21 Optimisation-Based Clearance: the Linear Analysis

Andras Varga

German Aerospace Center
DLR - Oberpfaffenhofen
Institute of Robotics and Mechatronics
D-82234 Wessling, Germany.
Andras.Varga@dlr.de

Summary. We present the clearance results of the HIRM+RIDE control configuration for the linear stability and handling criteria mostly used in the current industrial practice. The performed analysis is based on an optimisation-driven worst-case search. Two classes of linear stability related criteria are considered: the Nichols exclusion region based stability margin criterion and the unstable eigenvalues criterion. The considered handling criteria are the average phase rate and the absolute amplitude criteria. The analysis results clearly illustrate the high potential of the optimisation-based approach in reliably solving clearance problems with many simultaneous uncertain parameters.

21.1 Optimisation-Based Clearance of Linear Criteria

Our analysis addresses the linear stability and handling criteria defined in Chapter 10, namely

(a) the stability margin criterion,
(b) the unstable eigenvalues criterion,
(c) the average phase rate criterion,
(d) the absolute amplitude criterion.

Because of space limitations, we are forced to restrict our presentation to selected results which best illustrate the different aspects of an optimisation-based clearance approach. In this introductory section, we discuss some aspects which are common to all clearance tasks formulated above. Specific aspects and detailed analysis results are presented in separate sections dedicated to particular classes of criteria. The complete results, including results for the classical gridding-based approach, are presented in a GARTEUR AG11 Report [1].

The **definition of suitable distance functions** is a crucial step in an optimisation-based clearance approach (see Chapter 7). The main requirement for a satisfactory distance function is to enforce at minimum, worst cases where the clearance conditions are potentially most violated. Additionally, a satisfactory distance function to be used for an optimisation-driven

worst-case search must not introduce additional local minima, must be continuous and have continuous derivatives, and finally, must be easy to compute. Two categories of clearance problems can be identified, which lead to different types of distance functions. In the first category, for the clearance problems represented by (b) and (d), the clearance criterion is a mathematically expressible function $c(p, FC)$ depending smoothly on the parameters grouped in a vector p and flight condition grouped in a vector FC. Such a criterion usually serves for controller design, and therefore, let us assume that the design results are better for lower values of $c(p, FC)$. If c_o is the limiting acceptable performance level, then as a "distance" function, to be used by a function minimiser, we can define $d(p, FC) = -c(p, FC) + c_0$. By minimising $d(p, FC)$ (i.e., maximising $c(p, FC)$), worst-case parameter/flight conditions are determined. In the second category of clearance problems, represented by (a) and (c), exclusion regions are defined which, for satisfactory performance, must have empty intersection with certain sets of points used for graphical evaluations (e.g., frequency response Nichols-plots or performance plots). The boundary of the exclusion region can be associated with a fictitious limiting performance level c_0 for an appropriate criterion $c(p, FC)$ (to be chosen). If we define the signed "distance" function as $d(p, FC) = c(p, FC) - c_0$ then, by minimising $d(p, FC)$, we can determine worst-case parameter/flight condition combinations. Note that the term "worst-case" is always to be considered in connection with the chosen distance function. For both categories of clearance problems, if for a fixed FC, $d(p, FC)$ is negative for some parameter values of p, then the clearance requirement is not fulfilled in FC and the point FC is *not cleared*. Otherwise, we define FC as *cleared*.

A **fast and reliable criteria evaluation** is another aspect of paramount importance for the effectiveness of an optimisation-based worst case search. Typically, when evaluating criteria on the basis of linearised augmented aircraft models, all criteria evaluations can be done using a unique parameterised nonlinear model which describes the HIRM+ dynamics (see Chapter 8) in a feedback connection with the RIDE controller (see Chapter 9). A particular feature of HIRM+ is that for large values of AoA, a longitudinal/lateral coupling in the aircraft dynamics takes place. Therefore, contrary to the approach used in current industrial practice, we did not use separate linearised models for longitudinal and lateral axis dynamics. In our presentation, the terms "longitudinal axis analysis" and "lateral axis analysis" merely indicate that the analysis has been performed for the uncertain parameter set relevant to either the longitudinal axis or the lateral axis, respectively. As a consequence of using a unique nonlinear model to evaluate the defined clearance criteria (a)-(d), each function evaluation involves trimming and linearisation for a given flight condition and given parameter values. To speed up these computations, the HIRM+ and RIDE are trimmed and linearised separately and the final closed-loop linearised model is constructed by appropriate feedback coupling. The trimming of HIRM+ is done on the basis of special inverse

models, which ensure very fast and accurate trim computations (for details see Chapter 8). For criteria involving frequency-responses, model reduction techniques [2] have been used for further speeding up of computations. This technique is very effective, since the full HIRM+RIDE closed-loop model has order 61, while the reduced order models used to evaluate particular criteria have orders usually bellow 20. The numerically reliable evaluation of criteria means avoiding discontinuities originating from numerical computations (e.g., 360° phase jumps when computing the phase for the frequency response), improper tolerance settings (e.g., large truncation errors), or even failures of function evaluations (e.g., for points outside of the flight envelope).

A key aspect of the optimisation-based approach is the **choice of adequate optimisation software**. Here *adequate* means to employ the best suited software for each clearance task, taking into account the possible existence of multiple local minima, level of noise in function evaluation, possible discontinuities of functions or derivatives etc. Because of their expected higher performance, the gradient-based methods like the sequential quadratic programming (SQP) or projected quasi-Newton (e.g., L-BFGS-B) always represent the first choice. Still, these techniques are not always able to produce the best results, especially when discontinuities in function/derivatives and/or noise in function values are present. Alternatively, the derivative-free linear approximation based trust-region method COBYLA, or the slower but often more robust pattern search (PS) method can be employed. If the presence of multiple local minima is to be expected, global search methods like the genetic algorithm (GA) or simulated annealing (SA) can be employed, either to locate initialisation points for local search based methods or, although expensive, to perform a global search for selected flight conditions.

The analysis results presented in this paper have been obtained for the eight representative flight conditions FC_i, $i = 1, \ldots, 8$ specified in Chapter 10, for values of the angle of attack α, ranging from $-15°$ to $35°$ with a step size of $\Delta\alpha = 1°$. To define the "true" physical flight envelope, a preliminary analysis of the open-loop HIRM+ has been performed for all flight conditions for the nominal values of parameters to check if the conditions

$$-3\,\mathrm{g} \le n_z \le 7\,\mathrm{g} \tag{21.1}$$

are violated or not, and to check if the HIRM+ is trimmable within the allowed limits of the deflections of taileron and rudder actuators given in Chapter 8

$$\begin{aligned} -40° &\le \delta_{TS} + \delta_{TD} \le 10° \\ -40° &\le \delta_{TS} - \delta_{TD} \le 10° \\ -30° &\le \quad \delta_R \quad\; \le 30° \end{aligned} \tag{21.2}$$

The analysis revealed (see [1] and Fig. 21.27) that, because of violation of conditions (21.1) and (21.2), FC_1 is defined only for $\alpha \in [-9°, 35°]$, FC_6 is defined only for $\alpha \in [-9°, 29°]$ and FC_7 is defined only for $\alpha \in [-2°, 12°]$. Violations of conditions (21.1) and (21.2) have been incidentally detected

also for the computed worst-case parameter combinations. In both cases, the corresponding points are automatically not cleared, because they do not belong to the admissible flight envelope. To save computational times, the analysis automatically detects points where the conditions (21.1) and (21.2) are violated for the nominal values of parameters. For such points, no further analysis is performed. Another saving in computational time is possible by skipping further analysis in points where the clearance conditions for the nominal values of parameters are not satisfied. However, this was not done, just to show that even worse results can be computed by an optimisation-based search. It is worth mentioning that, because of the presence of an AoA-limiter at 29° in the RIDE controller, the "true" flight envelope must be probably further restricted to values of AoA $\alpha \leq 29°$, otherwise some analysis results for linear criteria are questionable for values of AoA $\alpha > 29°$.

The analysis of clearance criteria (a)-(d) has been performed for the both the small and the full parameter sets defined in Chapter 10 for both longitudinal and lateral axis analyses. For the longitudinal axis analysis there are 5 parameters in the small (most relevant) parameter set and 9 parameters in the full parameter set, while for the lateral axis analysis there are respectively 6 and 14 parameters in these sets. Complete results for the analysis of all criteria are presented in [1], where results for both optimisation as well as for gridding-based search are given. Due to space restrictions, we present in this contribution only a selection of the most relevant results.

All computations have been performed on a Pentium II 400 MHz machine running MATLAB 5.3 and Simulink 3.0 under Windows NT 4.0. The basis for the performed analysis was the analysis cycle described in Fig. 7.1. For each analysed criterion a corresponding procedure has been implemented as a MATLAB script which performs the whole analysis for all flight conditions, saves intermediary and final results, and evaluates and documents the results through appropriate plots. All scripts allow easy switching between different solvers available in a dedicated optimisation environment.

The employed optimisation software includes the updated SQP and PS software from the RASP library [3], L-BFGS-B using the implementation of [4], the COBYLA software implemented by Powell [5], and a binary coding based GA software adapted from David Carroll's code [6]. The open software architecture in Fig. 21.1 underlies our optimisation-based clearance. It allows us to easily add new solvers and to use the available solvers interchangeably for analysis. A unique MATLAB interface *mex*-function `nlpmex` offers a neutral interface to a generic solver for general nonlinear programming problems (NLPs). The solvers are launched as independent child processes (tasks) which communicate with the parent process (i.e., `nlpmex`) via a problem dependent process communication dialog. For example, each function and/or gradient evaluation involves transferring the current values of optimisation parameters from the solver (i.e., child process) to the *mex*-function (i.e., parent process) which calls the *m*-function typically used to implement the clear-

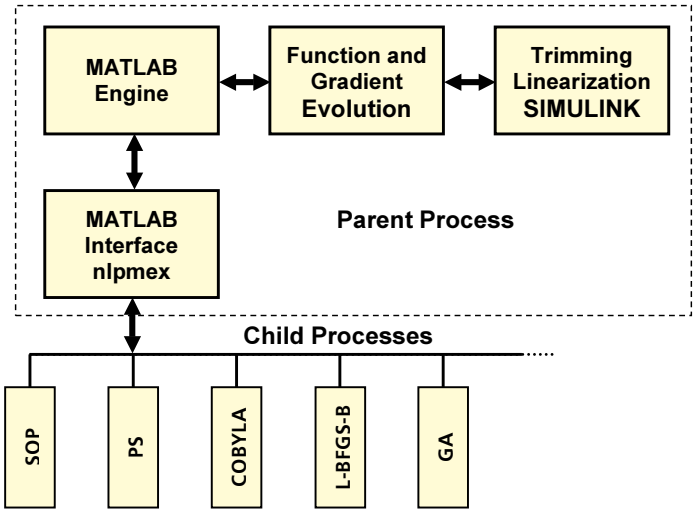

Fig. 21.1. An open software architecture for nonlinear programming

ance criteria. This function usually calls appropriate trimming and linearisation routines (partly based on Simulink). The computed function and/or gradient values are then sent back to the solver. This reverse-communication based software architecture completely separates the function evaluations from the optimisation software. This allows to easily integrate a heterogeneous collection of optimisation tools in a single flexible optimisation environment. Thus, the underlying tools can be written in different programming languages, can have different options and parameter lists, can use different function and gradient evaluation schemes, etc. This open architecture allows us to easily add new solvers as a need arises.

21.2 Results for the Stability Margin Criterion

The goal of the analysis is to identify *all* flight conditions in terms of Mach number M, altitude h, and AoA α, and *all* combinations of uncertain parameters where the Nichols plot stability margin boundaries are most violated. In Chapter 10 requirements for both single-loop as well as multi-loop analysis are defined. In our analysis, we consider only the *single-loop* analysis. The analysis requirement is to check if the open-loop Nichols plot of the frequency response obtained by breaking the loop at the input of the actuators for the symmetric taileron, differential taileron or rudder, avoids the exclusion region delimited by the polygonal boundary B_0 in Fig. 21.2. The analysis for the symmetric taileron loop has been performed only for the longitudinal parameter sets, while for the differential taileron and rudder loops only the lateral parameter sets have been considered.

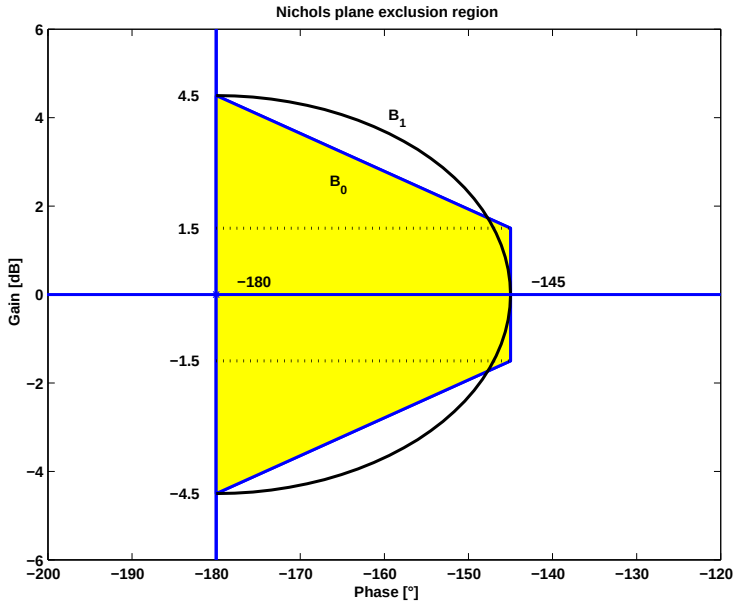

Fig. 21.2. Exact and approximate Nichols plot exclusion regions for gain and phase

To define an appropriate distance function, we approximate the boundary B_0 of the exclusion region in Fig. 21.2 by an arc of an ellipse B_1 defined by

$$\left(\frac{g}{4.5}\right)^2 + \left(\frac{\varphi + 180}{35}\right)^2 = 1$$

where g and φ are the gain and phase, respectively. Instead of using B_0, we can interpret B_1 as the given boundary for the limiting acceptable performance and we can use it to define a smooth distance function for an optimisation-driven worst-case search. Note that defining the distance function with respect to B_0 would lead to a non-differentiable distance function.

For a given frequency dependent gain-phase pair $(g(\omega), \varphi(\omega))$, a normalised signed "distance" to B_1 can be computed by using the formula

$$\tilde{d}(\omega) = \left(\frac{g(\omega)}{4.5}\right)^2 + \left(\frac{\varphi(\omega) + 180}{35}\right)^2 - 1.$$

Note that $\tilde{d}(\omega) > 0$ if the point $(g(\omega), \varphi(\omega))$ lies outside of the ellipse and $\tilde{d}(\omega) \leq 0$ otherwise. The minimum of $\tilde{d}(\omega)$ over a given frequency range $[\omega_{min}, \omega_{max}]$ defines the least distance $d(p, FC)$ to B_1 for a given uncertain parameter vector p and flight condition FC. Note that p and FC are the values used to determine the linearised model which serves for evaluating the frequency response.

The minimum value $\underline{d}(FC)$ of the distance function with respect to the parameters in p can be used to define the following stability margin for flight condition FC

$$\rho_s(FC) = \sqrt{1 + \underline{d}(FC)}. \tag{21.3}$$

With this stability margin, a flight condition FC could be categorised as *not cleared* if $\rho_s(FC) \leq 1$ and *cleared* otherwise. Since ρ_s is defined on the basis of a distance function with respect to the approximate boundary B_1, it could happen that values of ρ_s marginally less than 1 can be still cleared, because, in reality, no intersection with the Nichols exclusion region occur. Conversely, values of ρ_s marginally greater than 1, can be categorised as not cleared, because the Nichols plot intersects the exclusion region. This is why, we used the stability margin defined in (21.3) only for the visualisation of our computational results. Still, all reported clearance results rely on strict checks of intersection/no intersection of worst-case Nichols plots with the exclusion region defined by the boundary B_0 in Fig. 21.2.

For the success of an optimisation-based search, the function evaluations must be as fast as possible. There are two critical, time demanding computations to evaluate $d(p, FC)$, for given p and FC. The first critical computation is the trimming of the HIRM+. A fast trimming for HIRM+ is done using the inverse model approach, as described in Chapter 8 and we will not further discuss this aspect here. The second expensive computation in evaluating the distance function is the computation of frequency responses for the 57-th order (4 state components omitted) closed-loop SISO models formed from the linearised HIRM+ and RIDE. Taking into account the need to compute the minimum distance accurately, the number of frequency points to be used must be sufficiently high. For the optimisation-driven search, we employed 100 frequency values ranging logarithmically in the interval $[10^{-2}, 10]$.

Some timing results obtained by evaluating the distance function for the differential taileron loop on a 400 MHz Pentium II machine are interesting to keep in mind. The time necessary to trim, linearise and build the closed-loop SISO evaluation model is about 1.11 seconds, from which 0.06 seconds account for trimming HIRM+ using the inverse model, 0.13 seconds for the linearisation of HIRM+, and 0.9 seconds for the linearisation of the RIDE controller. Note that although the order of the controller is 9, about 90% of the time is necessary to linearise it using a SIMULINK model. The total times to evaluate $d(p, FC)$ for 100 and 1000 frequency values are 1.45 seconds and 3.9 seconds, respectively, of which, the times necessary to compute only the frequency responses are 0.36 seconds and 2.85 seconds, respectively. The times for frequency response computations can be significantly reduced, by observing that the resulting SISO closed-loop systems used were always nonminimal. To speed up the computations, a preliminary order reduction is performed by using the recently developed high quality model reduction tools [2]. The resulting minimal order models for the symmetric taileron, differential taileron and rudder loops have dimensions 14, 24, and 22, respectively. For

the differential taileron loop, the model reduction requires only 0.07 seconds, but has the effect of reducing the time to evaluate the frequency responses for 100 and 1000 points to 0.06 seconds and 0.61 seconds, respectively. Thus, the total times to evaluate $d(p, FC)$ for 100 and 1000 frequency points become 1.15 seconds and 1.7 seconds, respectively, with performance gains of 20% and 50% in these two cases. Note that the time saving is more substantial for a larger number of frequency values.

One main goal of our analysis was to compare the necessary computational efforts for the gridding-based approach and for optimisation-driven worst-case search methods. The number of maximum function evaluations required by the gridding-based approach for the 8×51 flight conditions can be easily computed for the small parameter sets for longitudinal and lateral axis analyses. The longitudinal axis analysis requires at most $2^5 \times 8 \times 51 = 13056$ function evaluations while the lateral axis analysis requires twice as many $2^6 \times 8 \times 51 = 26112$. About 10%–20% of function evaluations can be saved, by restricting the search only to those points which belong to the "true" flight envelope. In Table 21.1 we present some timing results for the performed analysis of the three loops on the small parameter sets: the symmetric taileron loop (STL) for the longitudinal axis parameters, and the differential taileron loop (DTL) and rudder loop (RL) for the lateral axis parameters. The solvers implementing the local search methods SQP, PS, COBYLA and L-BFGS-B (see Section 21.1) have been used for analysis, with the accuracy tolerance set to 10^{-5}. Table 21.1 also includes results for the classical gridding-based approach.

Table 21.1. Timing results for gridding and optimisation-based analysis

Method	Gridding	SQP	PS	COBYLA	L-BFGS-B
Times for STL (sec)	12943	19796	38188	20840	19426
Times for DTL (sec)	25074	23064	45333	38973	24024
Times for RL (sec)	27087	22374	41808	61657	20977

The time for the gridding-based solution for 5 parameters is always less than the times required by the optimisation-based search. However, this is not the case for 6 parameters, where gradient based methods like SQP or L-BFGS-B are more efficient than the gridding-based approach. The computational effort increases exponentially with the number of parameters and therefore a gridding-based solution is **not affordable** for the full sets of 9 longitudinal and 15 lateral axis parameters due to very high computational costs ($2^4 = 16$ and $2^9 \approx 500$ larger times, respectively). By using the optimisation-based approach it was possible to obtain complete clearance results for the full parameter sets. Table 21.2 summarises the number of function evaluations (NFE) and the computational times required by using the SQP method.

Table 21.2. Performance results for stability margin analysis

	Gridding		SQP (Small sets)		SQP (Full sets)	
	NFE	Time (sec)	NFE	Time (sec)	NFE	Time (sec)
STL	11382	12943	17519	19796	38649	42662
DTL	22710	25074	19220	23064	49596	55081
RL	23088	27087	18645	22374	43693	50903

Detailed analysis results for both gridding and optimisation-based search, for all three loops, are presented in [1]. The results for the differential taileron loop are summarised in Fig. 21.3, where we give, in each flight condition, the domains of α-values which are considered as "cleared" for three cases: (1) gridding-based analysis; (2) optimisation-based analysis for the small parameter set; and (3) optimisation-based analysis for the full parameter set. While for the small parameter set both gridding and optimisation-based search produce almost the same results in term of not cleared flight conditions, the optimisation-based search performed on the full parameter set revealed many additional not-cleared points. This clearly illustrates the power of the optimisation-based approach to simultaneously address many uncertain parameters.

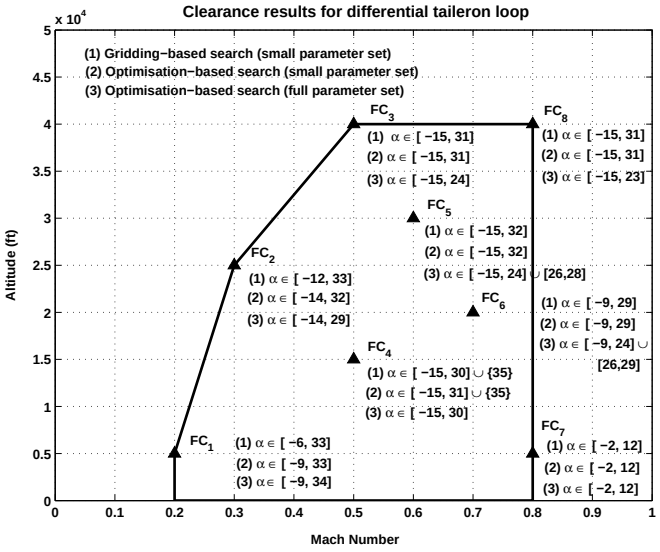

Fig. 21.3. Clearance results for the differential taileron loop

Fig. 21.4 shows the values of the stability degree ρ_s versus α for the small parameter set. These values are visibly greater than the corresponding

values of ρ_s for the full parameter set in Fig. 21.5. Thus, as expected, the HIRM+RIDE configuration has less stability margin in most of the points when more uncertain parameters are allowed to simultaneously vary.

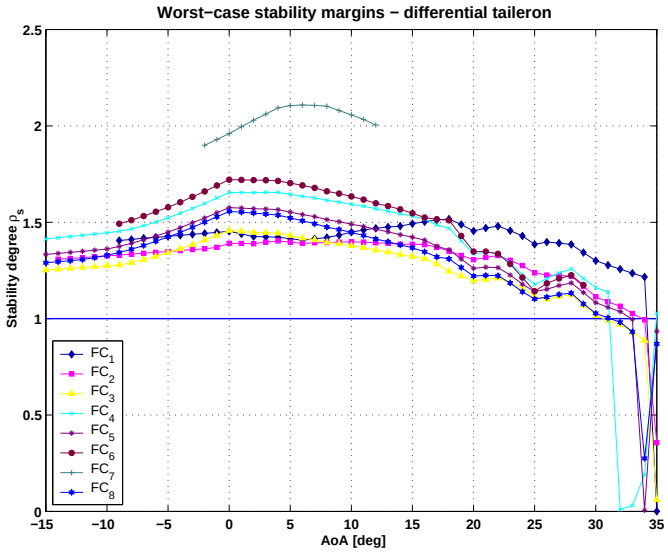

Fig. 21.4. Worst-case stability margins for DTL (small parameter set)

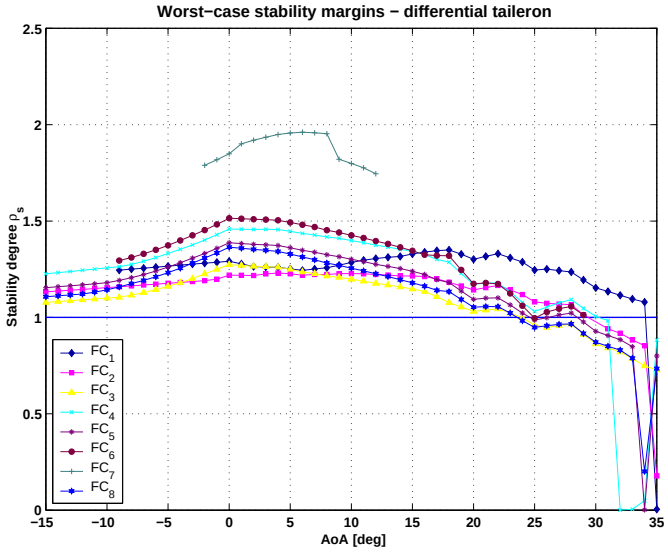

Fig. 21.5. Worst-case stability margins for DTL (full parameter set)

In each performed analysis we determined an "overall" worst-case parameter combination over all flight conditions. In Figs. 21.6 and 21.7, we present the worst-case parameter combination for the differential taileron loop and the corresponding frequency response, respectively. Note that in this case, the optimisation-based search was able to determine a global minimum which led to crossing of the exclusion region through its centre $(0, -180°)$.

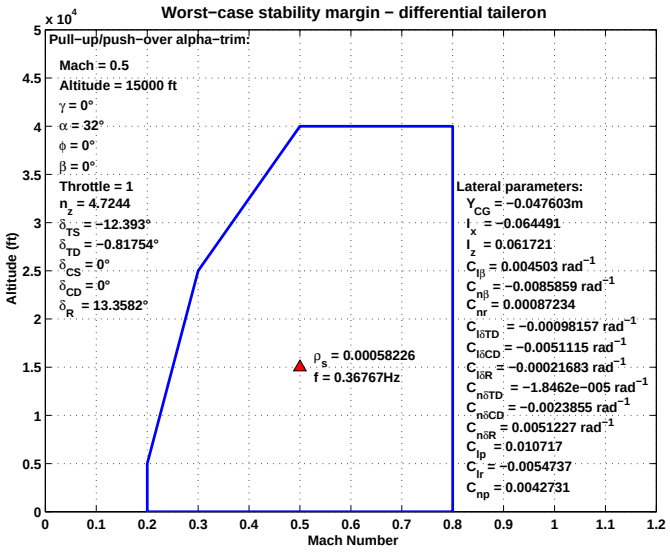

Fig. 21.6. Worst-case parameters for DTL (full parameter set)

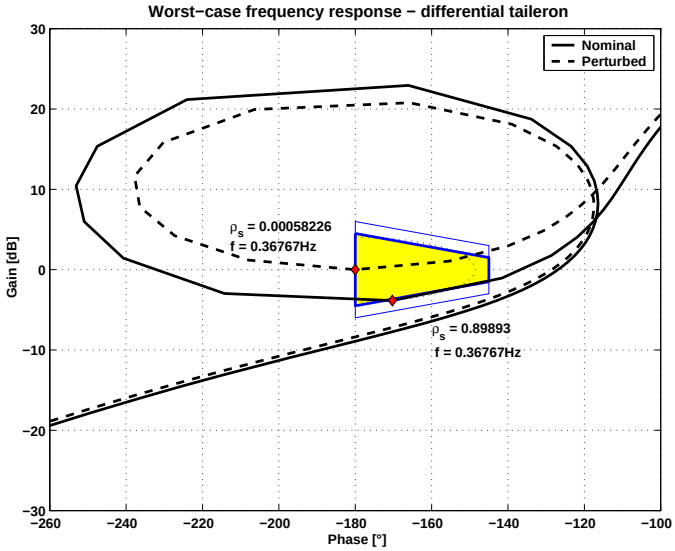

Fig. 21.7. Worst-case frequency response for DTL (full parameter set)

In Figs. 21.8 and 21.9 we present the clearance results for the symmetric taileron and rudder loops.

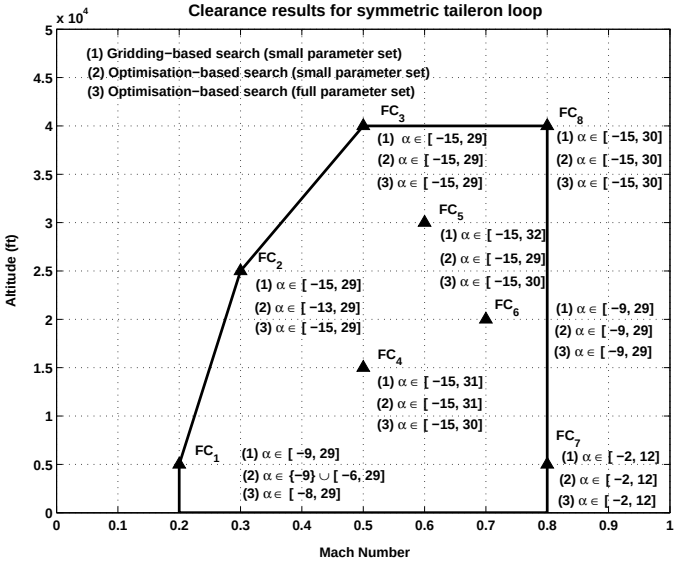

Fig. 21.8. Clearance results for the symmetric taileron loop

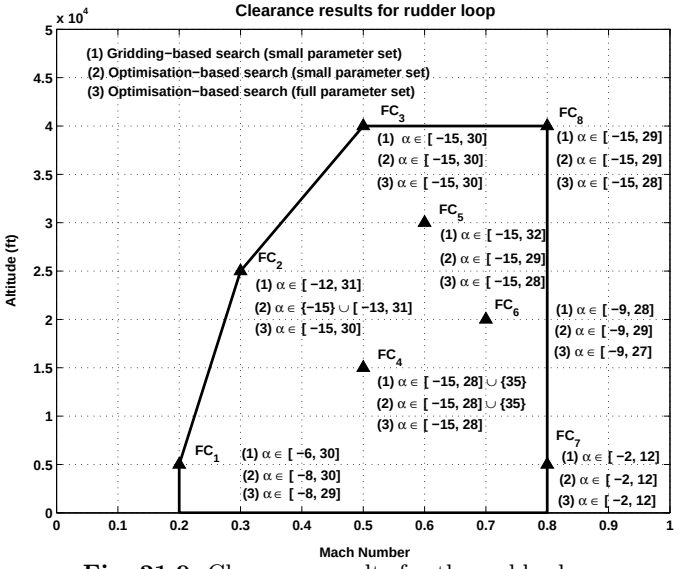

Fig. 21.9. Clearance results for the rudder loop

From the clearance results presented in Figs. 21.3, 21.8 and 21.9 we can see that in some cases, points which are not cleared by the classical

gridding-based search approach, appear to be cleared by the more powerful optimisation-based search approach for the small or even for the full parameter sets. The explanation of this apparent paradox is that the worst-case parameter combinations computed by different methods occasionally lead to violations of the condition $-3 \leq n_z \leq 7$ or of the trimming constraints (21.2). Note that such points, found only incidentally, automatically restrict the effective flight envelope of HIRM+ (see Fig. 21.27) and can be skipped in further analysis.

21.3 Results for the Unstable Eigenvalues Criterion

The goal of the analysis is to identify *all* flight conditions in terms of Mach number M, altitude h, and angle of attack α, and *all* combinations of uncertain parameters where unstable eigenvalues are found. The clearance task formulated in Chapter 10 is to check if the closed-loop eigenvalues lie to the left of the boundary defined in Fig. 21.10.

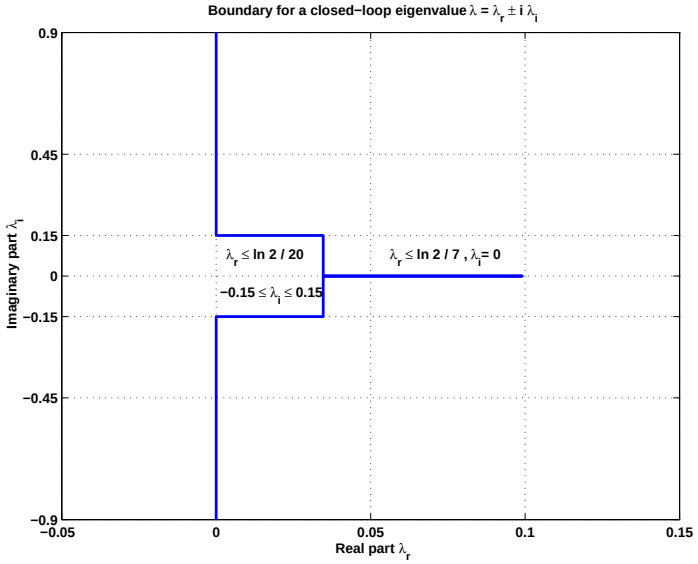

Fig. 21.10. Boundary for eigenvalues.

Mathematically, if $\lambda = \lambda_r + \mathrm{i}\,\lambda_i$ is an eigenvalue of the state matrix A_{cl} of the linearised closed-loop model, then the real part λ_r must satisfy the following conditions

$$\lambda_r \leq \begin{cases} 0, & |\lambda_i| \geq 0.15 \\ \ln 2/20, & 0 < |\lambda_i| < 0.15 \\ \ln 2/7, & \lambda_i = 0 \end{cases}$$

For a given flight condition FC and a parameter vector p, a straightforward "distance" function can be defined as

$$d(p, FC) = -\max\{\mathrm{Re}\,\lambda | \lambda \in \Lambda(A_{cl})\} \tag{21.4}$$

By minimising $d(p, FC)$, we determine the worst-case parameter combination leading to a maximal real part of the closed-loop eigenvalues. This real part corresponds either to a purely real eigenvalue or to a pair of complex conjugate eigenvalues. Each evaluation of $d(p, FC)$ involves the construction of the closed-loop state matrix A_{cl} using the linearised models of the HIRM+ and RIDE controller and the computation of eigenvalues of a matrix of order 61. To evaluate $d(p, FC)$ there is no need to separate the longitudinal and lateral dynamics to perform the analysis for the two categories of uncertain parameters. For both cases, the unique nonlinear parameter uncertain model of the HIRM+RIDE closed-loop system can be used to evaluate the distance function.

In general, the distance function (21.4) does not cover all possible cases of the analysis because, by minimising the maximum real part, the real axis segment in Fig. 21.10 is primarily favoured due to the presence of at least one unstable real pole in the linearised model. To address more specifically the regions delimited in Fig. 21.10, distance functions can be defined for each region according to the values of the imaginary parts. For example, to restrict the analysis strictly to purely real eigenvalues, the "distance" function

$$d_1(p, FC) = \ln 2/7 - \max\{\lambda_r | \lambda_r + \mathrm{i}\,\lambda_i \in \Lambda(A_{cl}), \lambda_i = 0\} \tag{21.5}$$

can be used. To restrict the analysis to the strips defined by imaginary values satisfying $0 < |\omega| \leq 0.15$ the "distance" function

$$d_2(p, FC) = \ln 2/20 - \max\{\lambda_r | \lambda_r + \mathrm{i}\,\lambda_i \in \Lambda(A_{cl}), 0 < \lambda_i \leq 0.15\} \tag{21.6}$$

can be used. Finally, to restrict the analysis to the regions for $0.15 < |\omega|$, the "distance" function

$$d_3(p, FC) = -\max\{\lambda_r | \lambda_r + \mathrm{i}\,\lambda_i \in \Lambda(A_{cl}), 0.15 < \lambda_i\} \tag{21.7}$$

is appropriate. It is possible to perform a simultaneous analysis for all three regions by considering as distance function

$$d(p, FC) = \min\{d_1(p, FC), d_2(p, FC), d_3(p, FC)\} \tag{21.8}$$

This function is not continuous. Due to the shape of the eigenvalue boundary in Fig. 21.10, the migration of eigenvalues from one region to another leads to abrupt changes in the values of $d(p, FC)$. Still, our experiments have shown that for the HIRM+RIDE configuration, the distance function d_2 was practically never active for the given flight conditions, thus the discontinuity in function values manifests practically never. Complete analysis results obtained with both distance functions (21.4) and (21.8) have shown that there

are no qualitative differences for the clearance when using the simpler, but continuous distance function (21.4) instead of the discontinuous one (21.8). This strongly supports our approach in [1] to use (21.4) as a distance function.

In what follows, we present complete clearance results obtained using the SQP approach for the small and full parameter sets for both longitudinal and lateral axis analyses using the distance function (21.8). Table 21.3 shows the number of function evaluations (NFE) and the times for the gridding and the optimisation-based search methods.

Table 21.3. Performance results for maximum real part analysis

	Gridding		SQP (Small sets)		SQP (Full sets)	
	NFE	Time (sec)	NFE	Time (sec)	NFE	Time (sec)
Longitudinal	11382	12053	15060	15975	34496	35421
Lateral	22710	22482	9628	10300	25159	25832

In Fig. 21.11, we present the clearance results for the longitudinal axis analysis. As it can be observed, the analysis for the full parameter set revealed many not cleared points, which were previously cleared on the basis of the analysis performed using the gridding-based approach or optimisation-based search on the small parameter set.

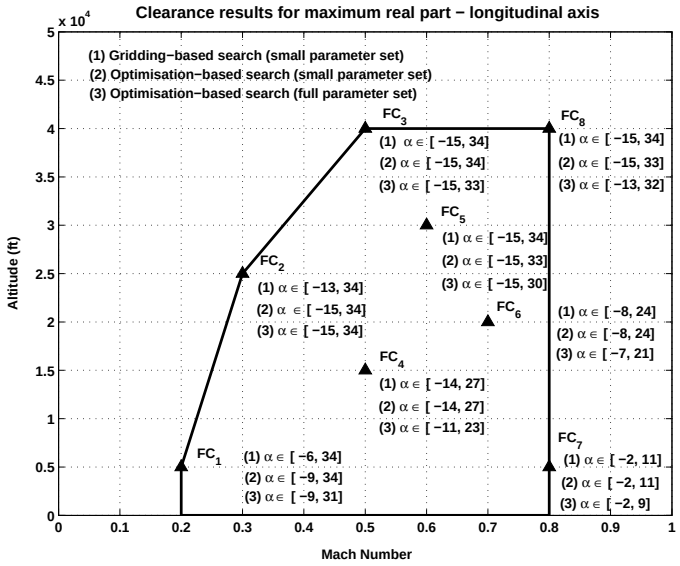

Fig. 21.11. Clearance results for the maximum real part - longitudinal axis

This can also be easily observed by comparing Figs. 21.12 and 21.13, which present the worst-case maximum real parts and corresponding imaginary parts versus AoA for the small parameter set, with the similar Figs. 21.14 and 21.15 for the full parameter sets. These results clearly illustrate the capability of optimisation-based worst-case search to solve clearance problems for many simultaneous parameters.

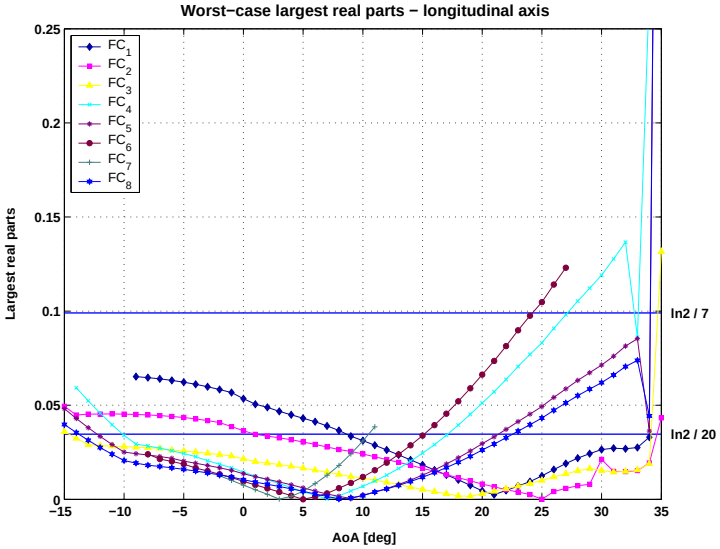

Fig. 21.12. Worst-case maximum real parts (longitudinal, small parameter set)

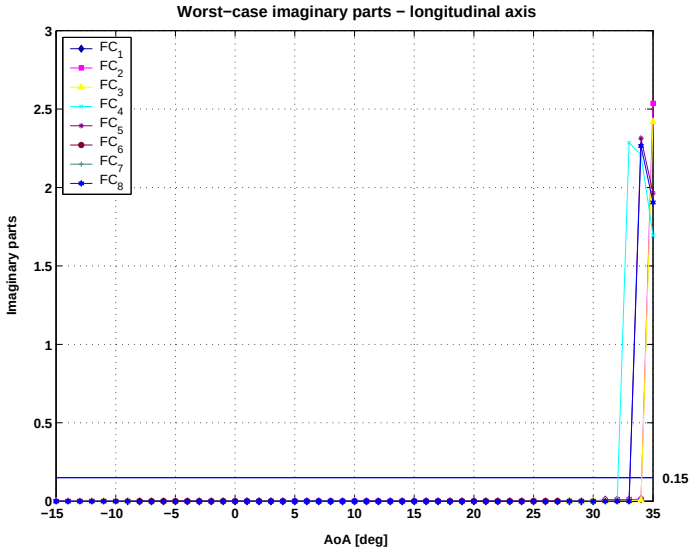

Fig. 21.13. Worst-case imaginary parts (longitudinal, small parameter set)

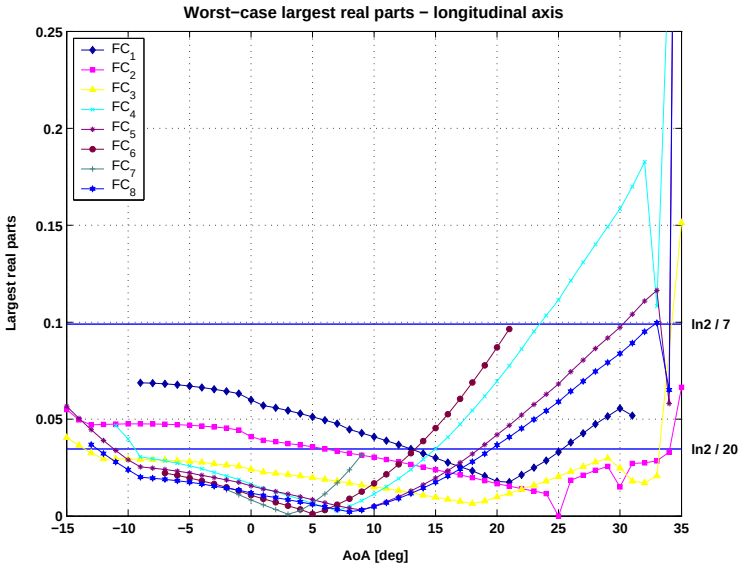

Fig. 21.14. Worst-case maximum real parts (longitudinal, full parameter set)

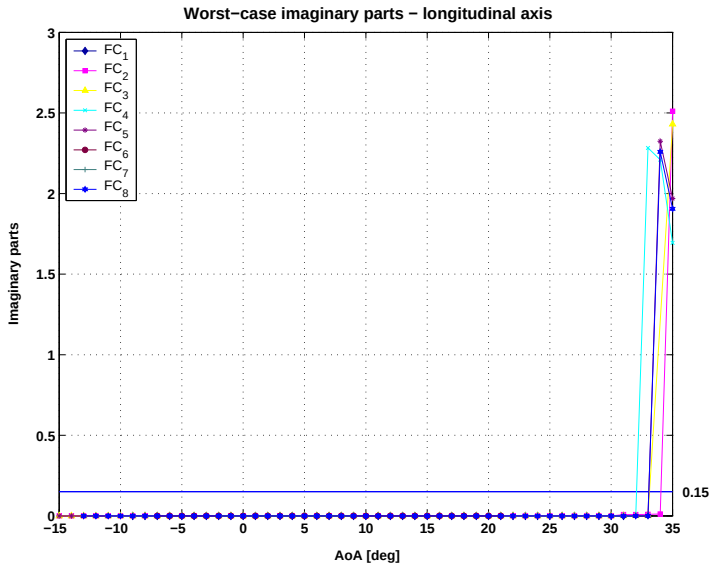

Fig. 21.15. Worst-case imaginary parts (longitudinal, full parameter set)

From Figs. 21.13 and 21.15 it is apparent that in most flight conditions d_1 was active, in very few cases d_3 was active and d_2 was never active. This confirms that the analysis for HIRM+RIDE configuration can be carried out reliably by using (21.4) as an alternative distance function.

In Fig. 21.16 we present the clearance results for the lateral axis analysis. As it can be observed, the results of gridding-based search and of the optimisation-based search are almost identical.

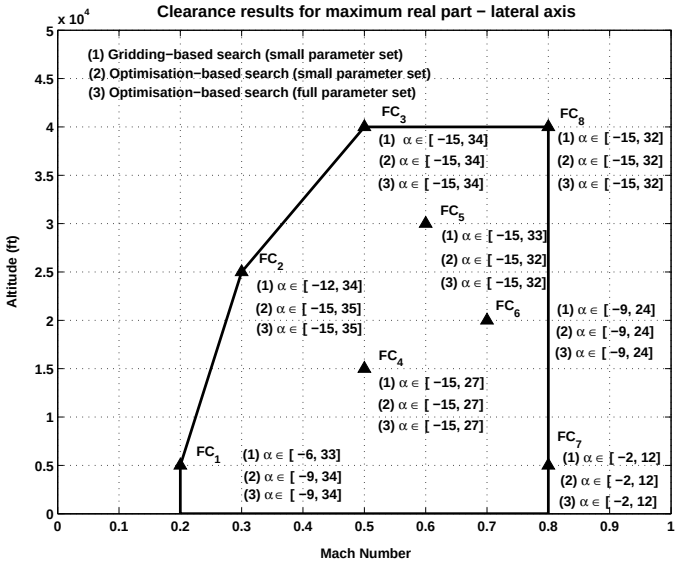

Fig. 21.16. Clearance results for the maximum real part - lateral axis

21.4 Results for the Average Phase Rate and Absolute Amplitude Criteria

The *average phase rate* (APR) and *absolute amplitude* (AA) criteria are intended to identify *pilot induced oscillation* (PIO) tendencies in the pitch and roll axis control loops. The precise mathematical definition of these criteria can be done on the basis of the Nichols plot in Fig. 21.17 for the transfer-function $g(s)$ of the transmission between the longitudinal or lateral stick force and the corresponding pitch or bank angle, respectively. For a given frequency f, let $g(j2\pi f)$ be the corresponding gain-phase representation of the frequency response

$$g(j2\pi f) = A_f e^{j\Phi_f}$$

where A_f and Φ_f are the gain (in dB) and the phase angle (in degrees) at frequency f, respectively. The APR is defined as

$$APR = \frac{\Phi_{f_c} - \Phi_{2f_c}}{f_c} = \frac{-180^\circ - \Phi_{2f_c}}{f_c}$$

where f_c is the (crossover) frequency (in Hertz) where the phase angle $\Phi_{f_c} = -180°$. The absolute amplitude is defined as the gain (in dB) corresponding to f_c

$$AA = A_{f_c}$$

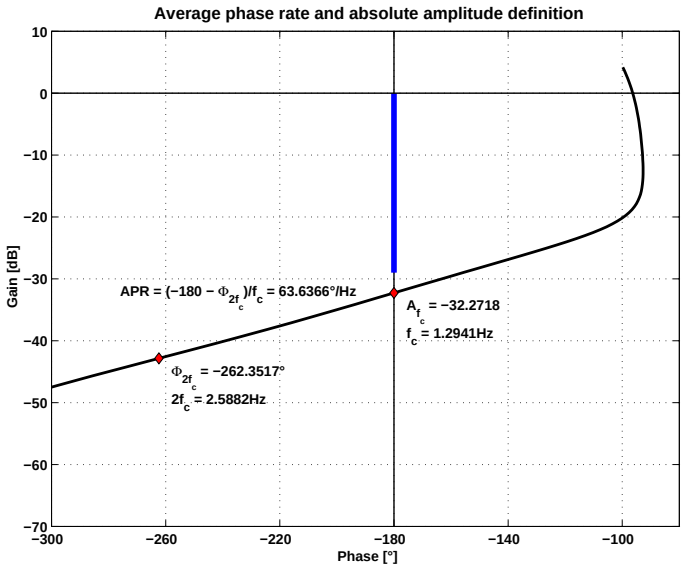

Fig. 21.17. Definition of average phase rate and absolute amplitude criteria

The goal of our analysis is to identify *all* flight conditions in terms of Mach number M, altitude h, and angle of attack α, and *all* combinations of uncertain parameters where the Level 1 boundary defined in Fig. 21.18 is violated or where the absolute amplitude exceeds -29dB (i.e., crosses the bold line in Fig. 21.17). Note that for the HIRM+RIDE control configuration the Level 2 specifications can be easily fulfilled, and therefore our analysis goal is more stringent than that formulated in Chapter 10.

For a given flight condition FC and a parameter vector p, a straightforward "distance" function for the analysis of the APR criterion is

$$d(p, FC) = 1/APR = \frac{f_c}{-180 - \Phi_{2f_c}} \tag{21.9}$$

Minimising this function is equivalent to maximising the APR (by minimising f_c) and thus the worst-case criterion values tend to exceed the boundary of the Level 1 region. For the analysis of the absolute amplitude criterion, the natural candidate for a "distance" function is

$$d(p, FC) = -A_{f_c} \tag{21.10}$$

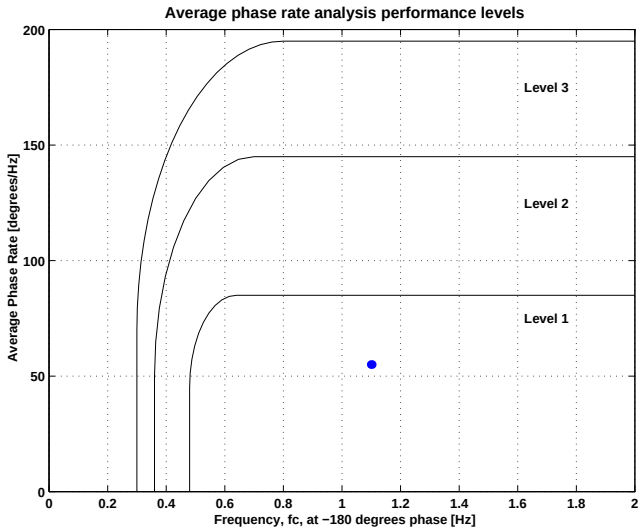

Fig. 21.18. Average phase rate criterion level boundaries

The evaluation of $d(p, FC)$ involves in each case, the construction of the closed-loop linearised model for the transfer function from the longitudinal or lateral stick force to the corresponding pitch or bank angle, respectively, the evaluation of the corresponding frequency response, and the determination of the crossover frequency f_c using, for instance, linear interpolation between two consecutive points to the left and to the right of the $-180°$ axis. To speed up the computation of frequency responses, model reduction techniques can be employed to get minimal order state-space representations for the respective transfer functions.

An important aspect for an optimisation-based search is the reliable numerical evaluation of criteria. For the APR criterion we encountered several difficulties which led to the need to experiment with several gradient-free methods. For example, because of random phase jumps of $\pm 360°$ in the initial phase values, we occasionally obtained completely erroneous values of the estimated crossover frequency f_c. To prevent such jumps, the phase matching approach usually employed when drawing Nichols plots has been extended to include an initial phase of about $-100°$ at the frequency $\omega_{min} = .1$. Another difficulty which we encountered was the occurrence of multiple crossover frequencies. In some cases, for values of $\alpha > 29°$, we encountered points where multiple frequency values satisfy the condition $\Phi_{f_c} = -180°$. To handle such cases, we defined f_c as the largest value of the frequency where a crossing occurs. This allowed us, in most cases, to compute an APR which was within the expected range of values. The cause of this difficulty probably lies in the presence of the α-limiter in the RIDE controller at $\alpha = 29°$. This is why the cleared points for values of $\alpha > 29°$ can only be cautiously accepted. Finally,

discontinuities in the gradient can be expected due to inaccurate localisation of the crossover frequency f_c. This could lead to noisy function evaluations and therefore, to difficulties when using gradient-based methods.

The above aspects can partly explain the somewhat poorer results obtained with gradient-based methods like SQP for the APR criterion, than those resulted from a gridding-based approach. To overcome the difficulties caused by noisy function evaluations, we ran the gradient based search with SQP in conjunction with the GA, where the GA has been used to compute initial points for the gradient-based search. Furthermore, we used the more accurate *central difference* approximation for the gradient instead of *forward difference* approximation, with the immediate consequence of higher computational costs. The best results for the APR criterion for the longitudinal axis parameters have been obtained by using the PS method, but the computational times were about 2.6 times larger than for the SQP method and 3.5 times larger than for COBYLA.

Complete clearance results have been obtained using the SQP approach for the small and full parameter sets for the APR criterion (both longitudinal and lateral axes) and the AA criterion (only longitudinal axis). Table 21.4 shows the *number of function evaluations* (NFE) and the times for the gridding and the optimisation-based search. Note the larger times resulted partly because of more expensive gradient computations.

Table 21.4. Performance results for APR and AA criteria analysis

	Gridding		SQP (Small sets)		SQP (Full sets)	
	NFE	Time (sec)	NFE	Time (sec)	NFE	Time (sec)
APR (long)	11382	13716	24539	32095	48711	61694
APR (lat)	22710	29935	24354	32256	68810	87008
AA (long)	10421	16202	11377	17696	29592	38022

In Fig. 21.19 we present the clearance results for the longitudinal axis analysis of APR criterion. As can be observed, there are practically no differences in the results for the gridding-based and optimisation-based approaches. Practically, all points in the flyable flight envelope exhibit Level 1 handling characteristics, as can be seen from the computed worst-case APR values shown in Fig. 21.20.

In Fig. 21.21 we present the clearance results for the lateral axis analysis of the APR criterion. This time, the analysis for the full parameter set revealed many points which do not fulfill Level 1 performance specifications, but which where considered as satisfying Level 1 handling characteristics on the basis of the analysis performed using a gridding or optimisation-based search on the small parameter set. These findings are also clearly visible from Figs. 21.22

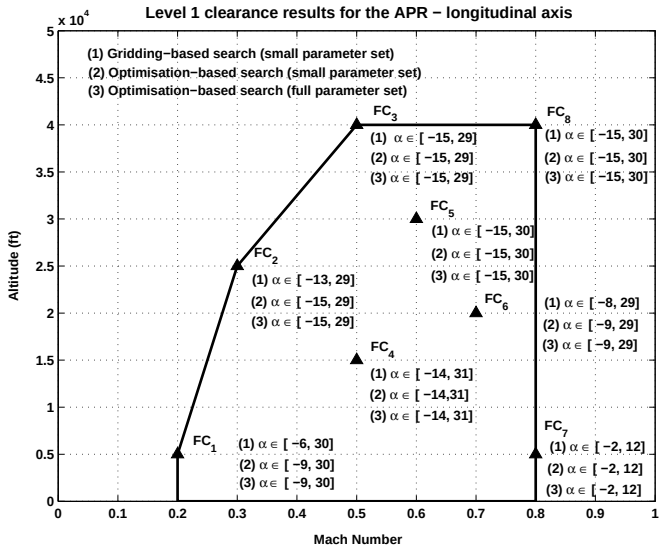

Fig. 21.19. Clearance results for the APR criterion - longitudinal axis

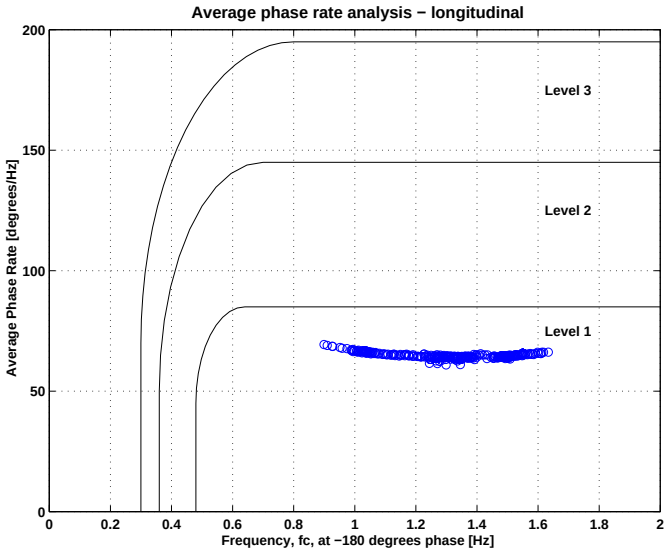

Fig. 21.20. Worst-case APR criterion analysis (longitudinal, full parameter set)

and 21.23, which present the worst-case APRs versus AoA for the small and full parameter sets, respectively. Once again, our results illustrate the power of optimisation-based worst-case search to solve clearance problems with many simultaneous parameters.

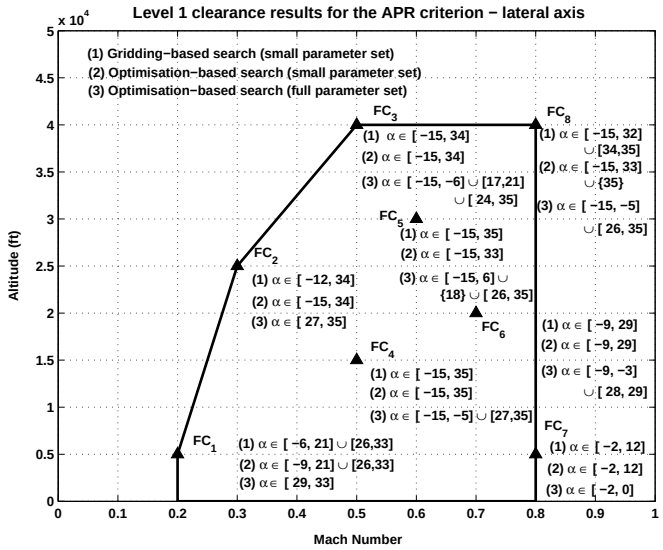

Fig. 21.21. Clearance results for the APR criterion - lateral axis

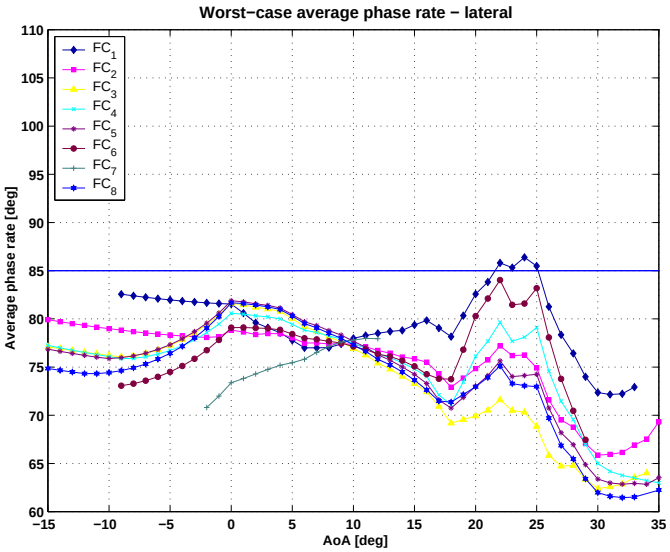

Fig. 21.22. Worst-case APR (lateral, small parameter set)

The analysis for the full parameter set revealed that most of the points in the flyable flight envelope manifest for the worst-case parameter combinations, only Level 2 handling characteristics. This can be clearly seen from the computed worst-case APR values shown in Fig. 21.24.

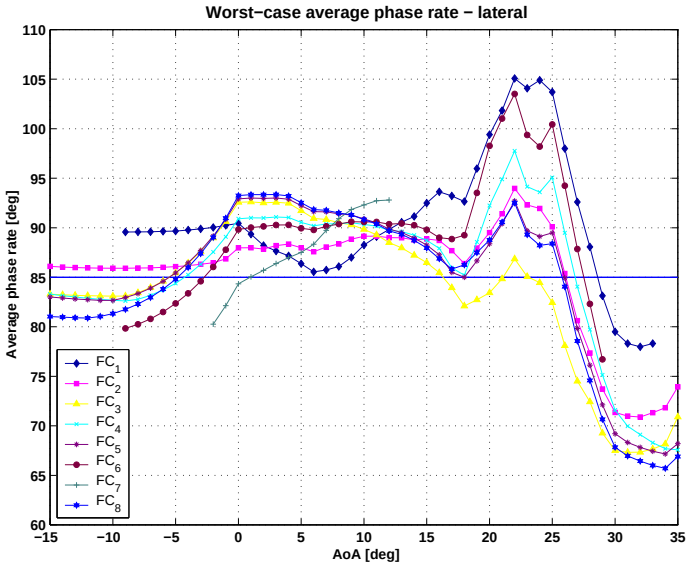

Fig. 21.23. Worst-case APR (lateral, full parameter set)

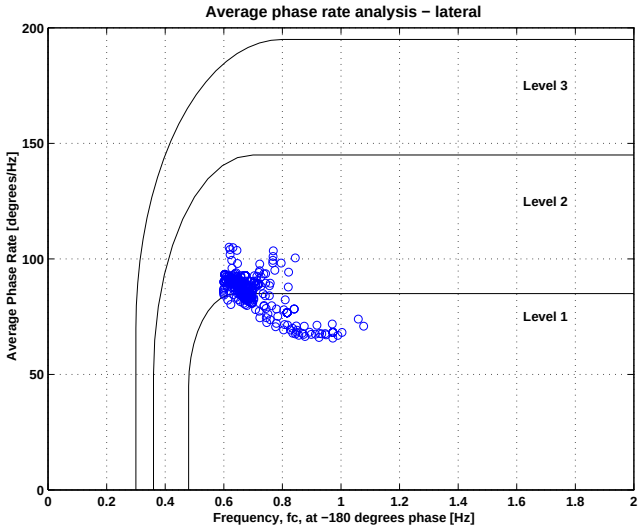

Fig. 21.24. Worst-case APR criterion analysis (lateral, full parameter set)

In Fig. 21.25 we present the clearance results for the longitudinal axis analysis of the AA criterion. Here, there are no notable differences in the results of the gridding and optimisation-based approaches. The only differences arise because of incidental detection of points which do not belong to

the flyable envelope for the computed worst-case parameter combinations. As can be seen in Fig. 21.26, practically, all points in the flyable flight envelope exhibit satisfactory handling characteristics in terms of the computed worst-case AA values.

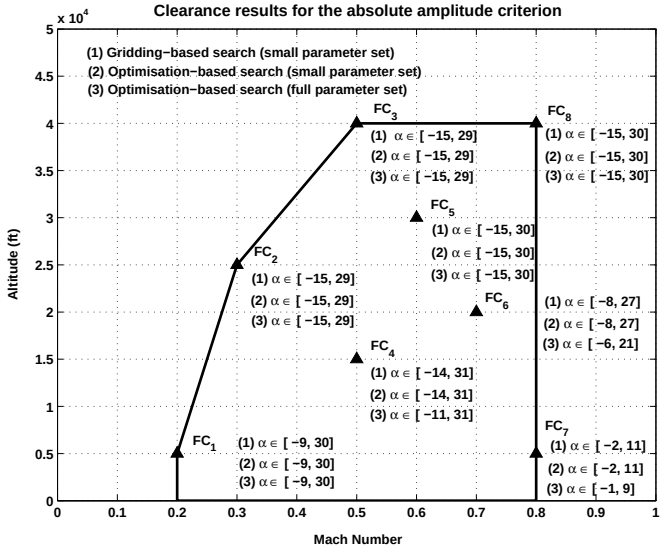

Fig. 21.25. Clearance results for the AA criterion - longitudinal axis

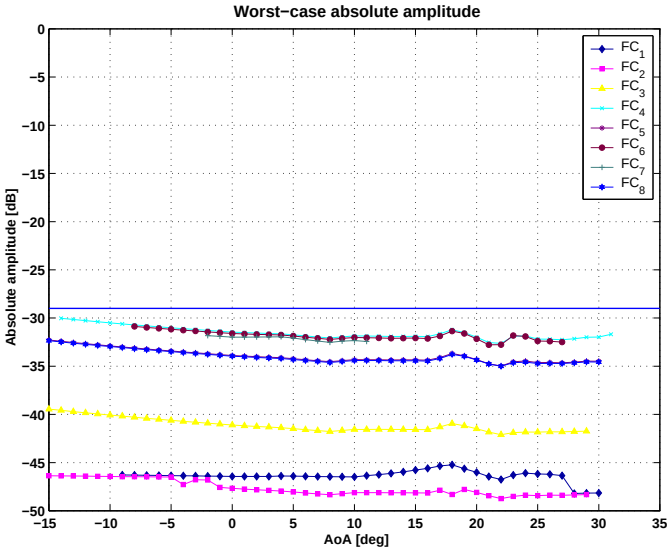

Fig. 21.26. Worst-case AA criterion analysis (full parameter set)

The analysis of the APR and AA criteria revealed several points which are not cleared because the controller gain resulting from the linearisation is zero. This is the case for values $\alpha > 29^\circ$ in FC_2 and FC_3, $\alpha > 30^\circ$ for FC_5 and FC_8, and $\alpha > 31^\circ$ for FC_4. The probable explanation of these cases is the presence of the AoA-limiter at $\alpha = 29^\circ$ in the RIDE controller.

21.5 Evaluation of Results and Conclusions

In this section we summarise the main results achieved by our analysis, discuss the main advantages of the optimisation-based approach to clearance, and some requirements to be fulfilled to use this approach. We also briefly indicate some directions for further investigations.

A first result which emerged from our analysis is an updated flight envelope of HIRM+ which covers the allowed parametric variations of the model. Both the gridding-based search as well as the optimisation-based search incidentally revealed points for which the worst-case parameter combinations led to violation of conditions (21.1) and (21.2). In both cases, the corresponding points are automatically not cleared. Figure 21.27 collects all results reported in [1] and shows comparatively the flight envelopes which can be flown by HIRM+ for nominal and worst-case values of parameters.

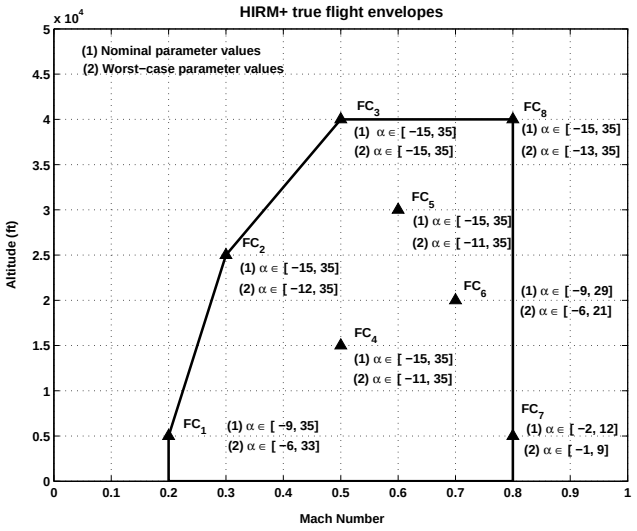

Fig. 21.27. Updated flight envelope for HIRM+.

The stability related clearance results for *full parameter sets* revealed many points where the HIRM+RIDE configuration is not cleared, but which

were cleared by using gridding or optimisation-based search on the small parameter sets. Note that this is a qualitatively new aspect of the optimisation-based clearance when compared with the classical gridding-based approach. Because of its exponential computational complexity, the gridding-based approach is strictly limited to handle only a small number of simultaneous parameters, and is not able to produce results comparable with those obtained by an optimisation-based search, with reasonable costs.

The clearance results for handling criteria revealed that the HIRM+RIDE configuration practically does not fulfill Level 1 performance specifications. Because the handling qualities provides merely indications of possible poor behaviour of the augmented aircraft[1], these criteria are not strict from the point of view of aircraft clearance. Note that for all additional *not cleared* points detected in handling criteria analysis, the achieved performance is still Level 2 (see Fig. 21.24) which can be acceptable in certain conditions. The analysis of the APR and AA criteria revealed several points where the controller gain resulting from the linearisation, for values of angle of attack above 29° is zero. Since such a behaviour can be explained only by the presence of the α-limiter, all results relying on linearisation of the non-linear HIRM+RIDE configuration for values of angle of attack exceeding 29° must be cautiously treated.

The optimisation-based approach to clearance has two **main advantages** over the classical gridding-based approach, both clearly illustrated by our analysis. While the classical approach is limited to analysis with at most $8-9$ simultaneous parameters, the optimisation based approach has no such limitations. Interestingly, in many cases the analysis with the full parameter sets determined parameter combinations where the control configuration is not cleared, although for the small sets, the system is cleared. Note that such cases can not be found by the classical approach. The second important advantage of the optimisation-based clearance is the increased reliability in locating worst-case parameter combinations. While the classical approach evaluates the criteria only in the min/max vertex points of the parameter space, the optimisation-based continuous search found many worst-case parameter combinations lying in the interior of the uncertainty region.

The main technical challenge for the applicability of the optimisation-based approach is an efficient and reliable identification of worst-case parameter perturbation combinations/flight condition for the given clearance criteria. Several requirements must be fulfilled for successful usage of this approach. First, adequate parametric aircraft models must be available which allow fast trimming, reliable linearisation and accurate simulation of the closed-loop system. Note that many of the existing models used for non-linear simulations by the industry, already satisfy these requirements or can be easily adapted for this purpose. Fast trimming and reliable linearisation are of paramount importance for the success of this approach, because these

[1] U. Korte, private communication

computations are necessary at each evaluation of the clearance criteria. A major requirement is the availability of adequate, numerically robust and computationally efficient optimisation software. Since the optimisation problems often involve noisy functions having multiple local minima, alternative methods to the usual gradient search techniques, as for example, derivative-free or global search method must be available.

The optimisation-based approach can be seen as a straightforward enhancement of the current industrial practice by replacing the traditional gridding-based search with an optimisation-driven search to locate worst-case parameter combinations. Since the application of this new clearance approach does not require special additional skills from the users, its acceptance by the industry must be a serious option to be considered to improve the clearance process for the next generation of aircraft.

There are many aspects of the optimisation-based clearance which need further investigations. One such aspect is the use of more powerful optimisation algorithms which are **best suited** to the class of NLPs which typically appear in clearance problems. A promising direction is the use of gradient-free methods with fast convergence rates, able to address the minimisation of noisy and expensive functions. For example, trust-region methods working mostly with surrogate function models to perform optimisation, are very efficient in terms of the required number of function evaluations. Therefore, using the recently developed trust-region methods as underlying optimisation tools can drastically improve the costs and reliability of the optimisation-based clearance.

Another direction is the use of optimisation algorithms for **mixed integer-continuous** problems, where some of the variables have discrete variation and others have continuous variation. By using mixed integer/continuous optimisation, it is possible to combine the discrete grid-based search (e.g., for those parameters with known monotonic effects) with a continuous exploration for the rest of parameters, thus increasing the overall efficiency of the optimisation-based search.

Using **global optimisation** techniques to solve clearance problems is worth investigating in depth. For example, for problems with possible multiple local minima, global search can be used in conjunction with local search algorithms to locate good initialisation points. New and very promising developments occur in parallel methods for global optimisation, where a large number of function evaluations can be done in parallel (e.g., in a genetic algorithm to evaluate a new population). With the advent of cheap parallel architecture machines, the high computational costs associated with global search methods can be significantly reduced, thus making their standard usage as clearance tools affordable.

References

1. A. Varga. Robust stability and performance analysis of flight control laws using optimization-based worst-case search: Linear stability and handling criteria. Technical report, GARTEUR FM(AG11)/TP-119-15, September 2001.
2. A. Varga. Model reduction software in the SLICOT library. In B. N. Datta, editor, *Applied and Computational Control, Signals and Circuits*, volume 629 of *The Kluwer International Series in Engineering and Computer Science*, pages 239–282. Kluwer Academic Publishers, Boston, 2001.
3. G. Grübel and H.-D. Joos. RASP and RSYST - two complementary program libraries for concurrent control engineering. In *Prepr. 5th IFAC/IMACS Symp. CADCS'91, Swansea, UK*, pages 101–106. Pergamon Press, Oxford, 1991.
4. C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal. Algorithm 778. L-BFGS-B: Fortran subroutines for Large-Scale bound constrained optimization. *ACM Transactions on Mathematical Software*, 23:550–560, 1997.
5. M.J.D. Powell. A direct search optimization method that models the objective and constraint functions by linear interpolation. In S. Gomez and J.P. Hennart, editor, *Advances in optimization and numerical analysis*, pages 51–677. Kluwer Academic Publishers, 1994.
6. D. L. Carroll. FORTRAN genetic algorithm (GA) driver. World Wide Web, `http://www.cuaerospace.com/carroll/ga.html`, 1999.