matrix
$\mathbf{D}$	direct transmission matrix
$\mathbf{e_y}$	output-error vector, various
F_{u_j}	cumulative standard deviation for output u_j, various
G	transfer function
h	altitude, m
$\mathbf{K}$	control-gain matrix, various
M	Mach number
$M_{T,x,y,z}$	thruster moment (x = roll, y = pitch, z = yaw), Nm
N	number of samples
n_s	subset number of samples
p, q, r	roll, pitch and yaw rate, rad/s
$\bar{q}$	dynamic pressure, N/m^2
$\mathbf{Q}$	state-deviation weighing matrix
$\mathbf{R}$	control-effort weighing matrix
$\mathbf{r}$	concatenated input vector
$\mathbf{T_{p,i}}$	weighting matrix (p = proportional, i = integral)
s	squared deviation from mean value, various
s_0	compensator parameter
t	time, s
$\mathbf{u}$	control vector (p = plant, m = model), various
V	relative velocity, m/s

*Assistant Professor, section Astrodynamics and Space Missions, e.mooij@tudelft.nl, Associate Fellow AIAA.

$\mathbf{x}$	state vector, various
$\mathbf{y}$	output vector, various
α	angle of attack, rad
β	angle of sideslip, rad
γ	flight-path angle, rad
δ	deflection angle (e = elevator, a = aileron, r = rudder), rad
δ	latitude, rad
Σ	integrated performance index
σ	bank angle, rad
σ	standard deviation, various
σ_i	integral-gain filter constant
τ	longitude, rad
χ	heading, rad
Φ	fitness

subscripts

0	sea-level condition
c	commanded
c	compensator
i	integral
m	reference model
p	plant
p	proportional
r	augmented state
u	input related
x	state related
y	output related
max	maximum condition

I. Introduction

For the design of re-usable launch systems, technology development will partly be based on flight experiments by means of experimental re-entry vehicles. Due to the relatively short development time, some uncertainties in flight characteristics will have to be solved during the flight and therefore a robust guidance and control system is required. A possible robust control concept is that of simple adaptive control, SAC.[1] This concept is a member of the greater family of so-called model reference adaptive control (MRAC) and was, amongst others, studied as attitude-control system of large, flexible space structures,[2] robots,[3] a space telescope,[4] a single-stage-to-orbit space plane,[5] and re-entry vehicles.[6,7]

In the design of MRAC systems, a reference model serves as the basis to generate the steering commands for the (unknown) plant. The parameters of the controller are adjusted in such a way that the difference between the model output and the plant output is minimised. The performance of the controller is in this way less sensitive to environmental changes, modelling errors and non-linearities within the system. Although the basic design appears to be simple and straight-forward, the design process relies heavily on the experience of the control engineer to tune the large number of design variables. These design variables mainly originate from the two weighting matrices that are required to compute the proportional and integral components of the adaptive gains. In addition, a number of parameters is needed to tune the potential feed-forward compensators around plant and/or reference model to increase the stability of the controlled plant and to reduce the tracking error, and first-order filters to prevent the integral part of the adaptive gain to reach very large values in non-ideal environments.

Previous work has focused on the evolutionary optimisation of the weighting matrices for the longitudinal control of a winged re-entry vehicle, prior to re-entry.[8] It appeared that the 20 weighting coefficients could be optimised to gain a 30% performance increment in state deviation and control effort with respect to the baseline design. However, due to the non-linear constraints (the two weighting matrices have to be positive

definite and positive semi-definite, respectively) the optimisation algorithm could not always evaluate the performance of the designed control system. Therefore, the current research focuses on the parameter ranges to guarantee a large number of allowable designs. Moreover, the performance sensitivity to variation of each of the weighting coefficients is addressed to possibly reduce the number of design parameters and hence to simplify the optimisation problem.

This paper will concentrate on an application of re-entry control for the longitudinal motion of a winged vehicle similar to the Space Shuttle, although the conclusions can easily be extended to the lateral motion. Due to the large variation in flight parameters (*i.e.,* orbital velocity at an altitude of 122 km, with a large angle of attack of 40°, successive hypersonic descent with decreasing angle of attack and bank reversals, resulting in bank-angle changes of more than 140°), the efficiency of the control actuators (reaction-control thrusters and aerodynamic control surfaces) will vary considerably. Therefore, it is not viable to use one set of weighting matrices and other design parameters for the complete mission. The control-actuator efficiency is addressed by eigenvalue analysis and gain-variation of classical feedback controllers, and the change in weighting matrices is studied for a number of trajectory points. Scaling of design parameters based on variation of dynamic pressure is addressed as well.

The layout of this paper is as follows. Section II describes the simulation model, consisting of the reference vehicle, an unpowered, winged re-entry vehicle, and the reference mission, a re-entry from orbital velocity at 122 km altitude down to Mach 3 at 25 km altitude. Section III provides some background information on simple adaptive control, which is successively used in Sec. IV for the adaptive controller design. In Sec. V, a sensitivity analysis on the design parameters is discussed, whereas in Sec. VI the influence of the large flight regime on these same design parameters is addressed. Section VII, finally, concludes this paper by summarising the design strategy and to give some recommendations for further research.

II. Simulation Model

A. Reference Vehicle

The design and verification of the SAC system is applied to a reference vehicle similar to the Space Shuttle, the HORUS-2B.[9] Initially foreseen as a fully reusable 2$^{\text{nd}}$ stage to the Ariane-5 launcher, HORUS was unpowered, although it had been equipped with a deorbitation engine and attitude-control thrusters. Later on, a rocket engine was added to the design, which then became the manned, 2$^{\text{nd}}$ stage of Sänger, the German TSTO reference concept. In this study, we use the original design being that of a winged, unpowered re-entry vehicle with an entry mass of 26,029 kg (Fig. 1).

For attitude control, HORUS is equipped with both reaction-control thrusters and aerodynamic-control surfaces. The control surfaces comprise two rudders (deflection angles $\delta_{r,l}$ and $\delta_{r,r}$, positive outboard), two elevons ($\delta_{e,l}$ and $\delta_{e,r}$, positive down) and one body flap (δ_b, positive down). The rudders are outward movable only, which means that for yaw control only one rudder is active at a time.

In Fig. 2, the entry control modes for HORUS have been depicted. The reaction control thrusters are operated in the early phase of re-entry, when dynamic pressure is too low to allow for aerodynamic control. Once the dynamic pressure has reached a minimum value of $\bar{q} = 100$ N/m^2, the aileron and elevator (combined into one control surface, the elevon) start working. The rudder, which effectiveness is low, is

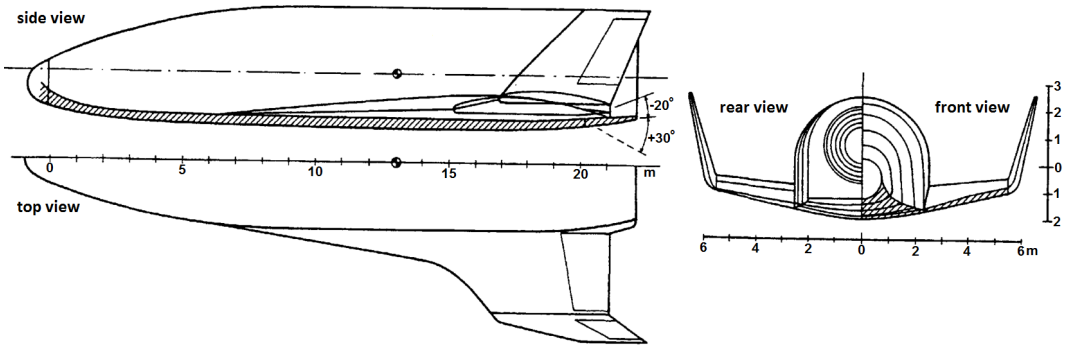

Figure 1. The HORUS-2B reference vehicle.[9]

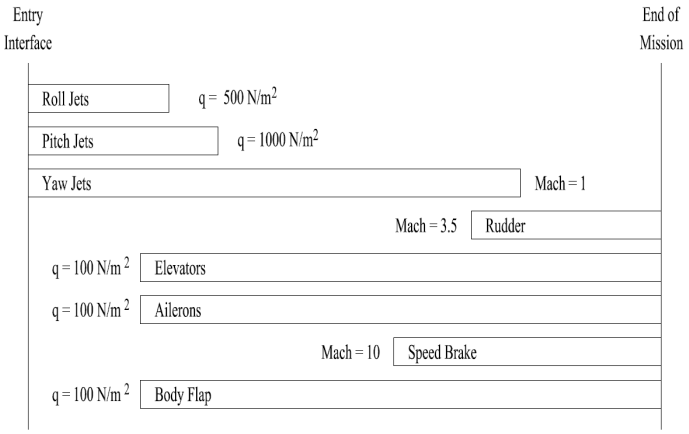

Figure 2. Entry control modes for HORUS.

operational from a dynamic pressure of 150 N/m^2 and above. To aid in yaw control, the yaw thrusters continue to operate down to a Mach number of 1. The body flap is used for trim only. Despite being active from a dynamic pressure in excess of $\bar{q}$ = 100 N/m^2 only, it is set in a non-zero position close to a value it will have when activated ($\delta_{b,0}$ = 15°)[a]. In this way it will contribute to the vehicle's stability, even though less efficient, and it is avoided that a large change in deflection is required when switched on.

B. Reference Mission

The re-entry mission of HORUS starts at the atmospheric entry interface with initial conditions: altitude h = 122 km, longitude τ = -106.7°, and latitude δ = -22.3° for the position, and velocity modulus V = 7435.5 m/s, flight-path angle γ = -1.43°, and heading angle χ = 70.75° for the velocity. The vehicle is assumed to be heading towards a runway in Kourou. Some 80 km from the landing site guidance switches to the Terminal Area Energy Management (TAEM) logic. In Fig. 3, top, the altitude-velocity profile is plotted.

The global task of the hypersonic guidance system is to steer the vehicle along a nominal trajectory and to adjust the steering profile of angle of attack and bank angle (Fig. 3, bottom) to satisfy the flight constraints and to arrive at a suitably defined TAEM interface. Indicated in the figure are the four bank reversals that take place to keep the heading error within bounds. Execution of these reversals is done by

[a]The applied trim law is only active when also the Mach number is lower than M = 20. This will occur at a significantly larger dynamic pressure than $\bar{q}$ = 100 N/m^2.

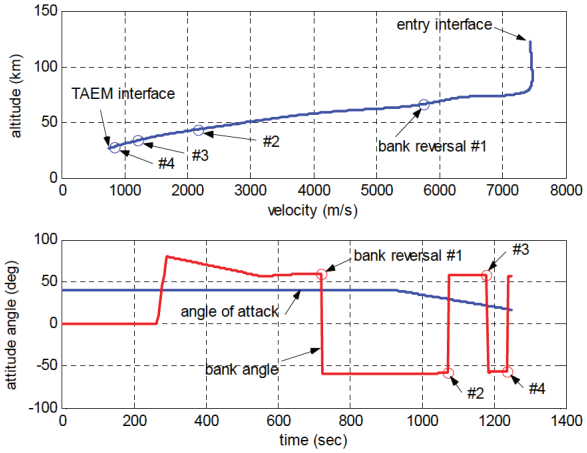

Figure 3. Altitude-velocity profile (top) and nominal control history (bottom).

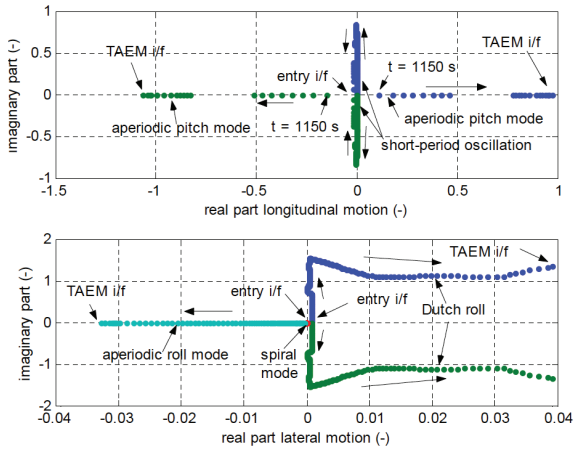

Figure 4. Eigenvalues for hypersonic entry and descent (rotational motion).

the lateral guidance logic, which keeps track of the difference between the current flight heading and the projected heading towards the runway.

To increase understanding the task of designing the control system, in Fig. 4 the eigenvalues of both longitudinal and lateral motion have been plotted. The longitudinal motion is characterised by a short-period oscillation with a Mach-number dependent frequency. Close to t = 1,150 s, the vehicle becomes statically unstable ($C_{m_\alpha} > 0$), and the complex pair of eigenvalues splits into two real ones, one being positive and one being negative. The lateral motion is mainly characterised by the phugoid or Dutch roll, which is marginally unstable at the entry interface but slowly becomes more unstable during the re-entry. From eigenvector inspection, two more eigenmotions can be discerned, *i.e.,* the (almost indifferent) spiral mode and the stable aperiodic roll mode.

III. Simple Adaptive Control

In this section, some background material will be presented on the set-up of an MRAC system, and in particular the class of direct adaptive control algorithms.[1] A favourable approach is that of simple adaptive control based on output feedback, requiring neither full state feedback nor adaptive observers. The algorithm is based on matching the response of the system that is to be controlled (the plant) to that of a reference model (the model).

A schematic overview of such a controller is given in Fig. 5, and the basic adaptive algorithm to compute the plant input $\mathbf{u_p}$ is given by:

$$\mathbf{u_p}(t) = \mathbf{K_r}(t)\mathbf{r}(t) \tag{1}$$

where $\mathbf{r}(t) = [\mathbf{e_y}(t)\ \mathbf{x_m}(t)\ \mathbf{u_m}(t)]^{\mathrm{T}}$ and $\mathbf{K_r}(t) = [\mathbf{K_e}(t)\ \mathbf{K_x}(t)\ \mathbf{K_u}(t)]$. It can be seen that the model input $\mathbf{u_m}$ and model state $\mathbf{x_m}$ are required to form part of the input signal $\mathbf{u_p}$ to the plant. Moreover, the so-called output error $\mathbf{e_y}$ serves as a feedback quantity to form the third element that composes $\mathbf{u_p}$. The three gains, *i.e.,* $\mathbf{K_x}$, $\mathbf{K_u}$ and $\mathbf{K_e}$, are adaptive.

To compute the adaptive gains, $\mathbf{K_r}$ is defined to be the sum of an integral and a proportional component:

$$\mathbf{K}_r(t) = \mathbf{K}_i(t) + \mathbf{K}_p(t) \tag{2}$$

with

$$\dot{\mathbf{K}}_\mathbf{i}(t) = \mathbf{e_y}\mathbf{r}^T(t)\mathbf{T_i} \tag{3}$$

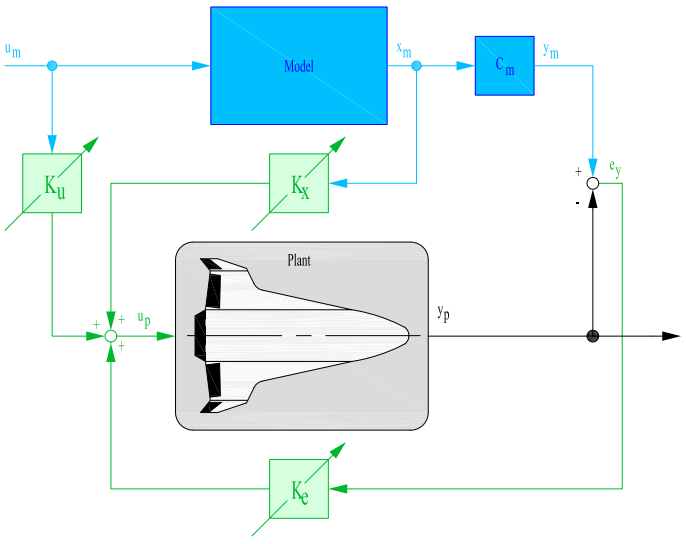

Figure 5. Basic architecture of a Simple Adaptive Control algorithm.

$$\mathbf{K_p}(t) = \mathbf{e_y}(t)\mathbf{r}^T(t)\mathbf{T_p} \tag{4}$$

The weighting matrices $\mathbf{T_p}$ and $\mathbf{T_i}$ are positive semi-definite and positive definite, respectively.

Note that the proportional-gain component has a direct influence on the transient tracking behavior, but is strictly speaking not required to enforce asymptotic tracking. This is guaranteed by the integral gain. To improve the transient response by only using an integral gain, a constant gain value can be added to $\mathbf{K_i}$. An advantage over the use of the proportional gain is that this constant value is independent of $\mathbf{e_y}$, and is therefore non-zero even if $\mathbf{e_y}$ is zero. In that case, the integral gain derived from Eq. (3) becomes

$$\mathbf{K_i}(t) = \mathbf{K_{i,0}} + \int_0^t \dot{\mathbf{K}}_\mathbf{i}(t)dt \tag{5}$$

One way to improve the damping of the system is to include the error derivatives in the output error vector. In that case, the error for output $\mathbf{y}$ becomes:

$$\mathbf{e_y}(t) = \mathbf{K_y^T}(\mathbf{y_m}(t) - \mathbf{y_p}(t)) + \mathbf{K_{\dot{y}}^T}(\dot{\mathbf{y}}_\mathbf{m}(t) - \dot{\mathbf{y}}_\mathbf{p}(t)) \tag{6}$$

with $\mathbf{K_y^T}$ and $\mathbf{K_{\dot{y}}^T}$ being a proportional and derivative output gain, respectively. However, to avoid calculating the numerical derivative of the outputs and to tune the related gains in multiple-output systems an alternative expression for the output error may be used. Adjusting the output matrix by pre-multiplying it with $\mathbf{K_{lqr,c}}$, the optimal control-gain matrix from solving the Algebraic Riccati Equation for a linearised version of the plant, sufficient damping is commonly introduced in the system to have a proper response.[7] The output error becomes in that case

$$\mathbf{e_y} = \mathbf{y_m} - \mathbf{y_p} = \mathbf{K_{lqr,c}}\left(\mathbf{C_m}\mathbf{x_m}(t) - \mathbf{C_p}(\mathbf{x_p}, t)\mathbf{x_p}(t)\right) \tag{7}$$

So far, an ideal environment has been considered. To cope with environmental disturbances that lead to a persistent non-zero error and therefore to a continuous change in the integral gain $\mathbf{K_i}$, a robust design can be applied to adjust the integral gain and preventing it from reaching very high values. The integral term of Eq. (3) is adjusted as follows:

$$\dot{\mathbf{K}}_\mathbf{i} = \mathbf{e_y}(t)\mathbf{r}^T\mathbf{T_i} - \sigma_i\mathbf{K_i}(t) \tag{8}$$

Without the σ_i-term, $\mathbf{K_i}(t)$ is a perfect integrator and may steadily increase (and even diverge) whenever perfect output following is not possible. Including the σ_i-term, $\mathbf{K_i}(t)$ is obtained from a first-order filtering of $\mathbf{e_y}(t)\mathbf{r^T}\mathbf{T_i}$ and, therefore, cannot diverge, unless the output error diverges. It is noted, though, that variation of both σ_i and $\mathbf{K_{i,0}}$ is currently not further studied.

To guarantee that all states and gains in the adaptive system are bounded and the output error is asymptotically stable, it is necessary that the plant is *almost strictly positive real* (ASPR) – or a non-linear plant *almost strictly passive* (ASP). This means, for practical purposes, that the plant can be stabilised by any constant or time-variable high output gain and the plant is strictly minimum-phase. Note that for non-linear systems it is not trivial whether a system is ASP or not,[7] so in some cases the controller is designed ad hoc and validated by extensive simulation.

To apply the adaptive algorithm to a much wider class of systems, various modifications have been developed. One modification is to augment the plant with a so-called feed-forward compensator, such that the class of ASPR systems is increased. If the compensator dynamics is given by the strictly proper transfer function matrix $\mathbf{G_c}(s)$ with state-space realisation

$$\begin{aligned} \dot{\mathbf{x}}_\mathbf{c} &= \mathbf{A_c}\mathbf{x_c}(t) + \mathbf{B_c}\mathbf{u_p}(t) \\ \mathbf{y_c} &= \mathbf{C_c}\mathbf{x_c}(t) \end{aligned} \tag{9}$$

such that the augmented output to be controlled is

$$\mathbf{y_a}(t) = \mathbf{y_p}(t) + \mathbf{y_c}(t) \tag{10}$$

then the augmented system

$$\mathbf{G_a}(s) = \mathbf{G_p}(s) + \mathbf{G_c}(s) \tag{11}$$

is ASPR provided that $\mathbf{G_c}(s)$ is such that the relative degree of $\mathbf{G_a}(s)$ is m (number of inputs), and ${\mathbf{G_c}}^{-1}(s)$ stabilises the closed-loop output feedback system with transfer function $[\mathbf{I} + \mathbf{G_p}(s){\mathbf{G_c}}^{-1}(s)]^{-1}\mathbf{G_p}(s)$.

In general, while using plant feed-forward compensators the model-following error will be bounded but not zero. This can be alleviated by a second extension to the adaptive algorithm, *i.e.,* to apply feed-forward not only around the plant, but also around the reference model.

IV. Adaptive Controller Design

A. General Layout

The general layout of the SAC system is shown in Fig. 6. Five distinct elements can be identified, namely a mission-management system (not part of the SAC system; in this study it has been replaced by a step-function generator), the reference model, the adaptive-gain algorithm, the feed-forward compensators, and the actual plant (HORUS). The *reference model*, which will be discussed in more detail below, is a linearised state-space model of HORUS with decoupled longitudinal and lateral motion. For either motion channel, the system is stabilised by a state-feedback controller. The *adaptive algorithm* is based on Eqs. (1) through (4), and the corresponding weighting matrices will be the subject of the sensitivity analysis later on. The *feedforward compensators* around reference model and plant each have the same basic form, to be discussed in more detail in Sec. V. They take the model control vector, $\mathbf{u_m}$, and plant control vector, $\mathbf{u_p}$, as input.

The three elements of the SAC system, *i.e.,* reference model, adaptive algorithm and compensators are all digitised using a zero-order hold discretisation scheme. To avoid algebraic loops in the loop for the feed-forward compensator[b], as well as for the feed-back control system are removed by adding one-sample delays. For a global, maximum 1% error, the sample frequency should be put to 50 Hz, although slightly larger errors and thus lower frequencies are still acceptable. The influence of the sample frequency of either reference model or adaptive-gain algorithm on the performance will be addressed in Secs. IV.B and V.C.

Concerning the use of the reference model there are two options to consider. The first one is to use the model as part of a tracking system. This means that the model provides a response that aims at removing a guidance error, *i.e.,* the difference between the commanded attitude (α_c and σ_c, with $\beta_c = 0$) and the actual attitude. The reference model will thus be excited with a sampled error and its control system should bring this state deviation (the "guidance error") back to zero. As a consequence, when the reference-model state is back to zero, it no longer provides a signal to the adaptive-gain algorithm and will only be active again in case of non-zero errors. This option has been implemented in, for instance, Refs. 6 and 8.

The second option is to use the actual guidance command, where the deviation from this command by the reference model is used to trigger its control actions. The moment the reference model has achieved the guidance state the deviation has become zero, so no more model control, $\mathbf{u_m}$. However, the difference with the first option is that the model state $\mathbf{x_m}$ is non-zero. As long as $\mathbf{e_y}$ is zero, this has no impact on the gain computation, but when $\mathbf{e_y}$ starts deviating the corrective action due to $\mathbf{x_m}$ is immediately much larger. It is expected that this will improve the tracking performance.

Despite relatively good performance with the first option in previous work, in the current paper the second option is pursued to see the effect on the performance.

B. Reference Model

The performance specifications and/or user requirements will enter the design through the definition of the reference model. Since this model is a simplified representation of the actual system – and more often than not a linearised time-invariant system – it becomes easier to meet with the specifications than it would be for the actual (often non-linear and time varying) system.

Although it is preferable to keep the reference model simple and computationally attractive, there are some points of attention. It must be absolutely clear that the reference model determines the vehicle response to a very large extent, since the adaptive controller will pursue model following. Messer *et al.* found that in case of a large difference between reference model and plant, the maximum control force as well as the

[b]The input to $\mathbf{G_c}$ is $\mathbf{u_p}$, and its output is used to actually calculate $\mathbf{u_p}$.

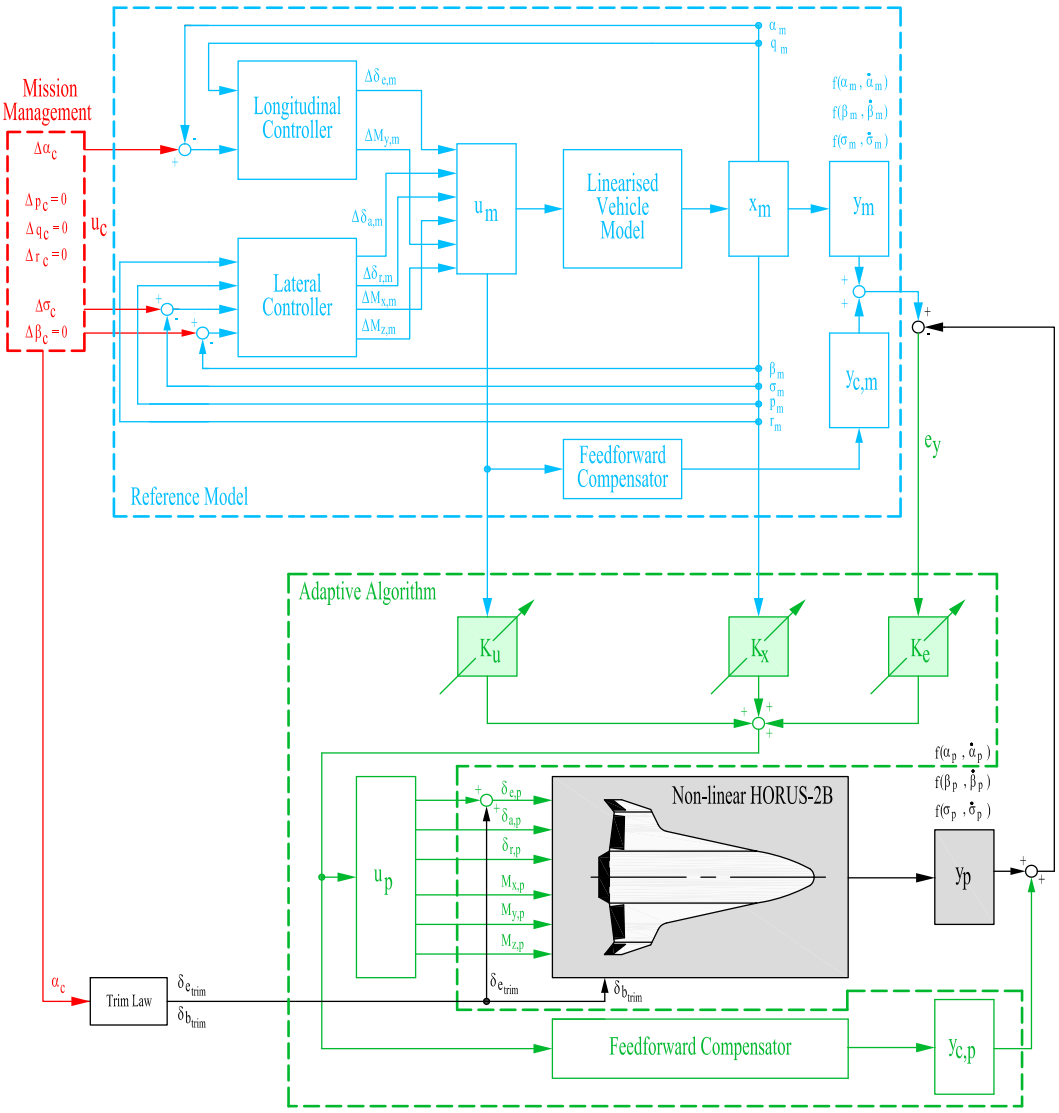

Figure 6. HORUS simple adaptive control system.

control effort increased.[10] It is therefore required to study the plant to some detail, such that a realistic reference model can be derived.

In the current application, we have selected a state-space model of HORUS, linearised around the reference trajectory and stabilised by a simple, linear state-feedback controller. The gains for this controller follow from optimal control theory, which makes this type of controller known as a Linear Quadratic Regulator (LQR), see, for instance, the book by Gopal.[11] The gains are obtained with an indirect method, in which a mathematically defined cost criterion is minimised on an indefinite control time interval:

$$J = \lim_{t_f \to \infty} \int_{t_i}^{t_f} (\Delta \mathbf{x}^{\mathrm{T}} \mathbf{Q} \Delta \mathbf{x} + \Delta \mathbf{u}^{\mathrm{T}} \mathbf{R} \Delta \mathbf{u}) \mathrm{d}t \tag{12}$$

where the term $\Delta \mathbf{x}^{\mathrm{T}} \mathbf{Q} \Delta \mathbf{x}$ represents the error-state deviation and the term $\Delta \mathbf{u}^{\mathrm{T}} \mathbf{R} \Delta \mathbf{u}$ the corrective-control effort. By varying the weighting matrices $\mathbf{Q}$ and $\mathbf{R}$ more importance can be attached to the state deviation, resulting in a faster response, or the control effort, giving smaller control signals. By varying each of the elements of $\mathbf{Q}$ and $\mathbf{R}$ the corresponding elements of $\mathbf{x}$ and $\mathbf{u}$ can be addressed. A good first choice for $\mathbf{Q}$ and $\mathbf{R}$ is given by *Bryson's Rule*:

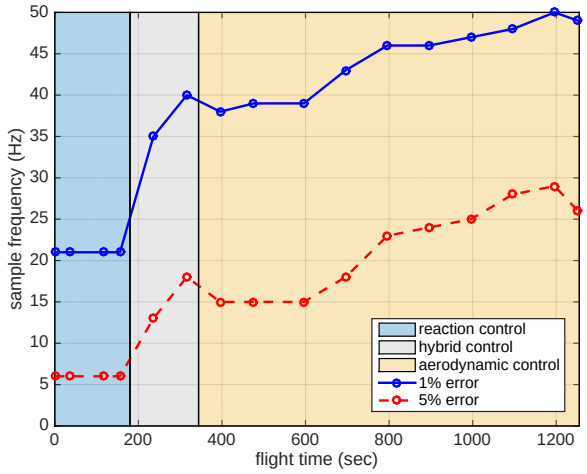

Figure 7. Minimum sample frequency for longitudinal reference model as a function of flight time.

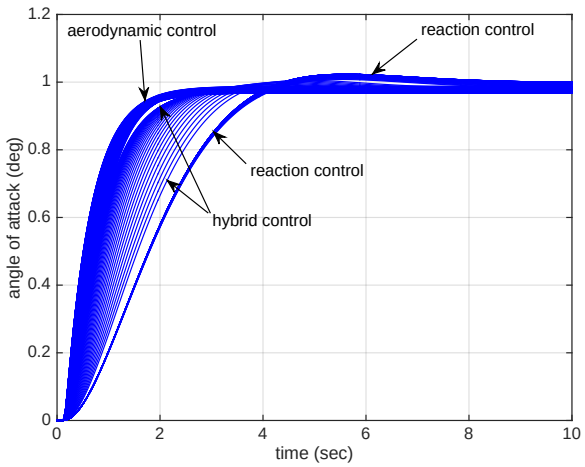

Figure 8. Angle-of-attack step responses for the reference-model along the re-entry trajectory.

$$\mathbf{Q} = \text{diag}\left[\frac{1}{\Delta x_{1,max}^2} \quad \frac{1}{\Delta x_{2,max}^2} \quad \cdots \quad \frac{1}{\Delta x_{n,max}^2}\right] \qquad \mathbf{R} = \text{diag}\left[\frac{1}{\Delta u_{1,max}^2} \quad \frac{1}{\Delta u_{2,max}^2} \quad \cdots \quad \frac{1}{\Delta u_{m,max}^2}\right] \tag{13}$$

where $\Delta x_{i,max}$ and $\Delta u_{j,max}$ are the maximum allowable state deviation and control effort, respectively.

In this study, the gains are computed using constant $\mathbf{Q}$ and $\mathbf{R}$ for each of the control modes during the trajectory. The open-loop state-space model is discretised using a zero-order hold method on the inputs, to keep the propagation of the reference model as simple as possible. Note that in a non-linear simulation, where the flight conditions vary rapidly and significantly, a reference-model update time – and thus also a gain recomputation – in the order of 10-20 s yields good results.

With respect to the sample frequency of the reference-model discretisation and update, in Fig. 7 the relation between sample frequency and accuracy has been plotted. Basis for comparison is a sample frequency of 65 Hz, which is considered to be the truth model. A step response in angle of attack has been simulated, and it is studied for a range of frequencies how much the maximum pitch rate would differ from the reference value at 65 Hz. Plotted are the frequencies for which the error would be 1% and 5%, respectively. Simulations have been repeated for a discrete number of flight times.

It is obvious that in the reaction-control mode a much lower sample frequency can be used. The moment also aerodynamic control is activated, the sample frequency increases, due to the faster dynamics of the closed-loop system, and hence the larger sensitivity to sample frequency. In the aerodynamic-control mode, the frequency keeps on increasing because of increasing dynamic pressure and hence a faster response. A final observation is that for the larger error of 5%, although the sample frequency is lower at the beginning, it increases more rapidly in the second half of the flight. This means that when the system dynamics are fast a small decrease in sample frequency leads to a large degradation of the accuracy.

For a bank-angle response, a similar relation exists. Two differences, however, should be noted. Due to the lower efficiency of the roll thrusters the sample frequency for roll control is lower. On the other hand, the smaller roll inertia coupled with larger controller gains enables a faster aerodynamic response, requiring a larger sample frequency. The maximum frequency at the end of the flight is the same, though, as for the angle-of-attack response.

From the above analysis it is clear that to guarantee a 5% error in reference-model propagation a baseline frequency between 25 and 30 Hz should be taken. A frequency of 25 Hz will be proposed for the design optimisation.

In Fig. 8, the responses to a step in angle of attack have been plotted for a large number of points in the trajectory, each of them separated by 4 s. In principle, three groups of responses can be discerned. The first one for the reaction-control mode, with a small overshoot and a slower response; the second one for the aerodynamic-control mode, characterised by a faster response and no overshoot; and the third region in between the previous two, *i.e.,* the one for hybrid control. Almost every response ends up with a small steady-state error, as is common for simultaneous feedback of α and q.

To obtain these responses, three sets of weighting factors, input to the design of the LQR, have been used: $\Delta\alpha_{max} = 1°$ and $\Delta q_{max} = 1.5°/s$ for the reaction-control mode, $\Delta\alpha_{max} = 1.5°$ and $\Delta q_{max} = 2.5°/s$ for the hybrid-control mode, and $\Delta\alpha_{max} = 2°$ and $\Delta q_{max} = 4°/s$ for the aerodynamic-control mode. Especially for the latter mode, the relaxation of Δq_{max} proved to avoid rapid oscillations of the elevator to keep the pitch rate within bounds. Note that to remove the constraint on q all together resulted in a larger overshoot of the angle of attack, since larger control commands could be issued.

For the lateral controller a similar analysis has been done. Summarising the results, the weighting factors for the reaction control mode are $\Delta\beta_{max} = 1°$, $\Delta\sigma_{max} = 4°$ and $\Delta p_{max} = \Delta r_{max} = 1.5°/s$, for the aerodynamic control mode $\Delta\beta_{max} = 1°$, $\Delta\sigma_{max} = 4°$, $\Delta p_{max} = 10°/s$ and $\Delta r_{max} = 4°/s$, and for the hybrid mode $\Delta\beta_{max} = 1°$, $\Delta\sigma_{max} = 4°$ and $\Delta p_{max} = \Delta r_{max} = 2.5°/s$. Note that the maximum roll rate of $10°/s$ (aerodynamic-control mode) corresponds with the maximum commanded bank rate of $\dot{\sigma}_c = 10°/s$.

It is further noted that only step responses in α and σ have been considered in this analysis, because during the mission of HORUS no sideslip angles other than zero are commanded. Analysing a step response in β (aerodynamic control mode), showed a significant induced σ that could not be controlled to zero simultaneously with β. In a more detailed design, this strong roll-yaw coupling should be further studied. As a possible design solution, an analytical model that decouples the roll and yaw motion, amongst others discussed by Raney and Lallman,[12] could be considered as reference-model controller.

V. Sensitivity to Design Parameters

A. Approach

The performance of a controller, or, expressed in terms from the field of multi-objective optimisation, the fitness Φ_i of an individual design i can be derived from its objective function(s). For the current control-system performance optimisation, we may look at the minimum state deviation of the plant with respect to the guidance commands[c]. Another objective in the design could be to minimise the control effort that is required to influence the plant's behaviour. For instance, in the case of longitudinal control-system design, these two objectives can be expressed as the integrated angle-of-attack deviation and the integrated RCS pitch-thruster activity (the total amount of fuel, or, equivalently, the required control moments), given by:

$$\sum_{\alpha_{err}} = \int_0^t |\alpha_c(t) - \alpha_p(t)|\mathrm{d}t \qquad \sum_{M_{T,y}} = \int_0^t |M_{T,y}(t)|\mathrm{d}t \tag{14}$$

A representation of the above metrics is shown in Fig. 9, represented by the grey areas, for a sequence of step commands in the angle of attack when only reaction control is available. It is obvious that both individual areas should be as small as possible for optimal controller performance, which means they can be used to evaluate different controller designs. In the given example, $\sum_{\alpha_{err}} = 7.221°$s and $\sum_{M_{T,y}} = 44{,}494$ Nm s. It is noted that in this example almost perfect model following was achieved, with $\sum_{\alpha_{err}} = 0.034°$s, in case $\alpha_{err} = \alpha_m - \alpha_p$ would be taken as the error.

To detect oscillations or otherwise discrete changes in the controls, the cumulative moving standard deviation can be used. For a subset j of n_s out of a total of N samples of an arbitrary control signal u, the moving mean is defined as $\bar{y}_j = \frac{1}{n_s}\sum_{i=j}^{j+n_s-1} u_i$. Here, j will run from $j = 1+n_s/2$ to N-$n_s/2$, so each subsequent subset will shift by only one sample. Let the squared deviation from this mean be defined as $s_{u,j} = (u_{j+n_s/2} - \bar{y}_j)^2$, which represents the value at the midpoint of subset j. The cumulative standard deviation, F_u, for subset j is then

$$F_{u_j} = \sqrt{\frac{1}{N - n_s - 1}\sum_{k=1}^{j} s_k} \tag{15}$$

Figure 10 gives two examples of the behaviour of the pitch thrusters, given the command profile of Fig. 9 and considered to be the nominal design case. In the left two figures the pitch-thruster activity is shown

[c]Also the difference with respect to the corresponding reference-model state could be used, assuming that the best control system will enforce an exact model following on the plant

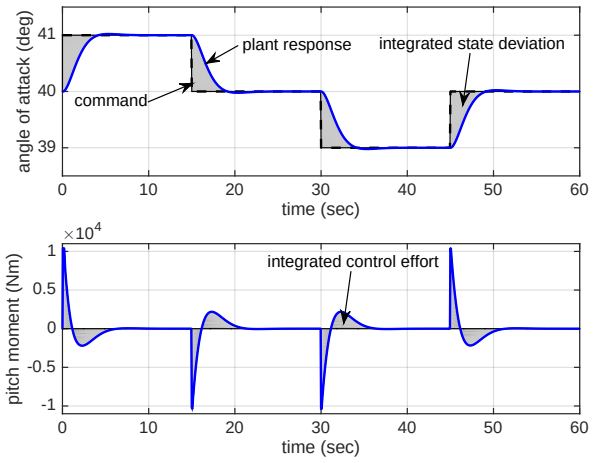

Figure 9. Angle-of-attack response for a reaction-control design point.

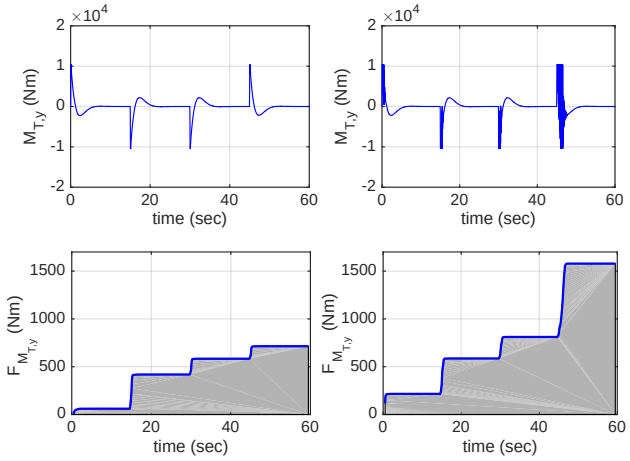

Figure 10. Pitch-thruster commands and cumulative standard deviation for a reaction-control design point: nominal (left) and off-nominal (right) control.

for this nominal response. Due to the discrete changes in attitude command, the pitch thrusters exhibit large jumps in required control moment. These discrete jumps contribute to a rapid increase in $F_{M_{T,y}}$. The corresponding (grey) surface under the curve, $\sum_{F_{M_{T,y}}} = \int_0^t F_{M_{T,y}}(t)\mathrm{d}t$, is a performance measure for these jumps, and should be as small as possible for smoother controls. For an off-nominal design case, where the controller performance has purposely been made worse, saturation periods can be discerned, as well as severe oscillations after $t = 40$ s. Both effects, most notably the oscillations, result in a significant increase of $\sum_{F_{M_{T,y}}}$.

From these results it is obvious that by comparing the numerical values of $\sum_{F_{M_{T,y}}}$, conclusions can be drawn towards the control behaviour, even though it paints a *global* picture only. For the examples shown, these values are $\sum_{F_{M_{T,y}}} = 26{,}244$ Nm s and 46,202 Nm s, respectively. It may be clear that the *gradient* of $F_{M_{T,y}}$ is indicative of *local* oscillations. It is stressed, though, that the two parameters are indicative only, because both the choice of sample interval and the progression of the mean control value have an effect on the actual values. However, these performance metrics will serve their purpose in an automated procedure to optimise the control-system design.

For aerodynamic control, as performance indicators the elevator usage, as well as the integrated cumulative standard deviation,

$$\sum_{\delta_e} = \int_0^t |\delta_e|\mathrm{d}t \qquad \text{and} \qquad \sum_{F_{\delta_e}} = \int_0^t F_{\delta_e}(t)\mathrm{d}t \tag{16}$$

can be added, but also the mean value, $\bar{y}_{\delta_e}$, of the elevator deflection. This indicator may serve as an indicator for a non-zero trim value.

For the lateral control-system design, similar performance indicators are defined for the angle of sideslip β and bank angle σ (state deviation), the roll and yaw moments $M_{T,x}$ and $M_{T,z}$, and aileron and rudder deflection δ_a and δ_r (control effort).

B. Baseline Design

To compare the results, while studying the sensitivity of a design with respect to its design parameters, a baseline design is defined. This design is related to a certain control mode (*i.e.,* flight condition) and is analysed through a (series of) step response(s) in either α or σ. Each of the performance indices, as discussed above, is evaluated and provides a baseline performance. For the reaction-control mode the first trajectory

Pitch (RCS thruster)		
element	T_p	T_i
e_y	0.018	0.007
q_m	$5{\cdot}10^4$	$5{\cdot}10^4$
α_m	$5{\cdot}10^4$	$5{\cdot}10^4$
M_{T,y_m}	$6{\cdot}10^{-6}$	$5{\cdot}10^{-5}$
$\sum_{M_{T,y}}$	46,457 Nms	
$\sum_{\alpha_{c,err}}$	7.3669°s	
$\sum_{\alpha_{m,err}}$	0.2878°s	
$\sum_{F_{M_{T,y}}}$	30,933 Nms	

Table 1. Weighting matrices reaction-control design point ($t = 0$ s).

Pitch (RCS thruster)		
element	T_p	T_i
e_y	1	1
q_m	10	10
α_m	10	10
M_{T,y_m}	5	5
$\sum_{M_{T,y}}$	45,082 Nms	
$\sum_{\alpha_{c,err}}$	7.2768°s	
$\sum_{\alpha_{m,err}}$	0.1419°s	
$\sum_{F_{M_{T,y}}}$	27,787 Nms	

Table 2. Normalised weighting matrices reaction-control design point ($t = 0$ s).

point ($t = 0$ s) is chosen, whereas for the aerodynamic control mode, a point at $t = 1{,}000$ s is selected. The sample frequency of the controller is set to 25 Hz for both control modes.

In Table 1 the corresponding weighting matrices have been listed for pitch-thruster control (no feed-forward compensators have been used). The table shows that there is a large difference between, for instance, the entry for q_m and M_{T,y_m}, and the performance appeared to be quite sensitive to the selection of the weighting matrices. The performance indices for this design point are $\sum_{M_{T,y}} = 46{,}457$ Nms, $\sum_{\alpha_{err}} = 0.2878$°s, and $\sum_{F_{M_{T,y}}} = 30{,}933$ Nms. Even though the response is quite good (comparable to the one shown in Fig. 9), which can be concluded from the small value of $\sum_{\alpha_{err}}$, a small but persistent high-frequency oscillation was observed in $M_{T,y}$ that drove the value of $\sum_{F_{M_{T,y}}}$ to quite a large value.

By inspecting the gain equations (3) and (4) it is obvious that the magnitude of the related elements of $\mathbf{r}$ is directly coupled with the entries in $\mathbf{T_p}$ and $\mathbf{T_i}$. By defining only diagonal matrices it is trivial to meet the positive (semi)definiteness conditions. However, these conditions are harder to meet when also off-diagonal elements are included, especially with the large differences between the diagonal elements. This plays an important role when using numerical optimisation to find the best weighting matrices, because a weighting matrix that is negative definite means a constraint violation. It is therefore better to somehow normalise $\mathbf{T_p}$ and $\mathbf{T_i}$.

Let us assume that a normalisation procedure has been found, such that we can vary each of the weighting coefficients in comparable ranges. To study the effect on the constraint, 5000 implementations with ranges as given in Table 3 are generated with a random generator (Matlab legacy generator, initialised with seed 12345). This will be done for two matrix sizes, *i.e.,* a 4×4 weighting matrix matrix required for longitudinal control and an 8×8 one for lateral control. For each of the matrices, it is determined whether the matrix is positive definite or not. The percentage of constraint violations (negative definite) is listed in the table as well.

It is clear that if we allow a constraint violation in 10% of the cases, such that a new design has to be generated in the optimisation algorithm, the normalised range for the diagonal elements should be in the order of [0,10], whereas the off-diagonal elements may vary in [0,1]. However, if we can narrow down the range for the diagonal elements – basically, this means we should have a baseline design with only diagonal elements and a reasonable performance – the ratio of ranges becomes more favourable. For instance, taking [1,2] and [0,1] for the diagonal and off-diagonal elements, respectively, gives a 4.1% constraint violation, whereas [0.5,2] and [0,0.5] results in only 0.1%. Allowing also negative values for the off-diagonal elements (and thus doubling the range) increases the number of constraint violations significantly, but not so if the diagonal elements are relatively larger.

For the lateral controller, more elements are involved, especially since for yaw control both rudder and yaw thruster are active at the same time. To compare with the previous results, we did a similar analysis for an 8×8 matrix, again for 5000 randomly generated matrices. The obvious conclusion is that the ranges should

normalised range diagonal elements	normalised range off-diagonal elements	percentage negative definite 4×4 matrix	percentage negative definite 8×8 matrix	normalised range off-diagonal elements	percentage negative definite 4×4 matrix
$[0, 1]$	$[0, 1]$	99.1	100.0	$[-1, 1]$	99.5
$[0, 2]$	$[0, 1]$	77.4	99.8	$[-1, 1]$	85.9
$[0, 3]$	$[0, 1]$	52.4	94.7	$[-1, 1]$	63.0
$[0, 4]$	$[0, 1]$	36.7	82.3	$[-1, 1]$	46.2
$[0, 5]$	$[0, 1]$	27.3	70.0	$[-1, 1]$	33.7
$[0, 10]$	$[0, 1]$	10.3	33.0	$[-1, 1]$	12.2
$[0, 25]$	$[0, 1]$	1.9	9.1	$[-1, 1]$	2.3
$[0, 50]$	$[0, 1]$	0.5	3.0	$[-1, 1]$	0.5
$[0, 100]$	$[0, 1]$	0.2	0.7	$[-1, 1]$	0.2
$[1, 2]$	$[0, 1]$	4.1	77.7	$[-1, 1]$	30.1
$[1, 3]$	$[0, 1]$	0.8	26.8	$[-1, 1]$	8.7
$[1, 4]$	$[0, 1]$	0.1	9.0	$[-1, 1]$	3.5
$[1, 5]$	$[0, 1]$	0.1	3.8	$[-1, 1]$	1.6
$[0.5, 2]$	$[0, 1.0]$	39.8	98.0	$[-1.0, 1.0]$	63.3
$[0.5, 2]$	$[0, 0.9]$	27.5	93.8	$[-0.9, 0.9]$	50.9
$[0.5, 2]$	$[0, 0.8]$	15.6	82.8	$[-0.8, 0.8]$	37.5
$[0.5, 2]$	$[0, 0.7]$	7.1	62.9	$[-0.7, 0.7]$	23.2
$[0.5, 2]$	$[0, 0.6]$	2.0	34.3	$[-0.6, 0.6]$	10.7
$[0.5, 2]$	$[0, 0.5]$	0.1	9.0	$[-0.5, 0.5]$	3.5
$[0.5, 2]$	$[0, 0.4]$	0.0	0.3	$[-0.4, 0.4]$	0.5

Table 3. Percentage of constraint violation as a function of normalised weighting-coefficient ranges.

be more well-defined for similar percentages of constraint violation, as compared with the 4×4 matrices.

The actual normalisation of the weighting matrices can be done at two levels. In the first place, the reference vector $\mathbf{r}$, containing elements of the reference-model state $\mathbf{x_m}$ and input vector $\mathbf{u_m}$, as well as the output error $\mathbf{e_y}$ can be scaled, by inspecting the maximum range for each of the related variables. Obviously, since the model response is also flight-time dependent, these maximum ranges may vary. In the second place, also $\mathbf{u_p}$ can be normalised with maximum control effort, *i.e.,* maximum thruster moment or maximum deflection angle. This means that $\mathbf{u_p}$ will be scaled between -1 and +1, and a (constant) gain will be added to convert $\mathbf{u_p}$ to the proper system input.

To begin with the scaling of $\mathbf{r}$, what does that actually mean? According to Eq. (1), the plant input is a gain $\mathbf{K_r}$ times $\mathbf{r}$, with $\mathbf{K_r}$ being decomposed as a proportional and an integral part, see Eq. (2). Further, Eqs. (3) and (4) give the implementation of these two gains. So, summarised:

$$\begin{aligned}
\mathbf{u_p}(t) &= \mathbf{K_r}(t)\mathbf{r}(t) \\
\mathbf{K}_r(t) &= \mathbf{K}_i(t) + \mathbf{K}_p(t) \\
\dot{\mathbf{K}}_\mathbf{i}(t) &= \mathbf{e_y}\mathbf{r}^T(t)\mathbf{T_i} \\
\mathbf{K_p}(t) &= \mathbf{e_y}(t)\mathbf{r}^T(t)\mathbf{T_p}
\end{aligned}$$

Once $\mathbf{r}$ is normalised, the corresponding elements of $\mathbf{T_p}$ and $\mathbf{T_i}$ can be scaled as well, without affecting the result of the multiplication in Eqs. (3) and (4). However, since $\mathbf{e_y}$ is part of $\mathbf{r}$, the gains will change if $\mathbf{e_y}$ is scaled. Moreover, an additional scaling takes place when $\mathbf{u_p}$ is computed according to Eq. (1). This means that the system response becomes different. To keep it simple, we will only scale the elements of $\mathbf{r}$, and when $\mathbf{e_y}$ is required independently, the original values will be used.

The basis for normalising $\mathbf{x_m}$ and $\mathbf{u_m}$ is logically chosen as the input parameters to the LQR design (maximum state deviation and maximum control effort), as the LQR is designed such that for normal operation $\mathbf{x_m}$ and $\mathbf{u_m}$ will remain within these limits. The output error $\mathbf{e_y}$ is determined from the combination of individual state errors multiplied with the LQR gains. The resulting values are thus as large as the corresponding actuator values. It will therefore make sense to scale them with the respective maximum allowable

actuator values, which, for a normal unsaturated control, will yield error values between -1 and 1. Allowing unsaturated values may lead to large values for e_y, especially when discrete steps in guidance commands can be expected. In that case it is better to limit the outcome of the error calculation to the maximum actuator values before scaling them, thus enforcing the range -1 to 1.

To illustrate the effect of normalisation, we will redefine the baseline design for longitudinal reaction control. In Table 2 the adjusted weighting matrices have been listed, in combination with the performance indices. It appeared to be easier to find proper values for the matrix elements that are all of the same order of magnitude, and even the performance improved. However, note that the results listed in Tables 1 and 2 were achieved without optimising $\mathbf{T_p}$ and $\mathbf{T_i}$, so no hard conclusions can be drawn. A similar response has been setup for an aerodynamic-control point (longitudinal motion): weighting matrices and results can be found in Table 4 and Fig. 11. The discrete changes in angle-of-attack command give rise to instantaneous elevator deflections and, combined with the large dynamic pressure, significant pitch accelerations. As a consequence there are relatively large overshoots, which should be reduced with proper weighting-matrix tuning. The sensitivity to varying weighting coefficients will be studied later in this section.

Pitch (elevator)		
element	T_p	T_i
e_y	2	2
q_m	0.2	0.4
α_m	3	4
δ_{e_m}	4	4
$\sum_{\delta_e}$	47.99°s	
$\sum_{\alpha_{c,err}}$	3.31°s	
$\sum_{\alpha_{m,err}}$	1.83°s	
$\sum_{F_{\delta_e}}$	95.81°s	

Table 4. Normalised weighting matrices aerodynamic design point (t = 1000 s).

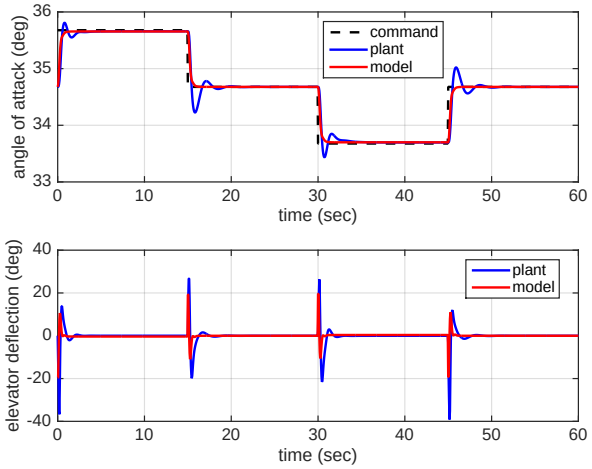

Figure 11. Baseline design for longitudinal aerodynamic control (t = 1000 s).

C. Sample Frequency

Messer *et al.* found that a high sample frequency in the order of 40–80 times the Nyquist frequency was required for a successful realisation of the adaptive controller.[10] In their particular application, frequencies up to 2000 Hz were necessary for an adequate response. Fortunately, the Nyquist frequency in the HORUS application is much smaller due to the physical limitations of manoeuvring speed. Therefore, sampling rates much lower than 100 Hz can be applied, which was already evident from earlier studies.[5,6,8] For the longitudinal controller several simulations with varying sample frequencies are executed. It was found that by simultaneously decreasing the frequency of the reference-model and the adaptive controller, the system became unstable just below the minimum required reference-model frequency (see again Fig. 7). However, by freezing the reference-model frequency at 25 Hz, the adaptive-controller frequency could be decreased to 3 Hz for the reaction-control point (t = 0 s), albeit at the expense of unwanted (yet stable) oscillations due to the bang-bang type of control. This is shown in the left plot of Fig. 12, where the response to a 1°-step in angle of attack is given.

However, by reducing the original weighting matrices $\mathbf{T_p} = \mathbf{T_i} = \text{diag}(1\ 10\ 10\ 5)$ by 99% (!), the response shown in the right plot of Fig. 12 is obtained. A very smooth response is shown for this particular command, albeit it is not guaranteed that such a response can also be obtained for different commands, or in the presence of uncertainties in vehicle and/or environment. What *is* clear, though, is that controller frequency and the settings of $\mathbf{T_p}$ and $\mathbf{T_i}$ are coupled. A further reduction to 2 Hz was also possible for an acceptable response, but now $\mathbf{T_p}$ had to be put to zero, and the e_y-entry in $\mathbf{T_i}$ had to be further reduced by a factor of 10. Note that a zero $\mathbf{T_p}$ is, from a theoretical point of view, still allowed, because it should

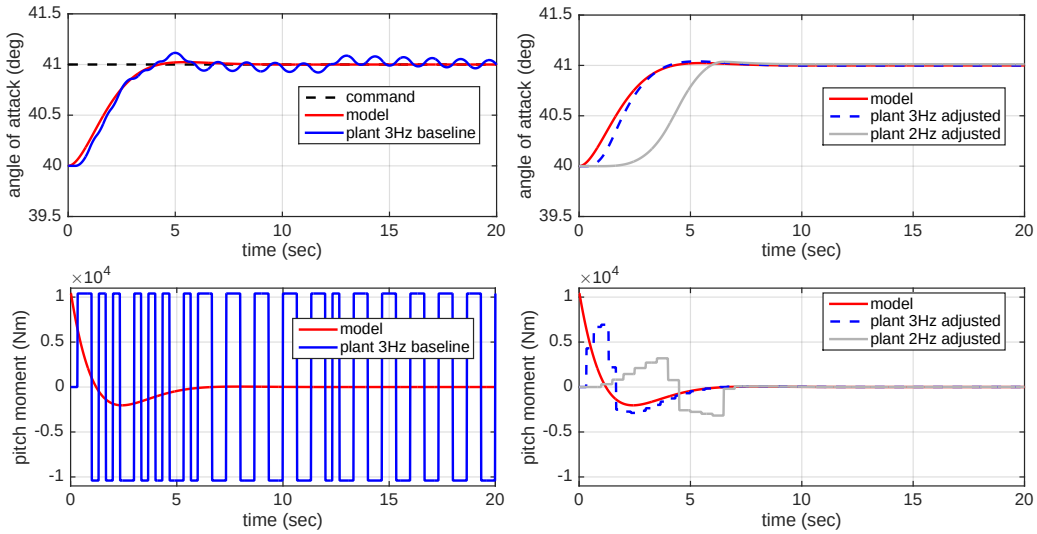

Figure 12. Influence of sample frequency on response characteristics (reaction-control design point).

only be positive semidefinite. Yet, from a control-accuracy point of view, this response is most likely not acceptable.

From the above results, it is concluded that the weighting-matrix optimisation is coupled with sample rate. Common practise in the design of digital controllers is, of course, to define the sample frequency based on on-board computer performance-characteristics and to design the control system accordingly. Only when the controller performance cannot be guaranteed, an increase in frequency might be considered. For the design optimisation, a controller-frequency of 25 Hz is proposed, which is the same as the one for the reference-model sampling.

D. Weighting Matrices

To study the effect of a change of weighting-matrix elements, symmetric matrices are assumed to keep the number of design parameters within limits[d]. For the longitudinal controller (reaction-control mode), let the general structure of each of the weighting matrices $\mathbf{T_p}$ and $\mathbf{T_i}$ be given by

$$\mathbf{T} = \begin{bmatrix} T_{e_y e_y} & T_{e_y q_m} & T_{e_y \alpha_m} & T_{e_y M_y} \\ & T_{q_m q_m} & T_{q_m \alpha_m} & T_{q_m M_y} \\ & & T_{\alpha_m \alpha_m} & T_{\alpha_m M_y} \\ & \text{symm.} & & T_{M_y M_y} \end{bmatrix} \tag{17}$$

With any variation of the individual weighting coefficients, the constraint of a positive definite $\mathbf{T_i}$ and a positive semidefinite $\mathbf{T_p}$ should not be violated. By using the normalised weighting coefficients listed in Table 2 as a starting point, we begin with a ±50% variation of the diagonal elements only (and thus never violating the constraints). For the sequence of step commands, the results obtained with a Monte-Carlo simulation ($N = 500$) are shown in Fig. 13. For this particular design point where the control response is inherently low, the variation in performance indices is not so large, despite the relatively large variation of the design parameters. This means that it is relatively easy to set up a baseline design, even though the results show that a performance improvement is still possible.

To understand which of the design parameters should be tuned first, a separate analysis should be performed. To keep an overview of the performance variation at most two to three variables should be changed at the same time. As an example, in Fig. 14 the performance variation with respect to the baseline, due to a variation in $T_{e_y e_y}$ and $T_{M_y M_y}$ of $\mathbf{T_p}$, is shown. These plots were obtained from a factorial variation of 10% to 200%, with steps of 10% of the nominal values listed in Table 2. Clearly, the non-linear behaviour

[d]This symmetry is not required from a theoretical point of view, which "only" dictates that $\mathbf{T_i}$ is positive definite and $\mathbf{T_p}$ is positive semidefinite. This can obviously also be achieved with asymmetric matrices.

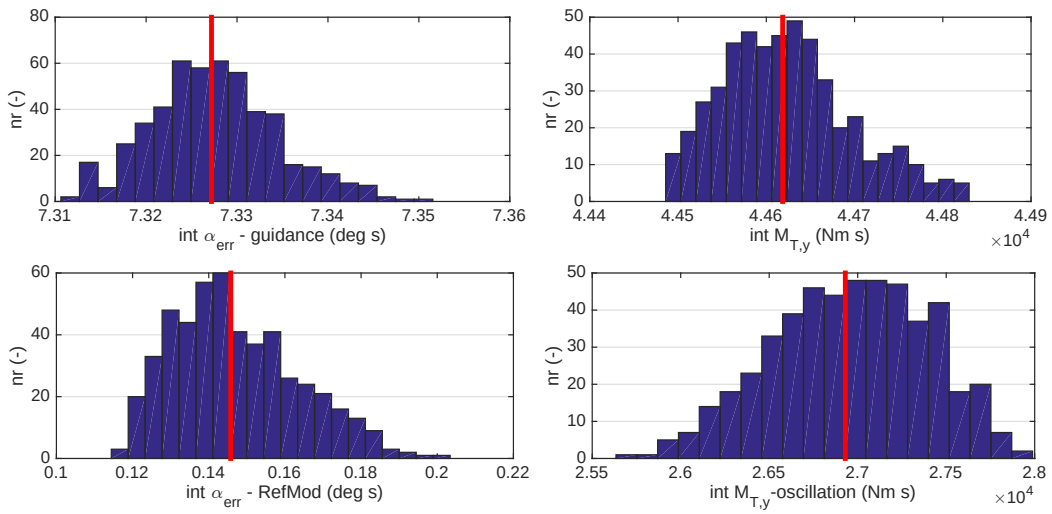

Figure 13. Uniform variation (±50%) of the diagonal elements of $\mathbf{T_p}$ and $\mathbf{T_i}$; reaction-control point, $N = 500$. The red line represents the baseline design.

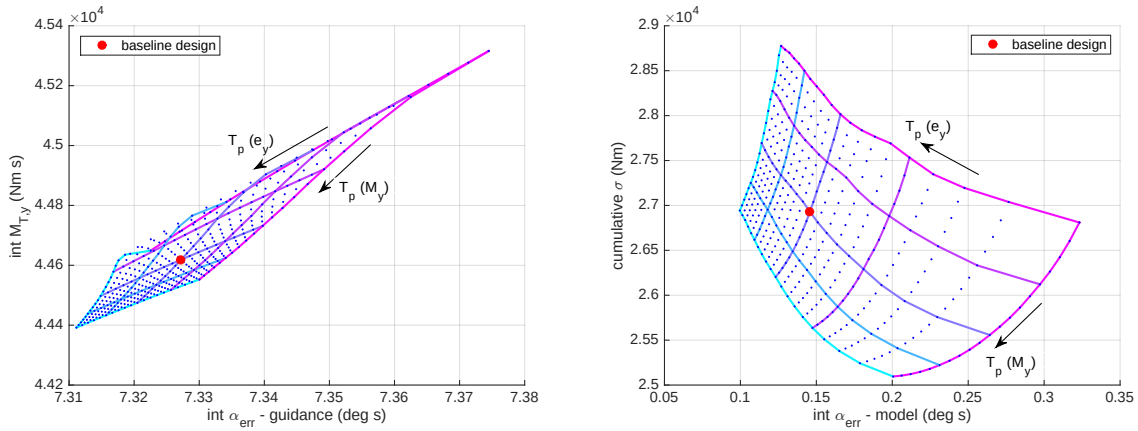

(a) $\sum_{M_{T,y}}$ versus $\sum_{\alpha_{err}}$ (guidance-related)

(b) $\sum_{F_{M_{T,y}}}$ versus $\sum_{\alpha_{err}}$ (model-related)

Figure 14. Influence of varying proportional weighting coefficients on response (reaction-control design point).

shows, as well as the interaction between the two parameters. Apart from a performance improvement in one of the two performance indices, also a combined improvement with respect to the baseline design is possible. It was found that this was easiest achieved by varying the weighting coefficients related to e_y and u_m. The influence of the latter variable is easily understood if one realises that this is some form of feedforward control: the plant should follow the (faster) response of the reference model by actuating the plant in the same way as the model.

The last variation of weighting coefficients that is studied, is one applied to the off-nominal coefficients. Assuming the symmetric matrix format of Eq. (17), the diagonal elements are set at the (fixed) values of Table 2, whereas the off-diagonal elements are varied up to 100% of the corresponding diagonal value (of the same row). Table 3 shows that there may be quite some constraint violations, but in that case a new matrix is generated until it meets the constraints. For a variation of $\mathbf{T_p}$ only, the results are shown in Fig. 15, whereas the results for a combined variation of $\mathbf{T_p}$ and $\mathbf{T_i}$ are given in Fig 16. With off-diagonal elements in either matrix the performance can be improved, usually at the expense of more control effort. The best performance can be obtained when both weighting matrices are tuned, but of course it all depends on the

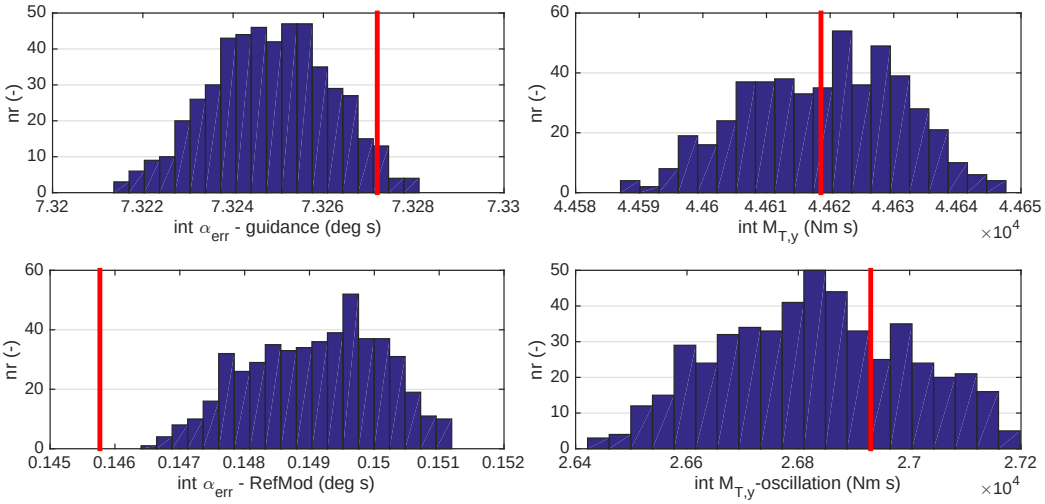

Figure 15. Non-zero off-diagonal elements in $\mathbf{T_p}$ (0-100% of the diagonal element); reaction-control design point, $N = 500$. The red line represents the baseline design.

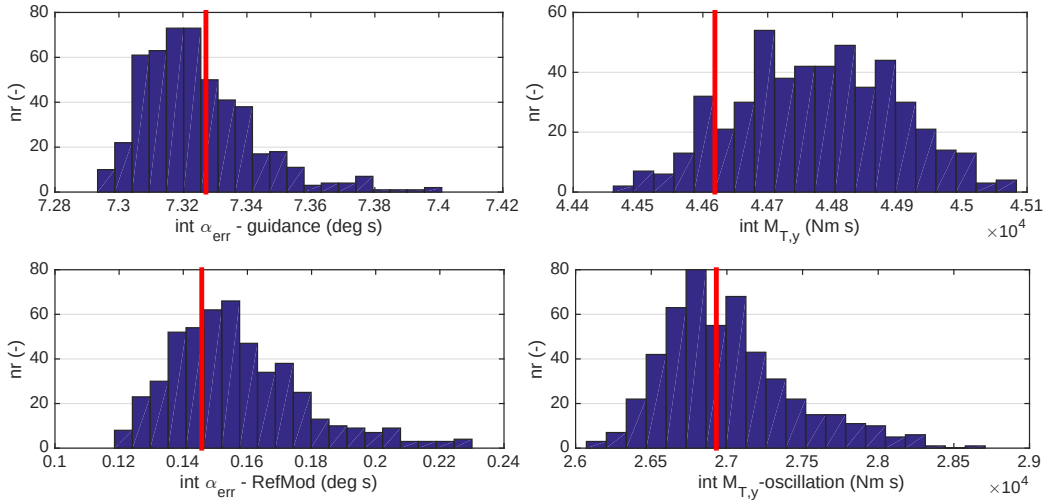

Figure 16. Non-zero off-diagonal elements in both $\mathbf{T_p}$ and $\mathbf{T_i}$ (0-100% of the diagonal element); reaction-control design point, $N = 500$. The red line represents the baseline design.

specific required accuracies to how far one should go. If the required performance can be obtained with diagonal elements only, this is preferred due to the much lower design effort.

To conclude the discussion on the weighting matrices, a Monte-Carlo batch is also executed for the aerodynamic-control design point (Table 4, with the baseline-design response shown in Fig. 11). The relative variation of the off-diagonal elements is the same as before. The four performance indices are shown in Fig. 17 and it is clear that there is more spread in the results than for the reaction-control design point. In particular, the elevator usage as well as the oscillations in the elevator deflection can be reduced significantly, indicating that the baseline design was far from optimal. Since both performance indices for the baseline design are always larger, and the integrated angle-of-attack error is smaller than the baseline for 20-30% of the cases, it means that there are quite some designs that improve on all four indices.

The larger spread of the results – or, in other words, the larger influence of the off-diagonal elements – is attributed to the much larger dynamic pressure and thus the much quicker response to an angle-of-attack command.

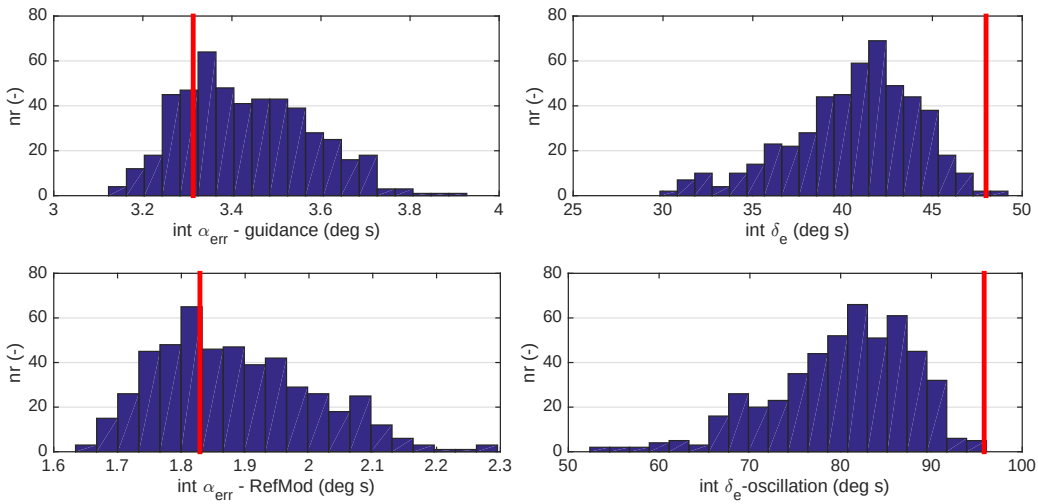

Figure 17. Non-zero off-diagonal elements in both $\mathbf{T_p}$ and $\mathbf{T_i}$ (0-100% of the diagonal element); aerodynamic-control design point, $N = 500$. The red line represents the baseline design.

E. Compensator Parameters

The last category of design parameters comes from the feedforward compensators. As has been stated in Sec. III, to augment the non-linear plant with a so-called feed-forward compensator, the class of ASP(R) systems is increased. A disadvantage is that the model-following error will be bounded, but not zero. This can be alleviated by adding a feed-forward compensator to the reference model as well. In case the system is already ASP(R), feedforward compensators can still be used to improve the performance. The effect of a range of implementations will be addressed in this section.

In this paper, the compensator suggested by Kaufman *et al.*[1] is used; its (generic) transfer function is given by:

$$G_c(s) = \frac{K}{1 + s/s_0} \tag{18}$$

and for the current implementation it is digitised with a zero-order hold on the input. This compensator was selected as the inverse of the output-feedback filter that (marginally) stabilises the plant.

Note that for both the reference model and the plant, the same compensator will be used, apart from the parameters K and s_0. In case compensators would be required for the integrated system (all actuators), an individual compensator is used for each element of $\mathbf{u_m}$ and $\mathbf{u_p}$, respectively. For now, both the reaction-control design point (longitudinal motion, pitch thruster) and the aerodynamic-control design point (elevator) will be analysed. First, the influence of adding a compensator around the plant alone is studied. The gains are defined to be

$$K_y = \frac{M_{T,y_{max}}}{\Delta\alpha_{m_{max}}} \qquad K_e = \frac{\delta_{e_{max}}}{\Delta\alpha_{m_{max}}} \tag{19}$$

where $M_{T,y_{max}} = 10{,}400$ Nm, $\delta_{e_{max}} = 40°$ and $\Delta\alpha_{m_{max}}$ is varied from 0.1° to 2.1° in steps of 0.1°. The parameter s_0 is varied from 0.5 to 10.5 in steps of 0.5. For the reaction-control design point the results are shown in Fig. 18. Adding the compensator will increase the tracking error, both with respect to the guidance signal and the reference model. However, the compensator has a positive effect on the oscillatory behaviour of the pitch thruster, meaning that the actuation has become less abrupt. In about 50% of the cases slightly less fuel will be required by the thrusters.

An additional batch of simulations was done where also a compensator around the reference model was implemented, with the exact same parameters as those for the plant compensator. The responses are now very close to that of the baseline design with little variation, which confirms that at least for this design point the plant output is now so close to the reference-model output, that it appears as if there is no difference with

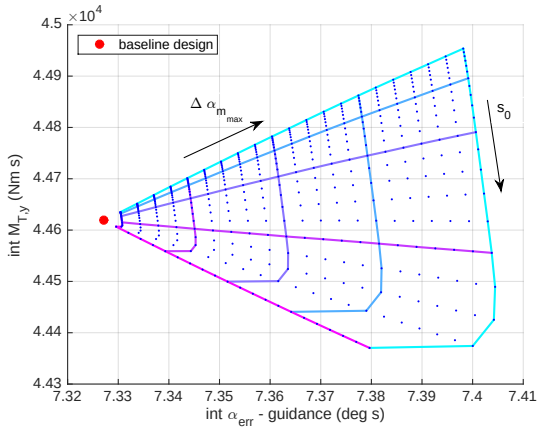

(a) $\sum_{M_{T,y}}$ versus $\sum_{\alpha_{err}}$ (guidance-related)

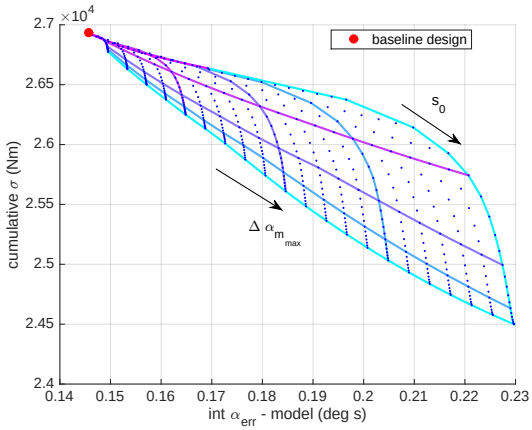

(b) $\sum_{F_{M_{T,y}}}$ versus $\sum_{\alpha_{err}}$ (model-related)

Figure 18. Feedforward compensator around plant: influence of varying parameters on the response (reaction-control design point).

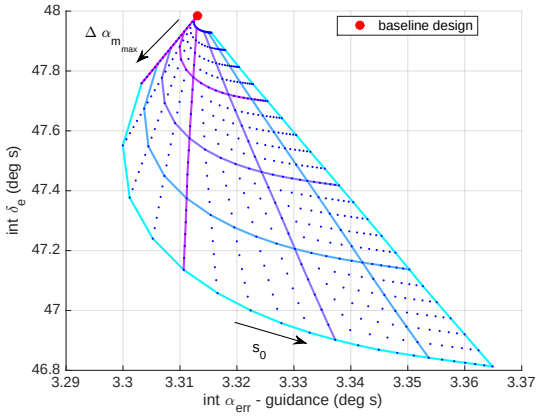

(a) $\sum_{\delta_e}$ versus $\sum_{\alpha_{err}}$ (guidance-related)

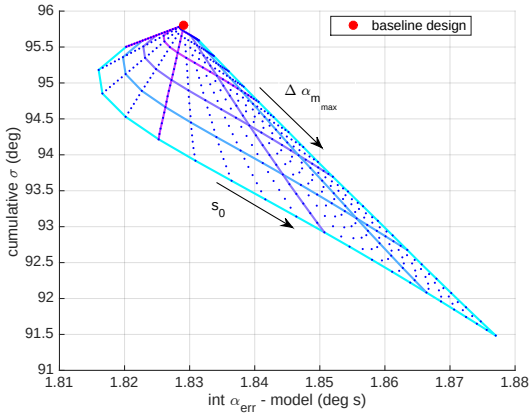

(b) $\sum_{F_{\delta_e}}$ versus $\sum_{\alpha_{err}}$ (model-related)

Figure 19. Feedforward compensator around plant: influence of varying parameters on the response (aerodynamic-control design point).

respect to the baseline design. For this situation it can thus be concluded that if the plant is not ASP(R), it was made so at the expense of a larger model-following error. And this error can be alleviated again by adding the compensator also around the reference model. However, that this is not always the case is shown with the following analysis.

For the aerodynamic-control design point the same variation is applied, *i.e.,* for a feedforward compensator around the plant only, $\Delta\alpha_{m_{max}}$ is varied from 0.1° to 2.1° in steps of 0.1°, whereas s_0 is varied from 0.5 to 10.5 in steps of 0.5. The resulting responses are shown in Fig. 19, again in relation to the baseline design (Table 4 and Fig. 11). Both the integrated elevator deflection and the oscillation index are always lower than those for the baseline design. For about 60-70% of the cases studied this comes at the expense of a small increase of integrated angle-of-attack error. Further – and possibly combined – optimisation of weighting matrices and compensators may lead to the best of both words: a guaranteed ASP(R) and a good performance.

In case feedforward compensators are also implemented around the reference model, with the same design parameters as for the one around the plant, the simulated responses are shown in Fig. 20. The variation

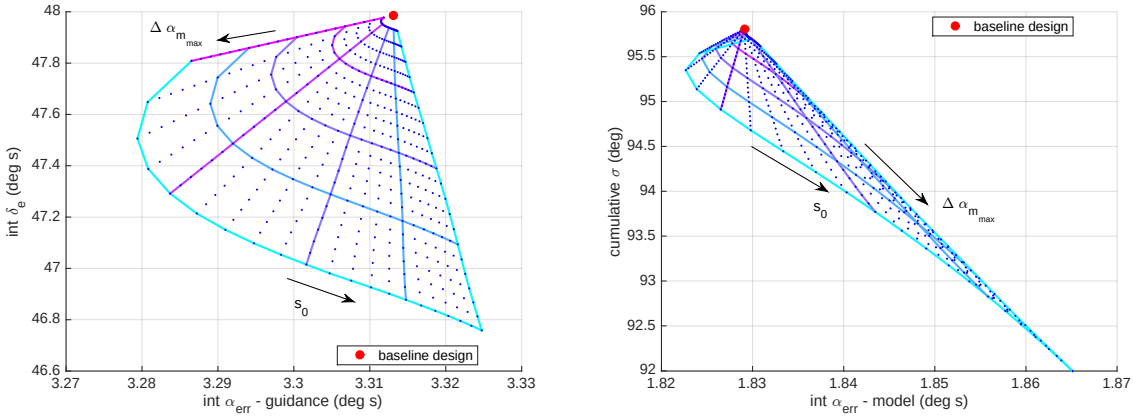

(a) $\sum_{\delta_e}$ versus $\sum_{\alpha_{err}}$ (guidance-related) (b) $\sum_{F_{\delta_e}}$ versus $\sum_{\alpha_{err}}$ (model-related)

Figure 20. Feedforward compensator around plant and model: influence of varying parameters on the response (aerodynamic-control design point).

patterns are similar to those shown in Fig. 19, but the variations are smaller. However, they are still different from the baseline design, *i.e.,* the influence of the combined compensators is clearly visible. It is most likely that this is due to the relatively large dynamic pressure, and thus a more rapid response than with reaction control. In such a case a small change in (augmented) output and difference between $\mathbf{y_p}$ and $\mathbf{y_m}$ can lead to noticeable differences in response. Also in this case, further optimisation of the design parameters will most likely improve each of the performance indices.

F. Lateral Control

Concerning the lateral control, several tests were executed for an initial sideslip error, as well as a bank reversal, similar to the one described in Ref. 6. Following the principle of signal normalisation discussed in Section V.B, a similar set of values is found for the weighting matrices. It is noted, though, that the nominal guidance command for the angle of sideslip is zero, so relatively larger values for the β_m-coefficients are required for a sufficiently fast transient repsonse.

The difference between the longitudinal and lateral controller is further driven by the smaller maximum reaction-control moments (1,600 and 7,600 Nm for the roll and yaw thrusters, respectively, versus 10,400 Nm for the pitch thrusters), and the difference between the respective moments of inertia. One complication that makes the design of the lateral controller more complex than that of the longitudinal controller is the hybrid yaw control, *i.e.,* the combined actuation of the rudders and the yaw thrusters. Without a moment allocator, *i.e.,* the controller commands a yaw moment that is then optimally divided over the two actuators, the system is not ASP(R), because there is no unique relation between the input *to* and the (multiple) output(s) *from* the plant. That means that by definition a feedforward compensator is required, as discussed in the previous section.

However, in summary, the variation of the relevant parameters (weighting matrices, sample frequency, and compensator parameters) shows similar variations as for the longitudinal controller. The analysis can be done in exactly the same way, albeit with a larger effort due to the substantially larger number of design parameters. Unfortunately, due to the space limitations, these results cannot be documented here.

VI. Flight Regime

The flight regime of HORUS is very extreme, as is the case for any entry vehicle returning from orbit. The vehicle enters the atmosphere at 122 km altitude with a velocity relative to the atmosphere of 7.4 km/s, whereas it will land with a velocity common for passenger aircraft. Moreover, to reduce the heating load it will fly at large angle of attack ($\alpha = 40°$), but later on in the flight a lower angle of attack is required

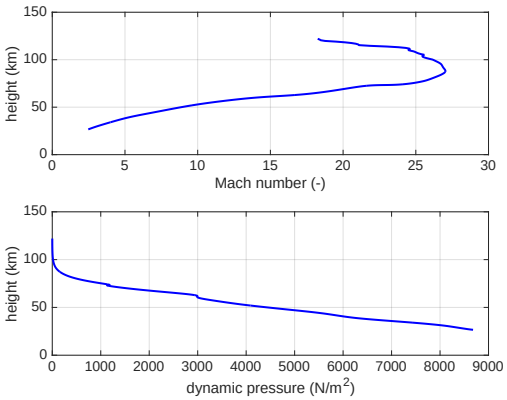

Figure 21. Mach number and dynamic pressure as a function of altitude.

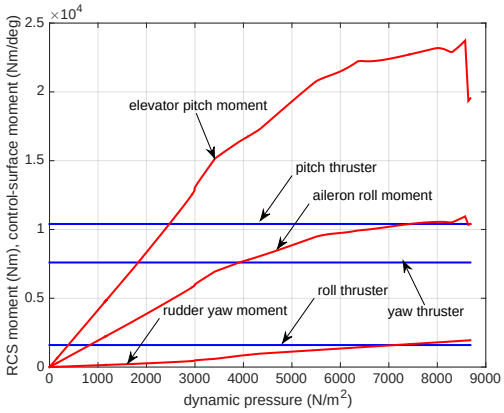

Figure 22. Maximum available moments (RCS thrusters) and moments per degree deflection (control surfaces).

to maximise its cross- and downrange capability. Last but not least, to avoid a skipping flight the lift force must be tilted towards the horizontal plane, but the induced lateral motion needs to be confined by so-called bank reversals, that leads to rolling manoeuvres over more than 130° at a rate of about 10°/s. In Fig. 21, the Mach number has been plotted as a function of altitude, as well as the dynamic pressure. Due to the large variation in operational regime, the aerodynamic properties of the vehicle and control surfaces will vary significantly during the trajectory.

The increase in dynamic pressure will result in an increasingly faster response of the vehicle to a control-surface deflection. To illustrate this, in Fig. 22 the available thruster moments, as well as the control-surface moments per degree deflection (linearised aerodynamics) have been plotted. It can easily be seen that for the elevons, three regions of a more or less linear increase with dynamic pressure can be found. The dip in elevator effectiveness is due to the fact that the elevator trim-deflection suddenly changes. The rudder effectiveness, though low, seems to be completely linear with dynamic pressure. This effect needs somehow to be included in the computation of $\mathbf{u_p}$, to avoid excessive angular rates due to too large deflection angles.

To put the control-surface moments per degree deflection in perspective, suppose we want to trim the vehicle by solely using the elevators (in normal operation, trim is guaranteed by the body flap). Up to a dynamic pressure of about 4,500 N/m^2, a deflection of $\delta_e \approx 30°$ would be required. The deflection would gradually decrease until the vehicle becomes statically unstable ($\delta_e = 0°$, $\bar{q} \approx 7{,}000$ N/m^2). Towards the end of the trajectory, the elevator would have to be deflected at its limit (δ_e = -40°) to guarantee trim equilibrium and would eventually become unstable.

To study the effect of changing flight environment on the controller performance, the gains for the reference-model LQR are inspected. The RCS gains are constant, which points at constant weighting matrices for the related part of the adaptive controller. This is confirmed by analysing the last reaction-control point at t = 184 s. Without changing $\mathbf{T_p}$ and $\mathbf{T_i}$ of the baseline design, the response is similar. Of course, some fine tuning could be done, because of the slightly larger damping effect of dynamic pressure.

With respect to aerodynamic control, some of the gains are constant for a large part (especially the angle-of-attack and pitch-rate gains), whereas some of the lateral-controller gains show larger variation (as a function of dynamic pressure). In particular bank-angle modulation induces dynamic-pressure oscillations, which tends to increase the gain variation.

To conclude the discussion on $\bar{q}$, we will use the baseline design of Table 4 (aerodynamic-control design point, t = 1,000 s) to study an angle-of-attack response at t = 680 s[e]. At this time point, $\bar{q}$ = 1,793 N/m^2 as compared with $\bar{q}$ = 4,401 N/m^2 at t = 1,000 s. The performance of the baseline design is $\sum_{\delta_e} = 47.99°$s, $\sum_{\alpha_{c,err}} = 3.31°$s, $\sum_{\alpha_{m,err}} = 1.83°$s, and $\sum_{F_{\delta_e}} = 95.81°$s.

[e]This point corresponds with a Mach number of 20, the moment that the trim law becomes active. For a lower dynamic pressure (but higher Mach number), the body flap would be fixed at 15°, which gives rise to an additional angle-of-attack error. In that case the comparison with the design point at t = 1,000 s would not be fair.

For $t = 680$ s, the corresponding values are $\sum_{\delta_e} = 70.29°$s, $\sum_{\alpha_{c,err}} = 5.05°$s, $\sum_{\alpha_{m,err}} = 2.70°$s, and $\sum_{F_{\delta_e}} = 62.86°$s. The lower dynamic pressure leads to larger errors, but at the same time increases the control effort due to the lower control-surface effectiveness for this design point. On the other hand, for the same reason the oscillation index reduces (slower thus smoother response).

Suppose we would scale $\mathbf{T_p}$ and $\mathbf{T_i}$ with the ratio of dynamic pressures, *i.e.,* for the point at $t = 680$ s the gain matrices are multiplied by 2.45, which should lead to a faster response. The performance indices have changed to $\sum_{\delta_e} = 67.93°$s, $\sum_{\alpha_{c,err}} = 3.89°$s, $\sum_{\alpha_{m,err}} = 1.55°$s, and $\sum_{F_{\delta_e}} = 96.65°$s. The response has become more similar to the response at $t = 1{,}000$ s, albeit at the expense of a larger control effort. More tuning would be required (as also the original $\mathbf{T_p}$ and $\mathbf{T_i}$ are not optimal), but a scaling with $\bar{q}$ could be used as a starting point.

It is important to realise that the performance of the adaptive controller is not depending on the weighting matrices alone, but also on the reference signals. From Fig. 8 it is obvious that these reference signals vary: as we have mentioned before, at higher dynamic pressures the response (of the reference model) becomes faster. Taking this into account, a more detailed study on how to use constant weighting matrices is required.

VII. Conclusions and Recommendations

In this paper, the sensitivity of a Simple Adaptive Control system with respect to its design parameters is studied. Such a design is centred around the use of a reference model and to avoid a large control effort, it is necessary to study the plant in more detail such that a realistic reference model can be derived.

From the simulation results, it became obvious that the weighting-matrix optimisation is coupled with adaptive-controller sample rate. For the design optimisation, a controller frequency reduction to 25 Hz for all phases is proposed, with the same frequency for the reference model.

Normalisation of the reference signals results in a more consistent range of the elements of the weighting matrices, such that it is easier to guarantee positive (semi)definiteness, while varying the off-diagonal elements. By applying different scaling vectors, the weighting matrices for both reaction and aerodynamic control are of the same order of magnitude. Moreover, after normalisation the performance appeared to be less sensitive to parameter variations.

For flight conditions with low dynamic pressure (reaction-control phase), or in other words, slow dynamics, it suffices to use diagonal weighting matrices. Adding non-zero off-diagonal terms may have a positive effect on the performance, but the effect is limited. For conditions with fast dynamics, *i.e.,* high dynamic pressure (aerodynamic-control phase), adding off-diagonal terms will have a larger effect and could thus be considered during design optimisation.

Besides guaranteeing that the plant is Almost Strictly Passive, the addition of compensators to the baseline SAC design may also have a positive effect on the controller performance, both with respect to the state deviation and the control effort. However, in any optimisation procedure, compensators around both reference model and plant should be taken into account, since both state deviation and control effort can then be reduced simultaneously. Variation of the compensator parameters exhibits a strong non-linear performance behaviour, with strong interactions between these parameters. Because σ_i and $K_{i,0}$ were not included in the parameter variation, it is left as future work to study how a variation of these parameters can contribute to more accurate (and preferably more robust) control.

Due to the analysed parameter variations and identified interactions, it has become clear that in principle all design parameters should be varied in one go. However, it may still be considered to forcibly omit certain off-diagonal elements to simplify the optimisation problem. It should be verified, though, that including these elements will not significantly improve the performance, both with respect to the performance indices and stability of the system in a non-ideal environment.

It remains to be studied in more detail what the influence of the changing flight environment on the different design parameters is. Scaling the weighting matrices with dynamic pressure may lead to a robust approach, but as of now the results are not yet conclusive. The response differences of the reference model should be taken into account in any further study.

Finally, the actual optimisation should be done for a selected number of trajectory points. (Linear) interpolation, coupled with a form of dynamic-pressure scaling may then result in a (consistent set of) controller(s) for the complete re-entry trajectory.

References

[1]Kaufman, H., Barkana, I. and Sobel, K., *Direct adaptive control algorithms: Theory and applications*, Second edition, Springer-Verlag, New York, 1998.

[2]Barkana, I and Kaufman, H., "Simple adaptive control of large flexible space structures", *IEEE Aerospace and Electronic Systems Magazine*, Vol. 29, No. 4, pp. 1137-1149, 1993.

[3]Barkana, I. and Guez, A., "Simple adaptive control for a class of non-linear systems with application to robotics", *International Journal of Control*, Vol. 52, No. 1, 1990.

[4]Mehiel, E.A. and Balas, M.J., "Adaptive Control for a Deployable Optical Telescope", AIAA-04-5222, From: *AIAA Guidance, Navigation, and Control Conference*, Providence, RI, August 16-19, 2004.

[5]Mooij, E., "Numerical investigation of Model Reference Adaptive Control for hypersonic aircraft", *Journal of Guidance, Control and Dynamics*, vol. 24, no. 2, March-April 2001, pp. 315-323.

[6]Mooij, E., "Simple Adaptive Bank-Reversal Control of a winged Re-entry Vehicle ", AIAA-04-4869, *AIAA Guidance, Navigation, and Control Conference*, Providence, RI, August 16-19, 2004.

[7]Mooij, E., "Passivity Analysis for Non-Linear, Non-Stationary Entry Capsules", *Int. J. Adapt. Control Signal Processing*, Published Online 2013.

[8]Mooij, E., "Evolutionary Optimisation of a Direct Model Reference Adaptive Control of a Winged Re-entry Vehicle", AIAA-2002-4922, *AIAA Guidance, Navigation and Control Conference*, Monterey, CA, August 5-8, 2002.

[9]MBB Space Communication and Propulsion Systems Division, "Study on re-entry guidance and control", ESA report reference: ESA CR (P) 2652, 1988.

[10]Messer, R.S., Haftka, R.T. and Cudney, H.H., "Cost of Model Reference Adaptive Control: analysis, experiments, and optimization", *Journal of Guidance, Control and Dynamics*, vol. 17, no. 5, pp. 975-982, Sep.-Oct. 1994.

[11]Gopal, M., *Modern control system theory*, Second Edition, John Wiley & Sons, Inc., Hoboken, NJ, 1993.

[12]Raney, D. and Lallman, F., "Control integration concept for hypersonic cruise-turn maneuvers", NASA Technical Paper 3136, NASA Langley Research Center, 1992.

[13]Vinh, N.X., *Optimal trajectories in atmospheric flight*, Elsevier scientific publishing company, 1981.

[14]Rusnak, I., Weiss, H. and Barkana, I., "Improving the Performance of Existing Missile Autopilot Using Simple Adaptive Control", *Proceedings of the 18th World Congress*, The International Federation of Automatic Control, Milano (Italy) August 28 - September 2, 2011.

[15]Ulrich, S. and De Lafontaine, J., "Development of a Novel Adaptive Control Algorithm for a Fighter Aircraft", AIAA-07-6649, *AIAA Guidance, Navigation, and Control Conference*, Hilton Head, SC, August 20-23, 2007.