High-Fidelity CFD Maneuver Simulation Using Blade Dynamics, Flight Mechanics and a Pilot Model

Manuel Keßler
Senior scientist
University of Stuttgart, Germany

ABSTRACT

High-fidelity CFD simulations for stationary flight have become state of the art for many years now, including fluid-structure coupling to blade dynamics and flight mechanics to reach a trimmed state. Unsteady maneuver simulations utilizing CFD are much more scarce, mainly due to the huge computational effort required to run for serveral tens or hundreds of revolutions, and thus more often rely on mid-fidelity aerodynamics. Continuous advances in High Performance Computing now allow the application of CFD to relevant macroscopic maneuvers, allowing an investigation of scenarios for example like vortex ring state, which are highly affected by interactional dynamics, and where mid-fidelity tools are less robust and reliable. The paper describes the challenges and possible solutions when progressing a well-established high-fidelity CFD framework from stationary flight to unsteady maneuvers, analyzing the physical causes of problems encountered, and further objectives to be addressed. Examples given for a Robinson R44 helicopter model include a slalom flight, a quickstop maneuver, and vortex ring state, including several evasive methods.

INTRODUCTION

Conventional comprehensive tools like CAMRAD II (Ref. 1), HOST (Ref. 2) or RCAS (Ref. 3) allow analyses of rotorcraft for various flight cases at low computational costs and reasonable accuracy. Actually, if experimental flight data exists, the corresponding models can be tuned to an amazing level of precision over the better part of the entire flight envelope. Taking rotor blade dynamics and deformation into account as well as flight mechanics is crucial to obtain such high-quality data, of course.

However, these tool reach limits when applied to new aircraft, beyond the safe flight envelope, and especially for unconventional configurations with multiple rotors. The numerous interaction phenomena of such rotorcraft are difficult to represent accurately before first flight, without the option to adjust modeling parameters to the specific layout. In contrast, high fidelity CFD simulations allow the evaluation of any configuration from first principles long before actual hardware exists, and with dependable results, if applied correctly. Even if there are still some inevitable modeling uncertainties remaining, for example regarding local phenomena in dynamic stall events, in general the accuracy of results is at least equivalent to a well tuned comprehensive model, without the need for any adjustments to a specific aircraft, even in extraordinary cases.

Of course, computational cost is larger by some orders of magnitude, but the benefit of avoiding an expensive redesign after prototype flight testing will easily outweigh the simulation investments. Certainly, blade dynamics and flight mechanics have to be considered as well in this application. Often, this is accomplished by coupling the CFD solver to a comprehensive tool, replacing low-fidelity aerodynamics by the CFD loads. In an iterative process the control angles, blade deformations and loads are successively exchanged between the two tools and converged to a consistent solution. This has been the state of the art for more than a decade (Refs. 4–6), with mostly quantitative improvements since then on accuracy, robustness and turnaround time for actual technical productivity.

Consequently, for new developments high-fidelity CFD has become an invaluable tool to reduce risk before first flight. However, with very few exceptions, e.g. (Refs. 5, 7–9), only stationary flight states have been considered yet, due to the extremely high computational cost of non-periodic cases. A typical maneuver flight takes several seconds corresponding to about a hundred revolutions or more. This is about an order of magnitude more than for a typical converged trim solution of a steady case. Even very recent investigations (Refs. 10–12) thus compose a maneuver of separate segments of stationary state, without considering the transitional phases in between. Correspondingly, the dynamic loads of the respective motions are not taken into account.

However, with continued growth of computational power, advanced rotorcraft frameworks now have reached the point where maneuver flight simulation at reasonable effort is possible for productive operation. The present work demonstrates the qualitative leap in capability towards such an application.

METHODS

Tool chain

Such an advanced rotorcraft framework and tool chain has been established at the Institute of Aerodynamics and Gas Dynamics (IAG) of the University of Stuttgart over more than two decades now. The fundament consists of the flow solver FLOWer, originally developed by the German Research Center DLR (Ref. 13) and significantly extended by IAG, for example utilising high-order WENO schemes and advanced grid deformation technology (Refs. 14, 15). Blade dynamics are considered using loose fluid-structure coupling to either one of the comprehensive codes CAMRAD II (Ref. 1) or HOST (Ref. 2). These tools also provide the fluid mechanics trim to adjust control angles and attitudes to obtain free flight conditions of vanishing global forces and moments in stationary flight.

Recap on Stationary Trim

For stationary flight cases, the air loads need to balance weight for vanishing residual forces and moments

$$\mathscr{L}(S) + W \overset{!}{=} 0$$

with loads $\mathscr{L}$ generated by an appropriate state S, and the weight W of the aircraft. In general, the state S consists of the four control angles $\underline{c}$, three attitudes $\underline{\theta}$, three velocities $\underline{v}$, and possible further internal variables, as the two flap angles $\underline{f}$, for example, adding up to a total of twelve degrees of freedom in the current analysis. The trim equation consists of the equilibrium of three forces and three moments, and possibly further internal constraints, such as two vanishing lateral hub moments for a teetering rotor, making up for eight equations in this example. Consequently, four degrees of freedom need to be prescribed, typically the three velocities plus yaw angle at zero, and the other ones can be solved for.

As the loads are a nonlinear functional of the state, an approximate state is evaluated, and the residual delta loads drive the iteration to an improved solution. For this, the loads evaluation functional $\mathscr{L}(S)$ is linearized by its Jacobian matrix $L = \frac{\partial \mathscr{L}}{\partial S}$. The iteration algorithm then reads

$$\begin{aligned} R_i &= \mathscr{L}(S_i) + W \\ S_{i+1} &= S_i - L^{-1} \cdot R_i \end{aligned}$$

In a common coupling scenario, the Jacobian L is typically computed using the low-order aerodynamics of the comprehensive code. Although often sufficient, this may be inefficient towards the boundaries of the flight envelope, as the approximation may lack physics in these cases, and sometimes even fail, e.g. at high load factors.

An improvement, albeit an expensive one, can be a CFD based Jacobian computation. Even if one-sided differences are utilized only, this requires $n+1$ simulations for n degrees of freedom in the state S. However, if a converged trimmed state is already available, for small variations of S, as necessary for the computation of L, usually a single further revolution or two is sufficient. Especially for flight states which are hard to converge otherwise this may be a worthwhile procedure.

In either case, such a trim iteration converges usually within a couple of iterations for common flight cases, each comprising about one CFD revolution to reach a reasonably periodic state, so total effort is in the region of approximately ten revolutions for a single flight state. Convergence can be considered sufficient, if residual forces fall below the weight of fuel burnt in one minute, for example, or residual moments below the corresponding change in CG.

Maneuver Handling

In maneuver flight, the state S needs to be enriched at least by the three rates of attitude change $\underline{p} = \dot{\underline{\theta}}$, such that we have three more degrees of freedom plus three additional equations. The static trim equation now extends to the dynamic

$$\mathscr{L}(S) - \dot{\mathscr{I}}(\dot{S}) + W \overset{!}{=} 0$$

with the change of (linear and rotational) momentum $\dot{\mathscr{I}}$ added, which depends on state variable derivatives. Some caution is needed, as $\dot{\mathscr{I}}$ also include gyroscopic contributions due to the rotor inertia. While the tail rotor's contribution is mostly negligible for reasonable maneuvers, the main rotor's is not. These terms considerably couple the dynamic motions of the airframe, complicating the analysis.

Now we again need to "trim" the aircraft. For a static flight case, we typically iterate towards the forces and moments (either measured in wind tunnel or zero for free flight) and adjust control angles and attitudes. In contrast, we could also apply measured controls and attitudes and evaluate residual loads to assess accuracy.

For maneuver flight even more options are possible to compare simulation to flight test, and again four degrees of freedom need to be described and the others solved for. Plausible selections include

A A target trajectory (plus one attitude, e.g. yaw) is given and controls and the other attitudes are driven to match that path.

B Control angles are prescribed and the aircraft moves appropriately.

C Attitudes plus forward velocity are forced and controls as well as trajectory will follow.

Alternatively, other combinations of individual components are possible, as long as the number of open degrees of freedom match the number of constraints. In theory, we could introduce external influences like crosswinds into the dynamic equation and solve for them in order to shift the modeling error to the experimental side, but this is – perhaps besides special cases – exaggerated. However, often only part of the relevant degrees of freedom is measured, which restricts useful choices. And of course, in a perfect simulation the selection would not matter and all variants agree consistently.

Concept B is easiest to implement, since the loads are evaluated for a given state and attitudes and accelerations can be computed accordingly. However, it can be difficult to devise *a priori* an appropriate control schedule to actually perform an intended maneuver. Concept A ensures the scheduled flight path, but requires a pilot model to actually find the correct control angles throughout. Indeed, the system of equations is actually quite stiff, when looking at the generalized eigenvalues, and for excessively aggressive trajectories a precise tracking may be beyond the physical limits, allowing only a more or less accurate approximation.

Depending on the selection of prescribed and open degrees of freedom, the system consists of either differential or algebraic equations or both. For the former, an integration scheme needs to be specified, for example Euler explicit in the simplest case, and the latter are handled by Newton iteration after the functionals $\mathscr{L}$ and $\dot{\mathscr{I}}$ have been linearized, Actually, $\dot{\mathscr{I}}$ already is a linear operator, consisting basically of the (inverse) of mass and rotational inertias, so no further linearization is necessary or possible.

Coupling

For the trim procedure in stationary flight usually weak (or loose) coupling is employed, meaning a loads evaluation over some period, e.g. one revolution, and adjusting controls and attitudes for the average residuals. This also works well with blade dynamics coupling, where the loads and deformations are also exchanged over such a period in terms of Fourier coefficients. For the dynamics, this guarantees periodicity of the solution by construction, which suppresses badly attenuated low-frequency oscillations.

Another option is strong (or tight) coupling, where loads are evaluated in each time step, and adjustments made instantly for the next step. This applies to flight mechanics and blade dynamics as well, and is in general less efficient for stationary flight with a periodic solution, as said low-frequency oscillations need long times to fade away sufficiently. In contrast, for maneuver flight this is clearly the correct approach, since the assumption of a periodic behaviour is inapplicable by definition. Unfortunately, such tight coupling comes at the cost of high-frequency data exchange beween flow solver and flight mechanics or blade dynamics, being detrimental for computational efficiency. The utilization of a low-latency communication library for IAG's simulation framework is described in (Ref. 16), but still in an early development stage, and was not yet ready for production for the current investigation.

In the meantime, although not being exactly periodic, flight mechanic motions are on a much slower time scale than the rotor revolution, blade dynamics, or even vibrations, for example. In this regard, it is plausible to separate attitude changes and linear accelerations and apply only short-term averages of them, similarly to weak coupling for stationary flight. This approach kind of resembles the separation of high-frequency fluid fluctuations into a turbulence model, and low-frequency content into the unsteady evolution of the RANS equations.

Instantaneous attitude rates can easily prescribed in the flow solver, and velocities anyway, so a continuous representation of position and attitude is ensured. For a two bladed rotor, for example, half a revolution is long enough to provide a reasonably smooth average, while reaction delay is sufficiently small for adequate control bandwidth. For the Robinson R44, which is investigated here, half a revolution corresponds to $0.075\,\mathrm{s}$, which is in the region of human response times, and thus kind of representative for the real world.

In the strong coupling, rotor loads are clearly different for the rotor blades at various azimuthal positions, e.g. in lateral or longitudinal orientation, so there are high load fluctuations over one period. In reality, this is perceived as vibrations in the cabin, although airframe elasticity plays a major role in this regard as well (Refs. 17, 18). However, these high-frequency linear and angular accelerations do not contribute to the macroscopic flight mechanical state, but only destabilize the numerical time integration process. Consequently, if vibration is not the goal of the investigation, it is much better to utilize the described weak coupling and the corresponding loads averaging over full periods, which suppresses such fluctuations by design.

The procedure is thus as follows:

1. CFD simulation is run for half a revolution and loads averaged.
2. In case the loads simulated within the last period and the targeted ones are too different, adjust state for this period and go back to 1 to recompute, until sufficient convergence is achieved. Actually, this never happened for the maneuvers investigated here, but may occur for more aggressive control inputs and maneuvers.
3. Loads are transferred to flight mechanics and the state for the next period updated accordingly.

Although not being exact in the mathematical sense and even some minor systematic phase shift may occur due to the state update being conducted slightly after the loads average being taken, this error is orders of magnitude smaller than the accuracy of CFD itself and thus easily negligible. For example, slight variations of similarly reasonable solver settings, e.g. one subiteration more or less, or tiny time step variations, can lead to global load discrepancies up to one percent of weight, depending on flight state, without any indication for the more accurate result. However, this is still better than in flight experiments, where the variance between individual data points even under nominally identical conditions is often easily double that figure, and worse for different aircraft of the same type.

Blade Dynamics

As is well known, blade dynamics play a major role in helicopter simulations, and thus definitely need to be accounted for in any reasonable approach. Unfortunately, CAMRAD II is not able to couple in maneuvering flight, so our well tuned CAMRAD II model for the Robinson R44 is of no use in this regard. CAMRAD II can either do weak coupling for stationary flight, or stand-alone transient computations, but even a steady rotation around pitch or roll axis is not possible in coupling mode. For HOST, the other comprehensive tool available at IAG, no R44 model exists, and our coupling process is still much less robust and validated with that software.

Consequently, at least blade flapping is taken into account explicitly, evaluating the in-plane hub moments over each period and adjusting the tip path plane accordingly. For the R44 with its teetering rotor, it needs to be assured that no hub moments are transferred to the airframe. In the stationary case, blade flapping follows the cyclic input nearly 1:1 in magnitude to fulfill that condition. In contrast, in maneuvering flight flapping lags behind somewhat, resulting in residual aerodynamic moments, which are then used up as dynamic loads for tilting the rotor plane.

However, this change of rotor orientation at the same time considerably changes the aerodynamic loads, and on a much shorter time scale than the evaluation window. According to Padfield (Ref. 19), the characteristic time (68% adaptation) for the flapping motion is about 120° azimuthal rotation, so the flapping adaptation is mostly accomplished during one period. In consequence, if not accounting for this flapping change over the period, hub loads are grossly off.

One remedy would be to adjust flap rates accordingly and do a second (and possibly third, and fourth, ...) CFD run, until convergence. However, this is obviously computationally very expensive. But as this accomodation is basically an exponential decay of the offset to the final state, it can be integrated analytically over the given period and – thanks to linearization of the loads operator $\mathscr{L}$ via its Jacobian L – the corresponding correction transferred to all the other loads, like lateral airframe forces.

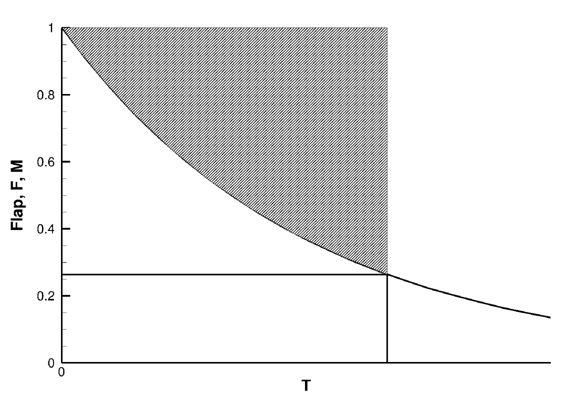

Figure 1. Flap alignment and loads correction

Fig. 1 visualizes this process. Blade flapping starts at some offset (normalized to 1) and the resulting aerodynamic hub moment drives it towards its final steady value. With decreasing discrepancy, the driving moment declines as well, resulting in an exponential decay for the blade flap offset. At the end of the time step, some value remains, which carries over to the next time step. The resulting hub angular momentum has thus been reduced by the hatched area, and the remainder corresponds to the change of rotor angular momentum by tilting the tip path plane. All other aerodynamic rotor loads (thrust, in-plane forces, drive torque) transferred to the airframe then need to be adjusted similarly by reducing the original value by the hatched area ratio and the respective Jacobian L contribution. As a consistency check, the equivalently adjusted lateral hub moments need to cancel out with the gyroscopic rotor tilt moments when handled this way.

By this method an expensive iteration of CFD recomputations with adapted flap schedules can be avoided at mostly equivalent accuracy, as the necessary flap adjustment is usually very small anyway between successive periods. Nevertheless, for fundamental stability reasons due to the stiff algebraic equation system, the effect definitely needs to be considered. Of course, this flap handling by tip path plane tilting is not an exact representation of the individual motion of a finite number of blades, but an approximation of a continous disc containing some angular momentum and corresponding inertia, including the gyroscopic effects. Nevertheless, this asymptotic perspective gives a sufficient estimate for the dynamic effects within the modeling accuracy obtained elsewhere.

Even then, blade deformation is not yet included, which is also a major contributor. However, we do not have any flight test data yet, to compare for example control schedules of our simulations to. An elastic blade twist, for example, would result in different cyclic and collective angles, the delta more or less equivalent to an average of the twist over the blade radius and its azimuthal distribution. However, this should mostly resemble an offset in the control schedule, the general trends over a full maneuver are expected to be quite similar. As soon as we will have flight test data available, the modeling will be upgraded accordingly to take blade elasticity into account. Of course, this will also include an individual blade flapping treatment necessarily.

Pilot Model

With the exception of concept B the control angles are to be computed by some means of a pilot model, either to follow a given trajectory (concept A), or attitude schedule (C). In each case, either the linear accelerations, and consequently forces, or rotational accelerations, accordingly moments, are prescribed and need to be matched by proper control adjustments. Although theoretically possible by inverting the proper part of the loads Jacobian L, this proved highly unstable, probably due to some inaccuracy within the CFD computed Jacobian or increasing nonlinearities in the full load operator $\mathscr{L}$, which gets evaluated by CFD. Even with considerable underrelaxation the system developed immediately high oscillations and diverged shortly.

Instead, some primitve means of "autopilot" was invented, somewhat related to the reaction of a human pilot. Each control angle was treated individually, meaning pitch, roll and yaw axis for cyclics and tail, and vertical axis for collective. For the latter, current thrust and vertical speed was included, respectively their deviation from the target values, basically the P and I components of a classical controller. Lateral cyclic included obviously the rolling moment and roll rate, horizontal force, and lateral speed. This corresponds to P and I components again, and two (I^2) and three (I^3) integrations, respectively, as the horizontal force corresponds to the roll angle (integrated roll rate) and horizontal speed to integrated force, at least as long as the roll angle is not too large. Similarly for longitudinal: pitching moment and rate, forward force and forward speed. Tail finally was assigned to a combination of yawing moment, yaw rate and yaw angle, relating to P, I and I^2 terms.

Identifying reasonable control parameters proved somewhat challenging, as a diverging CFD simulation wastes considerable computational resources, and all parameters depend more or less strongly on each other. So the Jacobian L was utilized as a surrogate model for adjusting parameters for various intended flight maneuvers, and the obtained values proved sufficiently robust to effectively handle the simulated maneuvers.

As it turned out, the actual value depends on the length of the period chosen for the weak flight dynamic coupling. In general, for short time windows, the stability boundaries are less stringent to keep the controller stable, while for larger periods, it may be impossible to find stable parameter combinations for the selected control architecture. Originally, a two revolution period was aspired, resembling about 0.3 s of time, which gives some averaging length for the loads and seemed plausible as a generic pilot reaction latency. However, no stable controller parameters could be found due to insufficient controller bandwidth in relation to the reaction time of the system. In part this explains why helicopters are notoriously difficult to fly – you need to react *very* quickly to keep the aircraft stable. Going down to half a revolution period considerably increased the controller parameter margin, to the point of being tolerant even to large load perturbations due to inadvertent input mistakes for intermediate periods, for example a full degree of wrong cyclic. Nevertheless, such disturbances can still cause considerable oscillations, so there is certainly room for improvement, especially on damping.

Of course, this hands-on approach is certainly *not* a replacement for a fully operational autopilot with various inner and outer control loops for stabilizing each individual axis and further ones to follow some specific mission profile, capable of handling highly varying flight conditions along. However, it proved useful enough to keep the helicopter in the intended operational state, even under the stochastic variation of CFD loads in each period, and intended changes of helicopter attitude and velocity.

Maneuver coupling

The flow solver FLOWer now runs for just over half a revolution (192 time steps at 1° azimuthal resolution) at one time, in order to ignore eventual initial transients when averaging the loads over a physical period of half revolution. These transients may arise from abruptly changing the motion or control angles between consecutive periods, but actually the changes are sufficiently small for all flown maneuvers that they are barely visible in the data. Nevertheless, the job length was kept as it matches favourably with an output every 48th time step, achieving about 50 Hz framerate for videos created from that.

The actual control of the individual runs is taken by a Python script, which reads the maneuver and control laws from its input, analyzes the loads and time from the last period, adjusts all degrees of freedom as position, velocity, controls, attitudes in the input file according to the pilot model and maneuver concept, and starts up the next run.

RESULTS

CFD Setup

The helicopter type used is a Robinson R44 Raven II, a lightweight helicopter with relatively low disc loading. The CFD setup used is somewhat simplified in geometric details – no rotor hub, for example – in order to keep cell count and thus computational effort at reasonable cost. The block structured grid system consists of body meshes in 9 structures, ensuring $y^+ \approx 1$ at the wall and a growth factor of 1.3 in normal direction. For Slalom and Quickstop, a local background mesh moving and rotating with the helicopter and resolving 2.5 cm near the body and coarsened up to 0.4 m at about one diameter away is used, inserted into a large background mesh up to 20 diameters in inertial system. The Vortes Ring State analyses utilize only the local background mesh, but extended to two diameters away. The different meshes are connected via Chimera grid embedding, with about 10 million cells in the body grids, 11 million in the local grid (41 million for VRS) and 33 million in the background grid, and totalling about 54 million (51 million for VRS) cells in just over 4000 blocks. The setup runs fairly well on 1024 cores of 16 nodes of the HLRS Hawk cluster (utilizing only half the available cores per node due to memory bandwidth), using about 17 s wall clock time for 40 subiterations at a time step of 1° main rotor azimuth. This makes for about two seconds of real flight time computable per day, demonstrating the considerable effort required for a full maneuver of, say, 20 seconds real time.

Coordinate systems follow DIN/LN9300 or ISO 1151 conventions, namely x forward, y to the right (starboard) and z down, fixed to the airframe, and corresponding roll φ right, pitch θ nose up and yaw ψ nose right angles. Controls are given as main (up), cos (lateral left), sin (longitudinal rear), tail (nose left), and blade flapping as cos (tilt forward) and sin (tilt left). For the controls, main rotor collective is positive

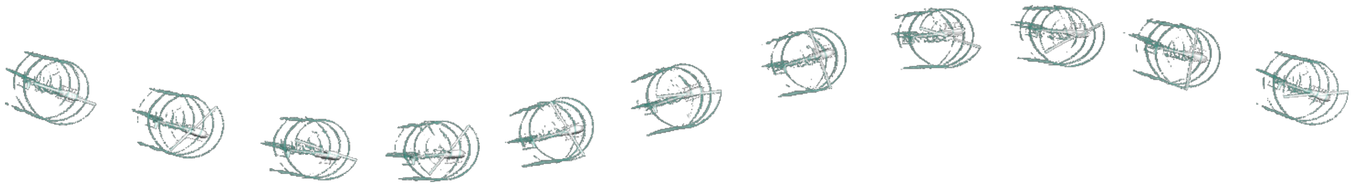

Figure 2. Slalom trajectory and attitude

up, as conventional, but the others cos/lateral, sin/longitudinal and tail/pedal are reversed to common pilot conventions. To be unambigous, the terms used are cos/sin/tail, with their exact mathematical definition.

Initial Trim

Before the intended maneuvers can start, a steady state flight condition is required. Slalom and Quickstop start at 20 m/s (approximately 40 kts), while Vortex Ring State is started from Hover. So two different stationary trims are accomplished using the method described in section Recap on Stationary Trim, getting a converged solution after 7 and 12 iterations, respectively.

After having reached steady flight, for both trim states the Jacobian matrix L is evaluated by disturbing the state S in all degrees of freedom by a small amount (1 m/s in velocities, 1° in all angles, 5°/s in all rates) for two more revolutions and taking finite differences.

Although the Vortex Ring State is admittedly quite different from the Hover situation, where the Jacobian has been evaluated, it is hardly possible to get robust derivatives in this flight state, as the revolution-to-revolution variance in this highly unstable state is already larger than the loads differences created by the small disturbances. Consequently, the "nearby" Hover situation is an approximation according to best effort, and actually works reasonably well.

Slalom

The first maneuver is a sine shaped Slalom at a forward velocity of 20 m/s with a lateral amplitude of ±10 m over a period of 10 s. Trajectory and attitude are computed analytically for this shape, with yaw angle matching the flight direction, roll angle to counter centrifugal forces, and keeping pitch angle constant. Control angles are adjusted to generate the necessary moments and thrust, according to concept A. Actually, it is not an exact representation of A, but somewhat simplified, as all attitudes are prescribed according to the analytical schedule. As a consequence, only moments are matched, and lateral as well as longitudinal forces may deviate slightly.

Fig. 3 shows the controls necessary for one full period of the continuous slalom trajectory, starting at the center position and moving to the right, while rolling left. As can be seen easily, both lateral and longitudinal inputs are required to follow the proper path. The tail rotor has to work harder for the counterclockwise main rotor over the first half period to yaw the nose left again, and then lets go after crossing the centerline again.

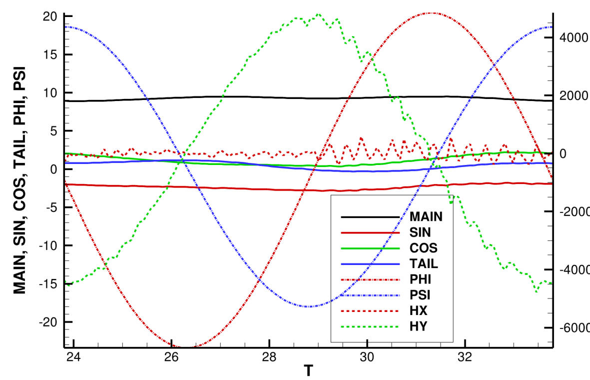

Figure 3. Control input (MAIN, COS, SIN, TAIL) for prescribed attitudes (roll/PHI, yaw/PSI) of slalom and aerodynamic hub moments (HX, HY)

What may be intriguing at first is the occurence of the large aerodynamic hub moment in y direction, pitching down at the beginning of the cycle, and up halfway through. This seems to be counterintuitive, as the helicopter does not pitch at all, and there should be also zero in-plane hub moments for an ideal teetering rotor. However, this is only the *aerodynamic* moment, and it is counteracted by the *dynamic* moment resulting from rolling to the left at the beginning, which means that the rotor with its angular momentum (pointing up for the counterclockwise rotor) is tilted to the left. This tilt requires a torque in negative y direction, which is provided by the aerodynamic loads. As already mentioned in section Blade Dynamics, this aerodynamic torque is "eaten up" by the rotor disc tilt (or rolling rate), and thus net zero in-plane mast moments are transferred to the airframe, as is consistent with the teetering hub.

Actually, there *is* a residual total moment on the airframe, as seen in Fig. 5, stemming from some blade flapping and thus thrust vector tilt with respect to the airframe (not through the center of gravity), which is indeed in the x direction. It provides the angular acceleration necessary to control the rolling motion. At the beginning of the cycle the helicopter rolls already to the left, which is then stopped a quarter period later at the maximum y elongation, requiring positive x torque on the airframe. Then the aircraft rolls back, getting level halfway through when crossing the centerline of the flight trajectory,

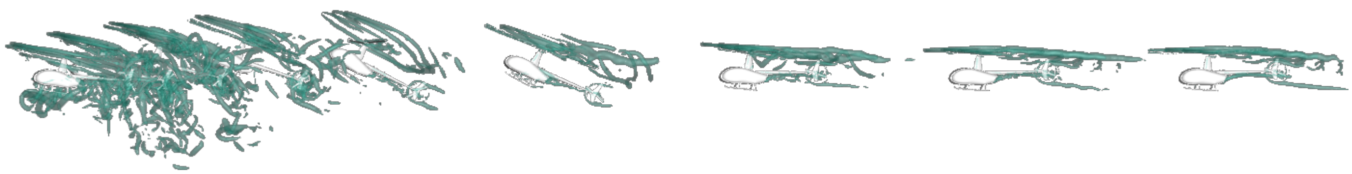

Figure 4. Quickstop trajectory and attitude

and rolling most to the right after three quarters, where the most negative *x* torque acts on the airframe.

One can also see some small wiggles in the control angles at this scale, where variances in the CFD solution – probably due to some stochastic vortex interactions or local separation phenomena, for example – produce small changes in airloads, which require adequate control reactions to generate the loads essential for the given trajectory. In a different interpretation, such control compensations could be related to atmospheric turbulence or gusts, if they were part of the simulation.

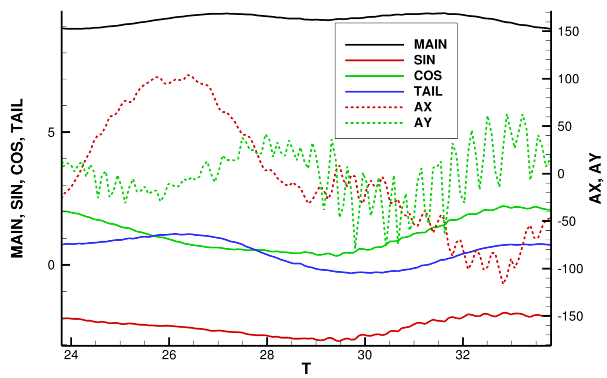

Figure 5. Control input for slalom and resulting airframe moments (AX, AY)

Quickstop

The Quickstop actually has been run in two variants, both starting out again from trimmed flight at $20\,\mathrm{m/s}$, and corresponding to concept B (control inputs prescribed, freely moving aircraft). In the first scenario, the pilot model of section Pilot Model was commanded a target velocity of $0\,\mathrm{m/s}$, and thus an initial speed deviation of $-20\,\mathrm{m/s}$. The model reacted quite reasonably by "pulling back", reached a pitch rate of $11\,{}^\circ/\mathrm{s}$ and a maximum pitch angle of $19\,^\circ$, slowing down to $10\,\mathrm{m/s}$ in about 5 seconds. Unfortunately, the pilot model then started to diverge, especially at the tail rotor, such that another approach was needed.

In order to overcome the stability issues with the pilot model and at the same time accelerate this maneuver and shorten the stopping distance, a different schedule was devised. Longitudinal cyclic was prescribed instead of computed after having reached stationary forward flight during the trim phase. At the starting point, longitudinal was increased from the trimmed

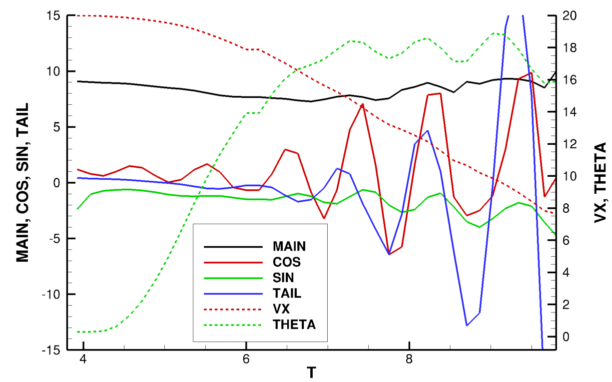

Figure 6. Control input for quickstop by the pilot model and resulting pitch attitude and forward deceleration

value by $+5\,^\circ$, causing a rapid pitch-up motion, which very quickly (within two revolutions) reached its limit value of approximately $30\,{}^\circ/\mathrm{s}$. After just over a second the longitudinal was decreased againg by $-7\,^\circ$, just below the trimmed value. The pitch-up stopped within a second, reaching a maximum pitch angle of about $31\,^\circ$, after which two seconds of the original trimmed longitudinal was applied. During that phase, pitch attitude decreased slowly, and then longitudinal was handed over to the pilot model.

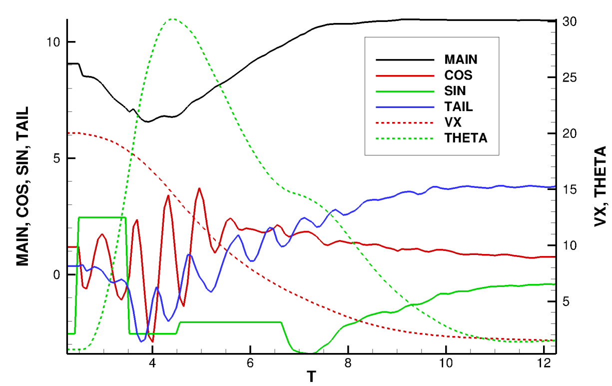

Figure 7. Prescribed longitudinal (SIN) input for quickstop and other controls by the pilot model, and resulting pitch attitude and forward deceleration

During this quite aggressive cyclic input, again significant oscillations emerged on the other controls, which where still handled by the pilot model, but fortunately quickly faded out during the pitch down phase.

The described control schedule can be certainly improved for a more exact stopping point instead of the creeping deceleration later on, but the general behavior is plausible. Looking more closely at the longitudinal flapping motion, the corresponding control input does not immediately result in an appropriate reaction, as it would do in a static simulation, but first remains in its offset state, creating the necessary hub moment to initiate the rotor pitch up moment in x direction. The lagging of the tip path plane behind its static orientation, suitable for the control input given, provides the aerodynamic moment to further tilt the plane toward its target. This gyroscopic effect increases, correspondingly to the pitch rate, and decreases when slowing down the pitch up motion. The pitch down rotation occurs at a lower rate, and the aerodynamic x hub moment is reversed during that time. While the y moment oscillates again due to insufficient damping on the pilot mode for lateral cyclic, its mostly symmetric around zero, not introducing any macroscopic roll rates or attitudes.

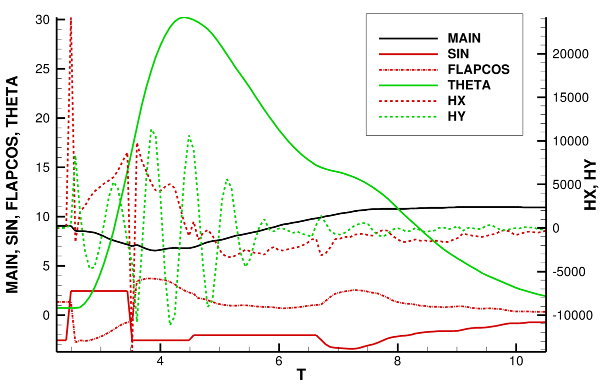

Figure 8. Primary control inputs (MAIN, SIN) for quickstop, and resulting pitch attitude, longitudinal flapping and hub moments (HX, HY)

Summing up, the Quickstop maneuver could be simulated qualitatively correct, with the precise control schedule leaving room for further optimisation. Furthermore, the primitive pilot model of section Pilot Model is far from being ideal and suitable for actual flight control, and only marginally stable. Nevertheless, it allows some reasonable control input to actually run specific maneuvers.

Vortex Ring State

As the final goal, the challenging task of Vortex Ring State recovery was attempted. At first, starting from a trimmed hover state, vertical rate was enforced to $-10\,\mathrm{m/s}$ and collective input reduced by $-1.5\,°$, and the helicopter quickly built up further vertical descent rate. According to the linearized Jacobian L a downward velocity of about $-8\,\mathrm{m/s}$ or $-1500\,\mathrm{ft/min}$ is to be expected for that collective setting, but it is the specific character of Vortex Ring State to be highly nonlinear. Actually, descent rate stabilized within about 7 s at pretty much the double value, namely $-15\,\mathrm{m/s}$. In order to verify actual VRS conditions, the streamlines in the airframe relative coordinate system show in Fig. 9 the expected pattern of heavy recirculation.

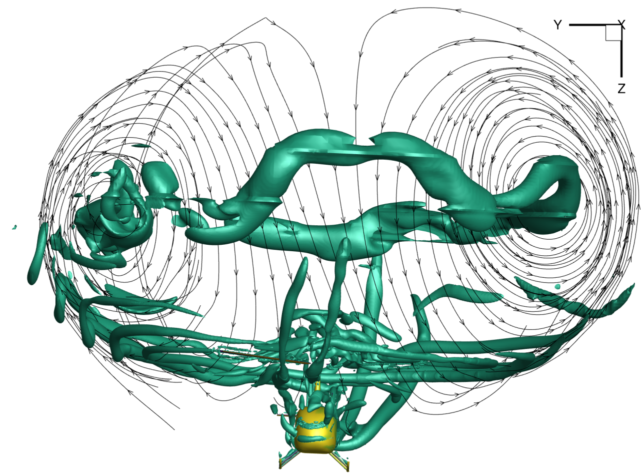

Figure 9. Streamlines and λ_2 isosurface in fully developed VRS

After having reached that state, in simulating a "panic" recovery collective was increased by $3.5\,°$, that is, $2.0\,°$ above the previously trimmed hover collective. As one can see, this gave a short push upwards (more thrust than weight) by up to 8 kN, which quickly fell back to zero after less than a second. Afterwards, rate of decent stabilized again at the slightly lower $-12\,\mathrm{m/s}$. Adding two more degrees of collective finally seemed to recover from VRS, stopping the simulation at $-8\,\mathrm{m/s}$ and still $4\,\mathrm{m/s^2}$ deceleration and probably two more seconds before reaching hover again. Estimated height loss would be about 90 m after the "panic pull" in this situation.

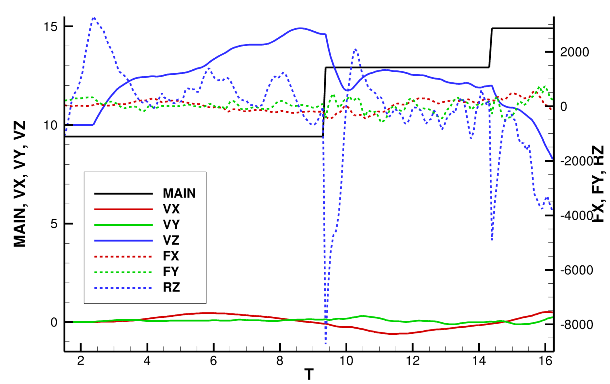

Figure 10. Collective pitch input, and resulting residual forces (FX, FY, RZ) and airframe velocities (VX, VY, VZ)

As can be seen from this graph, horizontal residual forces undergo heavy fluctuations as well. Correspondingly, the pilot model strives hard to keep the helicopter level, with considerable control input at the cyclics and pedal. Nevertheless, attitude fluctuations of several degrees occur, particularly in roll, where the moment of inertia is smallest. This happens especially in the phase with collective pulled up again, where ride quality severly suffers without much help to escape VRS.

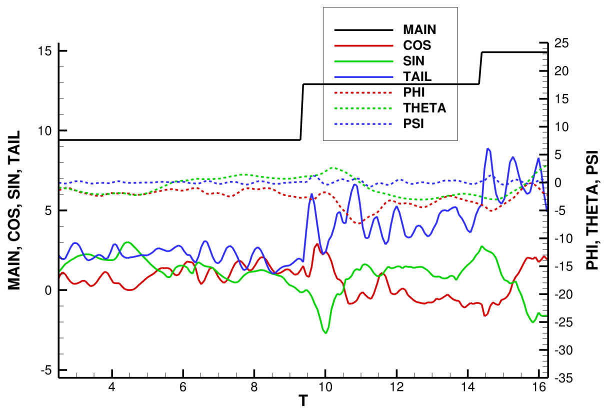

Figure 11. Control inputs by the pilot model, and resulting airframe attitude

Of course, helicopter pilots are trained extensively to recover from VRS not by pulling collective, but by leaving the recirculation zone. Indeed there are three common evasive maneuvers established within the community:

1. Classic: Reduce collective further and put forward cyclic, until sufficient forward velocity develops, then pull collective when reaching the edge of the recirculation zone
2. Airbus: Similar to Classic, but pull collective immediately together with forward cyclic
3. Vuichard: Pull collective immediately like Airbus, but execute lateral cyclic to the advancing blade side (right for a counterclockwise rotor) together with power pedal, effectively crossing lateral and pedal input

The idea behind 1. is to leave the recirculation zone downward while recovering control authority, and to the front, with the danger of loosing too much height. 2. tries to remedy that height loss to some extent by relying on (hopefully) still sufficient control authority. 3. was invented and heavily promoted by Swiss pilot and trainer Claude Vuichard in order to minimize height loss.

All three techniques have been simulated and are compared in the following. Unfortunately, no specific control inputs were available for all methods, so the intent was not to optimize individual schedules, but keep them comparable to the extent possible.

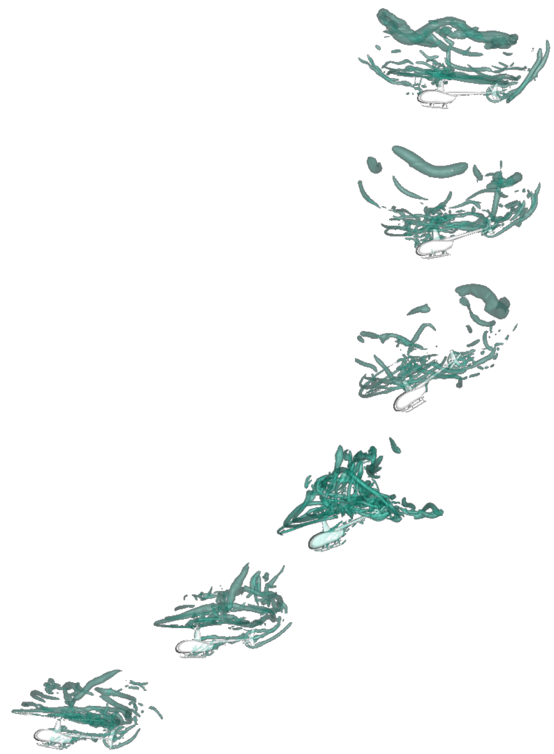

Figure 12. VRS classic recovery trajectory and attitude

Classic recovery In this case, collective was reduced by another -1.5°, and forward cyclic of 5° applied for one second, which initiated nose down pitch. After taking back the cyclic to neutral for another second a forward velocity of $8\,{}^{\mathrm{m}}/_{\mathrm{s}}$ had built up at nose down of -34°, and then collective was increased back to the hover value. Lateral cyclic and pedal wer controlled by the pilot model during the whole maneuver, and longitudinal after having reached hover collective as well. Fig. 12 shows snapshots of the trajectory, one second apart each after initiating recovery.

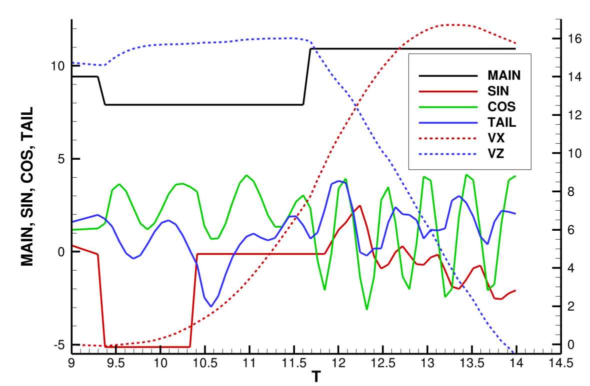

Figure 13. Control inputs for Classic recovery, and resulting forward and downward speed

Again, considerable oscillations developed, see Fig. 13, but kept stable during the maneuver. Zero descent rate was achieved after 4.6 s at forward speed of about 16 m/s (30 kts) and a height loss of 54 m/180 ft, significantly lower than the "panic" reaction of pulling collective alone. Yaw and roll attitudes were kept fairly low by the pilot model.

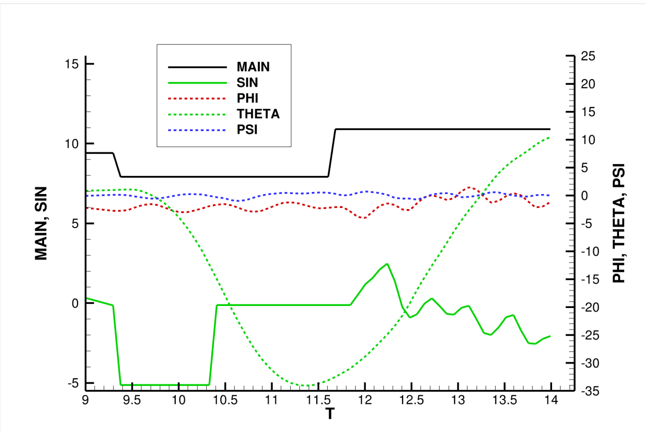

Figure 14. Scheduled control inputs for Classic recovery, and resulting airframe attitude

Airbus recovery Airbus recommends to pull collective immediately, so the hover value was established together with forward cyclic of 5° applied for one second, also causing nose down pitch. Afterwards, cyclic was taken back to neutral until having reached a forward velocity of 10 m/s, which happened 2.8 s after initiating recovery. Maximum nose down pitch achieved during that time was −25°. Lateral cyclic and pedal were controlled by the pilot model during the whole maneuver, and longitudinal after the prescribed schedule as well.

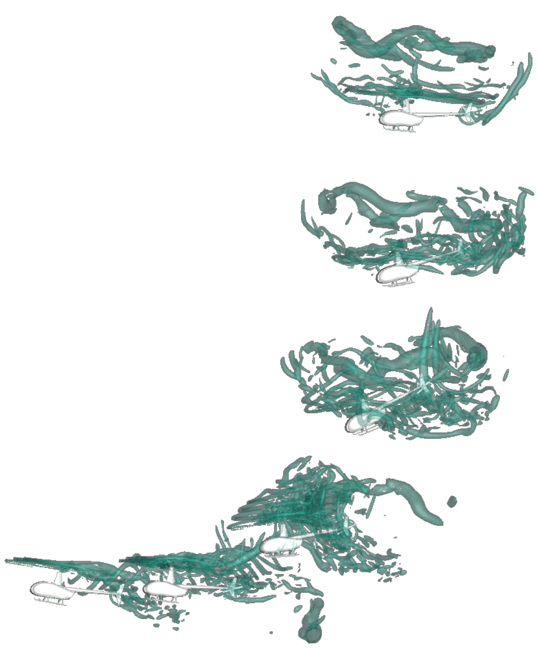

Figure 15. VRS Airbus recovery trajectory and attitude

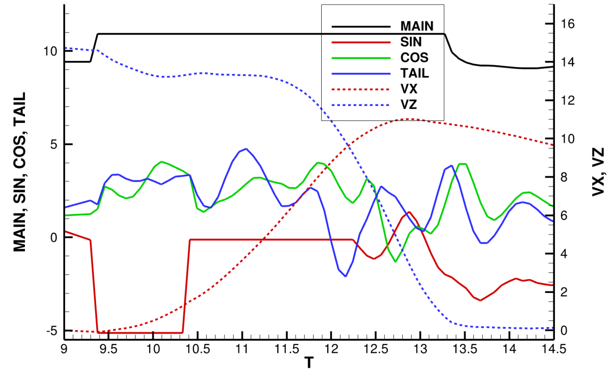

Figure 16. Control inputs for Airbus recovery, and resulting forward and downward speed

Oscillations were slightly smaller in this variant, and zero descent rate was achieved after 4.0 s at forward speed of about 10 m/s (20 kts) and a height loss of 43 m/140 ft. This indicates that indeed the Airbus recommendation to pull collective immediately seems to reduce the height loss by about 11 m/40 ft.

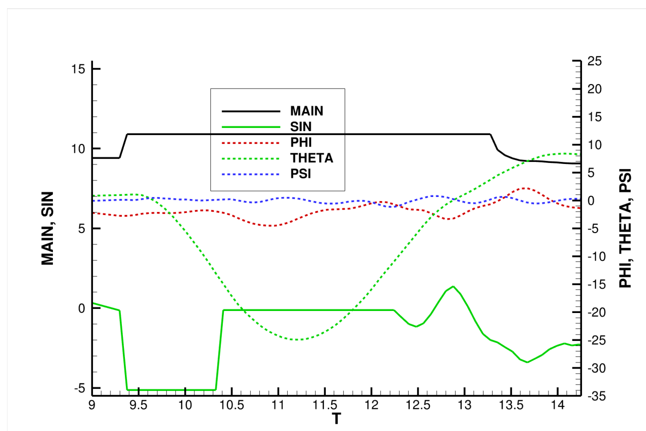

Figure 17. Scheduled control inputs for Airbus recovery, and resulting airframe attitude

Vuichard recovery The Vuichard method consists of pulling collective immediately "to the power limit". As this was hard to establish, again the hover value was chosen for collective, together with right cyclic of 3° applied for a half second. Due to the low inertia in roll direction this was already sufficient to achieve a roll angle of 25° to the right. Again, cyclic was taken back to neutral until having reached a lateral velocity of 8 m/s, which happened 2.0 s after initiating recovery. Then, to stop the rolling motion, 0.3 s of 3° left lateral cyclic were applied, and half a second of 1° to the right, before going back to neutral for another half second, followed by a half second

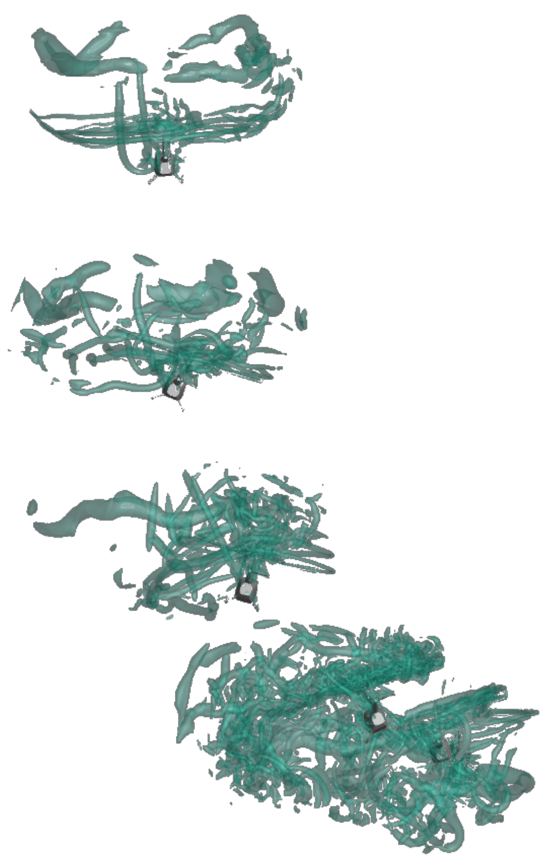

Figure 18. VRS Vuichard recovery trajectory and attitude

of 2° right and one more of 1.5° right. This complex lateral schedule can certainly be more streamlined and improved to quickly achieve zero roll rate and normal roll attitude more quickly and smoothly, resembling a standard side step, but this optimization does not change results considerably. Actually, to achieve zero descent rate at the later stages with low lateral velocity, main rotor collective had to increased as well to increase thrust to sufficient levels. Similarly to the other methods, longitudinal cyclic and pedal were controlled by the

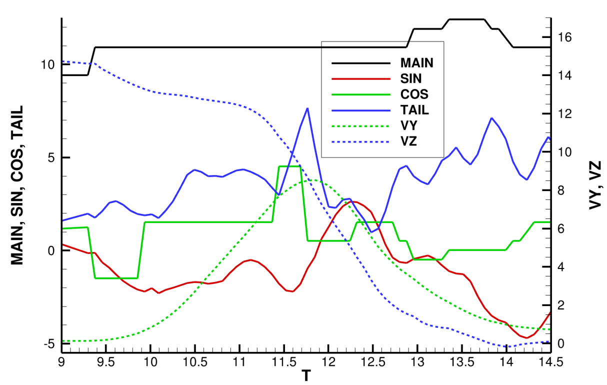

Figure 19. Control inputs for Vuichard recovery, and resulting right and downward speed

pilot model during the whole maneuver. The trajectory snapshots are taken from behind this time, again at one second intervals after initiation.

This time, oscillations had more low frequency content, possibly indicating macroscopic control issues instead of suboptimal pilot model parameters. Furthermore, to keep yaw angle low, obviously increased tail rotor pitch is required, which also increases power demand ("power pedal"). Zero descent rate was achieved after 3.9 s at lateral speed already reduced back again to 2 m/s (4 kts) and a height loss of 36 m/120 ft with lateral discplacement of 20 m. Compared to the classic recovery technique this reduces the height loss by about 18 m/60 ft, and is still somewhat better than the Airbus method, at least with the simulated parameters.

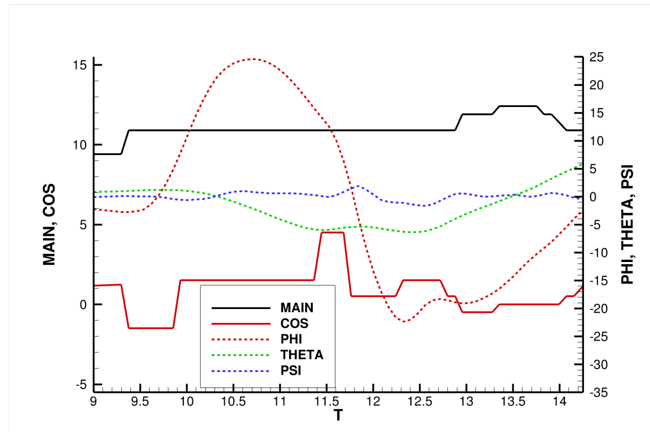

Figure 20. Scheduled control inputs for Vuichard recovery, and resulting airframe attitude

However, two challenges remain with Vuichard, as already indicated by the description. The increased tail rotor thrust adds to the power budget, which may or may not be sufficient at the time of VRS occurence. Furthermore, complex control schedules had to be applied in this case. While they may be simplified to achieve similar results, there is at least an indication of increased pilot workload, which is certainly unfortunate in a potentially life threatening flight situation. Consequently, there is not a definite command to prefer Vuichard to Airbus in case of VRS, but needs to be further investigated.

CONCLUSIONS

Summing up, a major step forward has been commenced towards high fidelity simulation of complex maneuver flight situations. In contrast to previous attempts of just putting together individual pieces of stationary flight, the dynamics of the instationary motion have been taken into account by means of inertial forces and moments, including gyroscopic effects. Main takeaways include

- Inertial loads can be large and fast changing.
- Gyroscopic effects of the rotors, especially the main rotor, need to be incorporated.

- Blade flapping changes during very short periods of time, altering airloads considerably.
- Tip path plane tilting moments required by the gyroscopic effects of the rotors are provided by a flapping motion lagging behind the control inputs, compared to stationary flight state behavior.
- As a consequence, it is probably to some extent sufficient to assume an instantaneous adjustment of flapping to the static value and at the same time neglect gyroscopic rotor moments, if not specifically interested in the time-accurate behaviour of flapping angles.
- Stochastic variance of CFD airloads during aggressive instationary motions poses some challenges.
- An advanced pilot model and/or flight controller needs to be developed, with high stability requirements due to the previous point of short-term airloads uncertainty.
- Realistic control schedules for target trajectories should be pre-computed using fast low-fidelity tools, possibly with continuous updates during simulation runtime.

Following all that recommendations, the introduced example maneuvers show the capability of current state of the art high fidelity CFD to address advanced flight mechanical questions. Current computing technology allows for a few seconds of real flight time within a day of simulation on a high performance computer, so a full maneuver may take one or two weeks. Within a couple of years, an order of magnitude improvement seems possible, such that one maneuver may run overnight. This allows for robust and reliable engineering assessment even for challenging flight states with highly interacting flow phenomena for complex configurations, such as future advanced air mobility. Other than mid-fidelity tools, which can be made very accurate by tuning according to flight test data, such first-principles simulations can be run long before actual hardware exists, and thus utilized for design improvements and risk reduction at very early development stages, saving time and money long before manufacturing, test and deployment.

REFERENCES

1. Johnson, W., *CAMRAD II Comprehensive analytical model of rotorcraft aerodynamics and dynamics*, fourth edition, 2009.
2. Benoit, B., Dequin, A.-M., Kampa, K., von Grünhagen, W., Basset, P.-M., and Gimonet, B., "HOST, a General Helicopter Simulation Tool for Germany and France," *American Helicopter Society, 56th Annual Forum, Virginia Beach, Virginia*, 05 2000.
3. Saberi, H., Khoshlahjeh, M., Ormiston, R. A., and J, R. M., "Overview of RCAS and Application to Advanced Rotorcraft Problems," 4th Decennial Specialist's Conference on Aeromechanics of the American Helicopter Society, San Francisco, California, 2004.
4. Embacher, M., Keßler, M., Dietz, M., and Krämer, E., "Coupled CFD-Simulation of a Helicopter in Free-Flight Trim," 66th AHS Annual Forum, Phoenix, Arizona, 2010.
5. Bhagwat, M. J., Ormiston, R. A., Saberi, H. A., and Xin, H., "Application of Computational Fluid Dynamics/Computational Structural Dynamics Coupling for Analysis of Rotorcraft Airloads and Blade Loads in Maneuvering Flight," *Journal of the American Helicopter Society*, Vol. 57, (3), 2012, pp. 1–21.
6. Garcia, A. J. and Barakos, G. N., "NUMERICAL SIMULATION OF THE ERICA TILTROTOR USING THE HMB2 SOLVER," 41st European Rotorcraft Forum, Munich, Germany, 2015.
7. Abhishek, A. and Baeder, J., "Prediction and Fundamental Understanding of Stall Loads in UH-60A Pull-up Maneuver," 66th AHS Annual Forum, Phoenix, Arizona, 2010.
8. Thomas, S., Ananthan, S., and Baeder, J. D., "Wake-Coupling CFD-CSD Analysis of Helicopter Rotors in Steady and Maneuvering Flight Conditions," AHS Specialists' Conference on Aeromechanics, San Francisco, California, 2010.
9. Crozon, C., Steijl, R., and Barakos, G., "Coupled flight dynamics and CFD – demonstration for helicopters in shipborne environment," *The Aeronautical Journal*, Vol. 122, (1247), 2018, pp. 42–82.
10. Tran, S. A. and Yeo, H., "Transient and Quasi-steady Numerical Simulations of Tiltrotor Conversion Maneuvers," *Journal of the American Helicopter Society*, Vol. 68, (2), 2023, pp. 18–32.
doi: 10.4050/JAHS.68.022003
11. Hong, S. H., Kim, Y. J., Park, S. H., Jung, S. N., and Kim, K. R., "Air and Structural Loads Analysis of a 5-Ton Class Rotorcraft in a Pull-Up Maneuver Using CFD/CSD Coupled Approach," *Aerospace*, Vol. 11, (2), 2024.
doi: 10.3390/aerospace11020102
12. Perez, D. G., Diaz, P. V., Sanguinetti, A., and Yoon, S., "Tiltwing Transition Flight Analysis Using High-Fidelity CFD," 80th VFS Annual Forum, Montreal, Canada, 2024.
13. Kroll, N., Eisfeld, B., and Bleeke, H., "The Navier-Stokes code FLOWer," *Portable Parallelization of Industrial Aerodynamic Applications (POPINDA)*, edited by A. Schüller, Vol. 71, Notes on Numerical Fluid Mechanics, Springer International Publishing, 1999, pp. 58–71.
14. Kowarsch, U., Öhrle, C., Hollands, M., Keßler, M., and Krämer, E., "Computation of Helicopter Phenomena Using a Higher Order Method," *High Performance*

Computing in Science and Engineering '13, edited by W. E. Nagel, D. H. Kröner, and M. M. Resch, Springer Verlag, 2013, pp. 423–438.

15. Kranzinger, P. P., Kowarsch, U., Schuff, M., Keßler, M., and Krämer, E., "Advances in Parallelization and High-fidelity Simulation of Helicopter Phenomena," *High Performance Computing in Science and Engineering '15*, edited by W. E. Nagel, D. H. Kröner, and M. M. Resch, Springer International Publishing, 2016, pp. 479–494.

16. Klauck, J., Keßler, M., and Krämer, E., "Time Accurate Fluid-Structure Coupling Employing a Lightweight Socket Based Data Exchange," 50th European Rotorcraft Forum, Marseille, France, 2024.

17. Schäferlein (né Kowarsch), U., Keßler, M., and Krämer, E., "Aeroelastic Simulation of the Tail Shake Phenomenon," *Journal of the American Helicopter Society*, Vol. 63, (3), 2018, pp. 1–17.

18. Wengrzyn, O., Frie, L., Keßler, M., Krämer, E., and Eberhard, P., "Tight Coupling of Helicopter Airframe Including Elastic Main Rotor Shaft," 50th European Rotorcraft Forum, Marseille, France, 2024.

19. Padfield, G., *Helicopter Flight Dynamics: The Theory and Application of Flying Qualities and Simulation Modeling*, AIAA education series, American Institute of Aeronautics and Astronautics, second edition, 2007.