CFD-based Simulation and Experiment in Helicopter Aeromechanics

R. E. Brown[1] S. S. Houston[2]

1 Department of Aeronautics, Imperial College, London SW7 2AZ,
corresponding author, r.e.brown@imperial.ac.uk
2 Department of Aerospace Engineering, Glasgow University, Glasgow G12 8QQ,
stewart@aero.gla.ac.uk

keywords: helicopter rotor aerodynamic simulation

Abstract

Accurate simulation of the aeromechanics of helicopters poses significant challenges. Satisfactory agreement between the predictions of numerical models and experimental data requires that the model successfully capture a broad range of complex physical interactions between the aerodynamics and the mechanics of the system. In most practical cases, the physical processes at work cannot be determined uniquely from the data at hand, reducing the process of constructing an accurate simulation to one of educated guesswork. On the basis that improved simulation fidelity might be obtained simply through improved modelling of isolated aspects of the physics of the problem, a helicopter aeromechanics simulation model was constructed in which the physical realism of the model of the helicopter wake could be varied. Experiences in validating this numerical test bed against real-world data, from both flight test and laboratory experiments, are generalised into set of more widely-applicable observations regarding the relationship between complex dynamic models and the physical reality that they are supposed to represent.

Introduction

Simulation of the flight dynamic performance of helicopters and other rotorcraft still poses significant difficulties to the analyst. The rotorcraft operates in a particularly complex aerodynamic environment that is dominated by the wakes generated by its rotors. These wakes have structure on length scales ranging from the blade chord to the rotor diameter and are strongly time-dependent. In addition, wake-wake and wake-rotor interactions introduce strong coupling between the dynamics of geometrically well-separated parts of the system. Even in the absence of these couplings, the dynamic behaviour of a rotorcraft can be very complex. Almost all rotorcraft can be characterised crudely as a collection of more-or-less flexibly-coupled bodies, all of which can undergo significant elastic deformations and can experience large excursions from their nominal equilibrium positions. The difficulties inherent in modelling such dynamical systems, together with the problems introduced by more subtle physical effects associated with powerplant, control systems and the human element, conspire to frustrate the analyst in his attempts to construct simulation tools which are valid over the broad range of timescales - from vibrational to quasi-static - that are of engineering interest.

The fundamental problem in rotorcraft simulation is one of model fidelity - usually expressed in terms of the bandwidth over which a set of representative transfer functions relating the excitation to the response of the simulated and real systems agree to within acceptable bounds. Historically, the resolution of inadequacies in simulation models has been approached by increasing the complexity of the model, either through direct addition of physical models for new phenomena, or through modification of the structure of an existing feature to encapsulate fundamentally more complex behaviour. Padfield,[1] in 1988, attempted to produce a practically implementable path leading to improved fidelity of rotorcraft simulation models by proposing a hierarchy of modelling 'levels', based on a direct enumeration of the physical effects accounted for during any given simulation and on some measure of the quality of the model used to capture each individual effect, towards which simulationists could strive. The underlying assumption of his strategy was that the incorporation or addition of specific modelling features could be related directly to issues of amplitude and frequency, and thus that modelling enhancements could be introduced in a systematic, step-by-step way with concomitant effect on the fidelity of simulations. For many years, limited progress was made in developing high-fidelity simulation tools for rotorcraft despite many serious attempts to implement this hierarchical strategy.

During the late 1990s the crisis in rotorcraft simulation focused on the resolution of an apparent contradiction in the predicted off-axis response of a rotorcraft to longitudinal or lateral control inputs. Although the focus provided by this particular issue motivated a significant body of work attempting to resolve whether the discrepancy was aeroelastic or aerodynamic in origin, no one study was ever able to provide a completely convincing physical argument explaining the underlying discrepancy between simulation and measurement. Many of the simulations used in these studies had at their core, though, a particularly simple representation of the wake dynamics that, although well-validated under benign flight conditions, was known (although not widely acknowledged) to be somewhat deficient in its representation of certain morphological and evolutionary features of the rotor wake. It was widely held within the simulation community that simplified representation of the rotor wake was unavoidable (and acceptable) if tractable simulation run-times were to be achieved on the computer hardware that was available at the time.

One of the current authors (Houston), through his access to data from the highly instrumented SA330 Puma helicopter formerly operated by the Defence Evaluation and Research Agency at Bedford and his experiences in correlating this data with numerical simulation,[2] hypothesized that the magnitude of the discrepancy between simulated and measured flight test data could be correlated with the extent to which modelling of the real-world aerodynamic environment of the rotors (in terms, for instance, of main-tail rotor wake interactions and blade-vortex interactions) was being sacrificed by simplifying the representation of the rotor flowfield. As an example of practice at the time, Fig. 1 compares the azimuthal variation of blade loading, generated by a contemporary flight dynamic simulation model, against measured data reconstructed from leading-edge pressure measurements on the full-size machine at moderate forward speed.[3] Given the bland character of the numerical predictions of the aerodynamic environment of the rotor blades, missing both qualitatively and quantitatively the features of the flight test data (in particular the significant distortion by vortex interaction effects on the forward half of the rotor and interference from the tail rotor on the rear half of the rotor disc) a strong case could be made that improved numerical treatment of the rotor wake would be a necessary (if not sufficient) contributor to any significant future advance in the fidelity of flight dynamic simulations of rotorcraft.

Realising (gambling, perhaps) that the very rapid expansion in cheaply available computational power would soon enable full CFD-type modelling of the rotor wake - even within the context of flight dynamic simulations - the two authors of this paper constructed a testbed for flight dynamic simulations of rotorcraft in which the effect of wake modelling fidelity on the validity of simulations could be investigated.[3] Their experiences in correlating this testbed against flight test data and laboratory experiment challenge the root assumption that improvements in modelling fidelity can be obtained via the building-block approach of incorporating, in sequential fashion, a range of seperable physics models. This paper describes some of the authors' experiences in validating their model against real-world data, and attempts to distil their experiences into a set of more widely-applicable observations regarding the relationship between complex dynamic models and the physical reality that they are supposed to represent.

CFD-based Flight Mechanics

The basic constituent of any comprehensive model for rotorcraft simulation is a numerical representation of the dynamics of the rotorcraft under the action of any applied aerodynamic (and possibly propulsive) forces. The RASCAL model, developed at Glasgow University, is one well-established such model and has been used previously for helicopter validation and simulation studies[4,5] and for the simulation of autogyros.[6]

In RASCAL, the equations of motion of the rotorcraft are cast into the standard state-space form

$$\dot{\mathbf{x}} = \mathrm{f}\left(\mathbf{x}, \mathbf{u}, \mathbf{F}\right) \tag{1}$$

which is then integrated numerically to obtain the unsteady motion of the rotorcraft as a function of time. The explicit form of f is constructed by assuming the helicopter to behave dynamically as a set of multiple rigid blades connected to a rigid-body fuselage via appropriate hinges.[4] The state vector $\mathbf{x}$ contains the air-frame translational and angular velocity components, blade flap, lag and feather angles and rates for each blade on each rotor, the angular velocity of each rotor, and the engine torques. The control vector **u** is aircraft configuration-specific, but for conventional single main and tail rotor configurations there are three main rotor controls and one tail rotor control. The all-important forcing vector $\mathbf{F}(v)$ is constructed from the aerodynamic loads on the system via the velocity field v surrounding the rotorcraft.

Up until very recently, the standard approach to incorporating the effects of the rotor wake on the aerodynamic loading on the helicopter during flight dynamic simulations has been to use some variation on Peters' so-called 'dynamic inflow' formalism.[7,8] The advantage of this formalism is that it is computationally efficient, affords some representation of the time-lags inherent in the rotor wake when calculating the aerodynamic loads generated on the rotors, and is compatible with the state-space form of the dynamic equations for the remainder of the rotorcraft model.

The dynamic inflow formalism relies on a modal expansion of the velocity surrounding the rotorcraft as

$$v(t) \; = \; \mathbf{a}(t) \; \cdot \; \mathbf{V} \tag{2}$$

where **V** is a *finite-dimensional* vector of velocity distributions in space. The evolution of the vector of inflow states **a** is governed by a nonlinear first order equation, forced by the aerodynamic loading on the system, of the form

$$[\tau(\mathbf{a},\dot{\mathbf{x}}_{rotor})]\;\dot{\mathbf{a}}+\mathbf{a}=[\mathbf{L}(\mathbf{a},\dot{\mathbf{x}}_{rotor})]\;\mathbf{F} \tag{3}$$

Wake distortions are modelled by allowing the matrix of time constants $[\tau]$ and the inverse gain matrix $[\mathbf{L}]$ to depend on the rates of change $\dot{\mathbf{x}}_{rotor}$ of the states defining the rotor orientation.[9]

On physical grounds, the main argument against the dynamic inflow formalism is that adoption *a priori* of a dynamic equation of the form of Eq. 3 yields a model that is too simple to represent properly the dynamics of the wake. In particular, if, as is implied by the form of Eq. 3, the vector of inflow states is forced solely by the aerodynamic loading on the rotor blades, then the formalism cannot retain sufficient off-rotor information to represent fully the dominant physical mechanism of vorticity convection in the wake.[3] Convection takes place largely independently of the instantaneous rotor loading, but governs very strongly the blade-vortex interactions and wake rollup that in turn dominate the rotor loading. In the absence of a proper treatment of wake convection, these physical effects are simply not captured by the approach. It is thus to be expected that a model which is capable of more faithfully and robustly representing the fluid dynamic processes occurring in the wake should yield 'better' simulations of rotorcraft flight dynamics in some sense, especially in the presence of strong aerodynamic interactions between the wake and the rotor blades.

Arguably the most complete (and certainly the most straightforward) way to model the vorticity-dominated aerodynamic environment of a helicopter rotor is to model the rotor wake directly as a time-dependent vorticity distribution in the region of space surrounding the rotor. If v is the flow velocity, then the associated vorticity distribution $\omega = \nabla \times v$ evolves according to the unsteady vorticity transport equation

$$\frac{\partial}{\partial t}\omega + v\cdot\nabla\omega - \omega\cdot\nabla v = S \tag{4}$$

This equation can be derived from the incompressible Navier-Stokes equation in the limit of zero viscosity. The differential form of the Biot-Savart equation then relates the velocity and vorticity fields throughout the flow:

$$\nabla^2 v = -\nabla\times\omega \tag{5}$$

In our vorticity transport-based approach to wake modelling,[3,10] a direct computational solution of Eq. 4 is employed to simulate the evolution of the helicopter flow field. After enclosing the helicopter and its surroundings within a three-dimensional, structured grid of computational cells, the vorticity distribution in the flow is advanced through time using a computational discretisation of Eq. 4, using Toro's Weighted Average Flux (WAF) algorithm[10] to construct the inter-cell vorticity fluxes. An interesting property of this approach is that, when a suitable flux limiting function is used in conjunction with the WAF algorithm, diffusion of vorticity can be controlled to the extent that vortical structures in the flow are preserved - even during very long computations.[11]

The operator $v = v(\omega)$ required to evaluate the velocity on which the rotor loads depend (and at which the vorticity is advected) is constructed by inverting Eq. 5 using cyclic reduction.[10] The wake evolution is coupled into the aerodynamic loading, and hence the dynamics, of the rotor system by defining the source term S in Eq. 4 in terms of the shed and trailed vorticity from the rotor blades as follows:

$$S = -\frac{d}{dt}\omega_b + v_b\nabla\cdot\omega_b \tag{6}$$

In RASCAL, the resultant blade loads are evaluated using a blade element approach, and the bound vorticity ω_b is explicitly dependent on the flow velocity v_b relative to the blades.

In practice the RASCAL model can be run with the rotor velocity field $v(t)$ calculated either using a three-state implementation of the dynamic inflow formalism (applied independently to each rotor) or with the velocity field calculated (in global fashion) using the vorticity transport apprach. As an initial test of the hypothesis that improved wake modelling lay at the root of improved flight dynamic simulation of rotorcraft, the predictions of RASCAL when run in either configuration were compared against flight test data for a full-scale Puma helicopter.[3] Figure 2(a), especially when compared with Fig. 1(a), gives some indication of the promising results that were expected from the RASCAL-vorticity transport combination, particularly given the appearance of certain features in the blade loading that were missing from the earlier calculations using the RASCAL-dynamic inflow combination. Particularly notable are the sharply-defined ridges in the blade loading induced by strong blade-vortex interactions on the forward half of the rotor disc - especially when it is borne in mind that the vorticity transport formalism was adopted specifically to allow features of this type to appear in simulations of the rotor loading.

Comparison with Flight Test Data

The aircraft used for the flight experiments was the SA330 Puma helicopter formerly operated by the Defence Evaluation and Research Agency at Bedford. This aircraft was fully instrumented with two separate inertial systems for measuring translational accelerations, angular rates and attitudes. A probe-mounted air data system measured airspeed, sideslip and angle of attack. Pilot control positions were also recorded. The flap, lag and feather angles of each blade were measured with potentiometers, and local blade angle of attack could be inferred from a suite of pressure sensors and strain gauges attached to one of the blades.[2] Steady flights were conducted between hover and 150 knots, at a nominal altitude of 3000 feet and at weights between 5100 and 5800 kg.

All numerical simulations were configured with the actual flight conditions and weights recorded at each test point. The resultant correlations between simulation and test data could be divided, fairly simplistically, into two distinct classes, as illustrated by the sample variations with forward speed of some of the state variables of the system shown in Fig. 3. In the first class could be placed those results where a marked difference (*e.g.* as shown in Fig. 3(a)) between the predictions obtained using the dynamic inflow or vorticity transport formalisms was observed (somewhat disappointingly, not always with correlation in favour of the vorticity transport formalism). Most surprisingly, a second class of results could be identified (*e.g.* as shown in Fig. 3(b)) where the choice of wake model had minimal to negligible effect at all on the predicted behaviour of the system - even in certain cases where the state variable in question was believed to be strongly influenced by the wake aerodynamics. Although the support for the importance of accurate wake modelling that could be extracted from these results was equivocal, certain important features of the data could be identified that did appear to support the hypothesis. The tail rotor operates within the wake generated by the main rotor, but this interference effect is captured only by the vorticity transport formalism. Predictions of the tail rotor collective pitch angle (Fig. 3(a)) obtained with this model appeared to correlate better with test data than the results obtained with the dynamic inflow formalism - especially at intermediate flight speeds where interference between main and tail rotors is strongest. Yet, in all cases, the comparisons between predictions and test data were characterised by good numerical matching between simulation and test at some flight speeds but poor matching of the trend of the data with flight speed, or, alternatively, by good matching of the trend of the data with flight speed but uniformly poor matching of actual numbers. The behaviour of the simulations was somewhat consistent with an error in the specification of the variation of drag with speed of the helicopter's fuselage (even though the drag variation was constructed from wind-tunnel measurements made on a model of an isolated Puma fuselage), but this discrepancy was not investigated with much vigour for reasons that should become clear very shortly.

Given the complexity of both the real and modelled systems, the nagging suspicion must be entertained that any one of a number of equally plausible rationalisations might have been put forward to explain any one of the observed discrepancies between simulation and experiment.

The truth is that it is extremely difficult to make any definite statements as to the origins of the deficiencies in any given model when attempts are made to validate it against the behaviour of a system that is as complex as a real-world, piloted rotorcraft. In validating against full-scale flight test data, it is always the case, of course, that the measured output of the real-world, physical system contains the effects of many unknowns. Some of these unknowns are on the level of unanticipated physical effects, but more pernicious are those unknowns introduced by physical (but unforseen, undetected or undocumented) defects in the system or in the measuring apparatus used to characterise the system. Even more difficult to quantify are the additional complications introduced into the behaviour of the (necessarily coupled) man-machine system, for example by the vagaries of the individual flying technique of the pilot, even though some deference is usually made to simplification of the system in this respect - either by prescribing the flight trajectory to be followed during the flight test or by defining in some fashion the pilot inputs that should be made to execute a given test manoeuvre.

A simple example of the confusion that can be introduced into the validation process by the possible presence of an undocumented physical defect within the system is illustrated in Fig. 4 where measured levels of vibration on the Puma main rotor are compared against RASCAL predictions using both the dynamic inflow and vorticity transport models.[3] The poor performance of the dynamic inflow model in predicting the measured vibration levels is expected since this model does not introduce any representation of effects on vibrational timescales into the azimuthal variation of the blade loading. The vorticity transport model does a somewhat better job at predicting the principal magnitude and frequency of the vibration, yet misses the significant once-per-revolution component of the measured signal. The most logical explanation for this feature in the measured data is presence of an undocumented error in the tracking of one of the blades. Undoubtedly this type of defect is very commonly encountered during practical helicopter operations. Yet, in some rotor systems, odd rotor vibrational harmonics are indeed encountered - usually as a result of some very obscure, but nonetheless intriguing (read “academically-exploitable”), physical effects!

The point is that the physical system cannot readily be simplified, and any complexity in the behaviour of the aircraft in the real world must usually be accepted by the analyst as an inherent characteristic of the system. When seen from the opposite point of view of the analyst trying to construct a valid simulation of the system, a full characterisation of the real system is never available either, and hence the physics that must be included to enable good validation is always a matter of (informed) guesswork. The simulation model thus runs the risk of being driven towards maximum complexity by the simple fact that, to obtain even a remote chance of agreement between computation and test, as complete a physical model of the real system must be constructed as is feasible. It is, of course, arguable that such an approach runs completely counter to the reductionalist underpinnings of currently accepted scientific method. Some element of validity might be retained in the process if each successive layer of physical modelling that was introduced into the simulation did not introduce with it a new set of model-dependent parameters (such as the fuselage drag mentioned earlier) with which to tempt the closet curve fitter. The danger, of course, is that the process of validation of any comprehensive rotorcraft model might degenerate into a festival of tweaked parameters, fudged physics and faulty inference as the simulationist drives his or her pet model to fit, to ever and ever greater precision, a finite data set, perhaps even selected more for its ready availability than for the accuracy of its representation of the system at hand. (Of course, any claims by the authors to immunity from this process may be judged from the presentation of their own results in this paper!)

Nevertheless, despite the objections set out above, it is not yet entirely clear how this situation can be circumvented, especially given the practical requirement for codes that can be used *a priori* to reproduce flight test - for instance to expose the possibility of pathological dynamics before it is actually encountered in the air. The real danger though is obviously not posed by well-correlated models *per se* - after all, even curve fitting is useful if extrapolation is not pursued too far - but by the possibility that ‘model enhancement’ merely introduces additional degees of freedom that, although allowing an improved fit of model to the data, yield a simulation model that has increasingly tenuous grounding in physical reality. This abuse of the validation process is especially worrying given the modern tendency to supplant flight test with simulation - an issue to be discussed further in a later section of this paper. Perhaps the sanest approach might involve a benefits-weighted application of Occam’s razor - or in other words an approach where the simplest explanation consistent with the data is traded off against the value of an explicit number-for-number matching of experimental results - on the principle that an understandable model is in many cases preferable, and is certainly more educational, than a complex, opaque or unpredictably idiosyncratic ‘black-box’ simulation tool.

A somewhat obvious alternative strategy might be to apply the correlation process in simpler circumstances, then to extrapolate the results of these correlations to support the validity of the model when simulating systems as complex as a real-world rotorcraft. Whether or not this strategy achieves any useful purpose will be explored in the following two sections of this paper.

CFD-based Rotor Aeromechanics

Almost five years of development since the original integration of the vorticity transport-based wake model into RASCAL has resulted in a parallel model that can be used for rotor performance calculations, known simply as the Vorticity Transport Model (VTM). This code was initially developed to allow exploration of the properties of the vorticity transport approach in isolation from RASCAL's comprehensive flight mechanics capabilities, but has now been used independently in studies of wake interactions,[12] wake instability[13] and rotor vibration.[14] Apart from the lack of any fuselage dynamic model within the VTM, the principal difference between the implementation of the vorticity transport formalism in this model and in RASCAL is in the treatment of the vorticity source term. In contrast to the explicit model used in the original RASCAL formulation, in the VTM the blade loading and the vorticity source are implicitly coupled, and the blade loading is calculated using the Weissinger-L approximation to lifting line theory instead of using a blade element type approach. These modifications were found necessary to represent properly the unsteady loading on the rotor blades, but arguably the resultant increase in computational effort is only justified at reduced frequencies that are somewhat higher than those that need to be resolved in flight mechanics computations. The latest variants of the VTM now invert the vorticity-velocity relationship using a very efficient algorithm based on the Fast Multipole Method rather than using cyclic reduction, and evolve the flow on a mesh that is no longer structured, but can adapt in form to encapsulate only those regions of the flow that actually contain vorticity.[14] Nevertheless, the basic underlying formalism remains common to both RASCAL and VTM.

Laboratory-Based Validation

An interesting situation arises if the validation exercise described earlier is repeated on a system that is far removed in complexity from full scale flight test. By running a wind-tunnel test on an isolated four-bladed rotor with fixed control angles, Harris,[15] in 1972, produced a set of data for the flapping response of a rotor as a function of forward speed that was soon appropriated by the rotorcraft simulation community as ammunition in the debate surrounding the prediction of the off-axis response of a rotor to longitudinal or lateral control inputs that was mentioned earlier. Figure 5 compares Harris' data to predictions obtained using the three-mode dynamic inflow model (implemented within the VTM's dynamic model), an early version of the VTM, and the latest version of the model. The plot shows how the dynamic inflow approach struggles to represent the magnitude of the lateral disc tilt, especially at low forward speed, and how the vorticity transport formalism comes into its own in this regime because of its sensitivity to the rather detailed and complex structure of the wake at low forward speed. This and other examples are at odds with the rather equivocal results from comparisons against flight test data, since validation in the context of simplified and rather sterile systems such as this one appear to vindicate the hypothesis that improved physical modelling of the wake does indeed yield an improvement in simulation fidelity. The unanswered question as to why the validation process should yield such non-uniform results will be returned to in the conclusion to this paper.

Close examination of the correlation between the predictions of the VTM and Harris' data shows some rather subtle effects. The variance between simulation and theory at high forward speed indicates an obvious deficiency in the blade aerodynamic model, but the variance between simulation and experiment at very low forward speed in the early versions of the VTM was initially thought to be insignificant given the error bounds on the measured data. The source of this variance in the spurious effect of the grid boundaries on the geometry of the wake only emerged on adopting the adaptive grid structure used in later versions of the model. This is a good example of how, even in the context of relatively good correlation (or, arguably, as a *result* of good correlation) a potentially important source of error can slip through the validation process through lack of sufficient care or perhaps even through wishful thinking. That the situation need not be quite as subtle as in this case is best illustrated in a second example.

Early on in the history of modern rotorcraft, Carpenter and Fridovich[16] produced an interesting paper drawing attention to the fact that, if the controls are varied suddenly, significant overshoots in the resultant loading on a helicopter rotor could be produced during the time that it takes for the inflow through the rotor to reach a steady state. Spring-scale measurements of the forces at the head of a rotor on a whirltower were taken at (fairly long) intervals while the collective pitch of the rotor was quickly ramped from one steady state to another, and some indication of the strength of the inflow through the rotor was obtained using smoke visualisation and balsa wind-vanes. This experiment, conducted in the late 1940s, is still accepted as the standard against which direct numerical prediction, or at least inference of the effects, of the unsteady development of the rotor wake should be validated.

Figure 6 shows a comparison between Carpenter and Fridovich's experimental thrust data, and predictions obtained using both the vorticity transport and dynamic inflow formalisms. The good fit between experiment and the dynamic inflow results is not surprising, given that the dynamic inflow formalism was derived largely in response to the specific findings of this experiment. The VTM also shows very good agreement on a superficial level, yet there are features in the predicted response that are qualitatively different to the experiment and that hint at the presence of unmodelled physics in the simulations. Interestingly, the features exhibited by the VTM are common to a fairly wide variety of other wake-modelling formalisms of varying degrees of equivalence. This leads to the curious phenomenon of the 'traditional' or 'accepted' explanation for an observed discrepancy between model and experiment that, in practical terms, really amounts to little more than an excuse for poor corrrelation to which the community is prepared to turn a blind eye. A prime example here is seen in the prediction of the thrust response of the rotor during the initial ramp input to the collective pitch. The overprediction of the thrust and, in particular, the kink in the variation near its maximum, is 'traditionally' attibuted to unmodelled torsional deformation of the rotor blades - yet nowhere in the literature has the effect of blade torsion on this set of data actually ever been quantified, nor, arguably, could it ever be since the structural characteristics of the real system were never fully quantified at the time of the experiment.

An alternative, and simpler, explanation could be merely that that the interval between experimental measurements was simply too long to resolve the worrisome features in the simulations, and, given that the tests were conducted outdoors, random contamination of the measurements by atmospheric gusts cannot be discounted - especially given that the reported data represented the results of isolated tests rather than being assembled as averages over repeated observations. In any case, over-analysis of this experimental data set is really tantamount to abuse, especially since its perceived 'deficiencies' are more a consequence of it never originally being intended for the validation of codes rather than of any particular fault on the part of the experimenters. Unfortunately, practitioners in the field seem, on occasion, all too ready to exhume data that should be left quietly to rest in dignified peace, or to ignore the basic tenet of experimental design that validation should begin with the model followed by experiment conducted in such a way that it is capable of disproving *a posteriori* the assertions of the theory behind the model. A re-stated version of this argument would hold that experimental data obtained without its end-use in mind is incapable of *disproving* the validity of a model, and hence will always be of use in *supporting*, even if only partially, the validity of the model, no matter how bad the model might actually be! In practical terms, this instance is a particular case where a clear need exists to repeat the experiment with more modern measurement techniques and with the needs of the simulationist clearly in mind.

Extrapolation

Given our observations on the apparent non-uniformity of the performance of helicopter simulation models when applied to systems with inherently different complexity, it is interesting to see if there are other situations in the field where similar problems have, or are likely to be, encountered. The final example to be presented in this paper is a rather extreme one illustrating the practical dangers of extrapolating the validation process from one system to another, even when both systems are characterised by much the same degree of physical complexity.

Currently under evaluation is the idea that helicopters might be used to exploit airspace that is inaccessible to fixed-wing aircraft, and thus that they might be used to alleviate congestion as our airports become clogged with ever more aircraft that are forced to adhere to well-defined traffic patterns. A major concern, though, is that helicopters might be exposed to the same hazards posed by encounters with wake turbulence that currently concern the operators of fixed-wing aircraft. Because of the inherently different dynamics of a rotor compared to a wing, a helicopter responds in a fundamentally different way to a wake encounter than does a fixed wing aircraft.

Surprisingly, especially given the practical importance of a proper quantification of the severity of wake encounters, very little experimental work in this field has been done. In the open literature there exists a single data set from an experiment conducted in the early 1970s when a helicopter was flown through the wake of an aircraft, the size and weight of which is hardly representative of current practice. Quite surprisingly, most of our ‘understanding’ of the severity of wake encounters has been obtained through a long series of essentially unvalidated computer simulations, conducted over the years using a variety of ‘independently’ verified and reliable models for the behaviour of the helicopter and for the dynamics of the wake of the fixed-wing aircraft. On closer examination, all of these studies have relied on the underpinning, yet unstated and implicit, assumption that interactions occur so quickly that the wake of the interacting aircraft does not have time to deform during the course of the interaction. Yet a recent set of studies[11] shows that if this so-called ‘frozen vortex’ assumption is relaxed, then the interaction is significantly different both in character and, indeed, in severity (see Fig. 7).

This example goes to show how the overt use of computer simulations, especially in this rather incestuous and self-referential way, cannot protect the analyst against the contaminating effects of unmodelled or poorly represented physics, and that verification of models is no guarantee of adequate performance especially when model validity is extrapolated from marginally related situations. Particularly worrying in this vein is the current trend to supplant flight test with computer simulation - usually on the basis that most of the underlying physics is ‘understood’ and hence will emerge naturally from a suitably complete physical model of the system. Unfortunately the notion of ‘completeness’ in this context would seem to rely more on faith in some cases than on a completely logical scientific foundation. We would go so far indeed as to argue that such a notion is logically indefensible in the absence of fundamental advances in our understanding of such basic physical processes as turbulence (particularly in a rotating, highly accelerated frame of reference) and viscous separation, the emergence of large-scale effects from unresolvably small-scale effects (rotor wake instability springs to mind), and in the absence of a global understanding of the non-linear mechanics of multidimensional dynamical systems.

Conclusions

The reason why, at least in our experience, a reductionalistic approach to validation via a series of successful correlations against simple, well-defined, experimental systems translated into rather unsatisfactory prediction of the behaviour of real-world systems has yet to be resolved satisfactorily. We believe, though, that the answer might lie in the unforseen emergence of highly-coupled, interactive, essentially non-separable physical effects as the complexity of the system is increased from laboratory to flight scale. These intermediate-scale physical effects have remained essentially unmodelled in most flight dynamic simulations of rotorcraft - largely as a consequence of the widespread acceptance of the validity of the building block approach to the enhancement of modelling fidelity that was described in the introduction to this paper. That this approach would not yield a viable path to improvements in model fidelity is not entirely unsurprising (in retrospect!) given the inherent nonlinearity of the rotorcraft system, the hints given by the almost continual emergence of structure, particularly in the rotor wake, as one steps back from a component-based view towards a more holistic view of the system, and given what little is indeed known about the physical processes cited at the end of the last section.

In our particular case, we had demonstrated that the enhancements that we introduced into our simulation were indeed able to embody some of these interactive, intermediate scale physical effects - but only as far as the wake was concerned. The most likely reason why correlation against data such as those from Harris’ and Carpenter and Fridovich’s experiments was reasonably successful is that those systems were simple enough dynamically for strong interactions to be contained, or isolated, *within* a fairly limited part of the physical system - in these particular cases, within the wake itself. In general, though, the dynamics of the wake is strongly coupled to the local aerodynamic loading on the blades (via the source term in the vorticity transport equation) and, thus, to the aeroelastic response of the blades and hence back to the rotor dynamics and loading (in response to the velocity of the blades relative to the air). Admittedly this description is still rooted in a building-block type appreciation of the complexities within the system, but it seems plausible that, in the context of validation against the flight test data presented earlier, the strong coupling *between* relatively isolated parts of the system that would have resulted from this wake-induced feedback might have have been responsible for drawing the accuracy of the entire calculation down to, at best, the accuracy of the poorest physical model contained within the simulation. Interestingly, this analysis begs the question as to whether simulations containing less comprehensive representations of the interactions within the system might not yield, on average, ‘better’ correlations with real data through their not being as susceptible to this ‘bootstrapping’ type of catastrophe than more comprehensive models.

If the rotorcraft system is indeed characterised by highly-coupled, non-separable physics, then the building-block type approach that has been adopted up to now in the construction of comprehensive simulation models will not yield much further progress. The resistance of complex systems to revealing their behaviour through reductionalist-type dissection does not bode well for our understanding of these systems and, indeed, argues for simulation models based on a very fundamental level of physical modelling. In the previous section we have argued that fundamental progress in the understanding of some basic physical processes still needs to be made before this type of simulation becomes a reliable practicality. In addition, the complete verification of these models, starting from their basic physics, through intermediate-scale correlations in simplified but representative environments, through to full flight dynamic validation will require a hierarchy of high-quality verification data to be created to fill the gap between the information gleaned from overly sterile laboratory based experiments on the one hand and the overly complicated opposite extreme of the full-scale flight test. The construction and exploitation of experimental systems and numerical test cases that meet the requirements of code validation in terms of rigour of definition, portrayal of relevant physics and isolation from extraneous effects will yield an extreme test of the synergy between experimentalist and simulationist.

References

[1] Padfield, G.D., "Theoretical Modelling for Helicopter Flight Dynamics: Development and Validation," ICAS-88-6.1.3, 1988, pp. 165-177.

[2] Houston, S.S. and Tarttelin, P.C., "Validation of Mathematical Simulations of Helicopter Vertical Response Characteristics in Hover," *Journal of the American Helicopter Society*, Vol. 36, No. 1, 1991, pp. 45-57.

[3] Brown, R. E., and Houston, S. S., "Comparison of Induced Velocity Models for Helicopter Flight Mechanics," *AIAA Journal of Aircraft,* Vol. 37, No. 4, 2000, pp. 623-629.

[4] Houston, S.S., "Validation of a Non-linear Individual Blade Rotorcraft Flight Dynamics Model Using a Perturbation Method," *The Aeronautical Journal*, Vol. 98, No. 977, 1994, pp. 260-266.

[5] Houston, S.S., "Validation of a Blade-Element Helicopter Model for Large-Amplitude Manoeuvres," *The Aeronautical Journal*, Vol. 101, No. 1001, 1997, pp. 1-7.

[6] Houston, S.S., "Validation of a Rotorcraft Mathematical Model for Autogyro Simulation," submitted to the *AIAA Journal of Aircraft*.

[7] Chen, R.T.N., "A Survey of Nonuniform Inflow Models for Rotorcraft Flight Dynamics and Control Applications," *Vertica*, Vol. 14, No. 2, 1990, pp. 147-184.

[8] Gaonkar, G.H. and Peters, D.A., "A Review of Dynamic Inflow Modeling for Rotorcraft Flight Dynamics," *Vertica*, Vol. 12, No. 3, 1988, pp. 213-242.

[9] Krothapalli, K.R., Prasad, J.V.R. and Peters, D.A., "A Generalized Dynamic Wake Model with Wake Distortion Effects," *Proceedings of the American Helicopter Society 54th Annual Forum*, Washington, May 1998, pp. 527-535.

[10] Brown, R. E., "Rotor Wake Modeling for Flight Dynamic Simulation of Helicopters," *AIAA Journal*, Vol. 38, No. 1, 2000, pp. 57-63.

[11] Whitehouse, G. R. and Brown, R. E., "Helicopter Rotor Response to Wake Encounters in Ground Effect," *Proceedings of the American Helicopter Society 57th Annual Forum*, Phoenix, May 2003.

[12] Whitehouse, G. R. and Brown, R. E., "Modeling the Mutual Distortions of Interacting Helicopter and Aircraft Wakes," *AIAA Journal of Aircraft*, Vol. 40, No. 3, 2003, pp. 440-449.

[13] Newman, S., Brown, R., Perry, J., Lewis, S., Orchard, M. and Modha, A., "Predicting the Onset of Wake Breakdown for Rotors in Descending Flight," *Journal of the American Helicopter Society*, Vol. 48, No. 1, 2003, pp. 28-38.

[14] Line, A. J. and Brown, R. E., "High-Resolution Wake Modelling using a Semi-Lagrangian Adaptive Grid Formulation," to be presented at the *29th European Rotorcraft Forum*, Friedrichshafen, September 2003.

[15] Harris, F. D., "Articulated Rotor Blade Flapping Motion at Low Advance Ratio," *Journal of the American Helicopter Society*, Vol. 17, No. 1, 1972, pp. 41-48.

[16] Carpenter, P. J. and Fridovich, B., "Effect of a Rapid Blade-Pitch Increase on the Thrust and Induced-Velocity Response of a Full-Scale Helicopter Rotor," *NACA TN 3044*, November 1953.

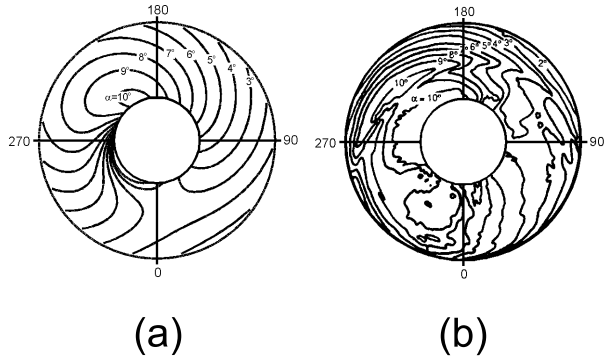

Figure 1: *Azimuthal variation of loading on a Puma main rotor in forward flight, expressed in terms of the angle of attack experienced by a single blade as it rotates once around the rotor hub. (a) RASCAL simulation with a dynamic inflow representation of the rotor wake - the state of the art in the 1990s. (b) A typical angle of attack variation as inferred from flight test measurements.*[1]

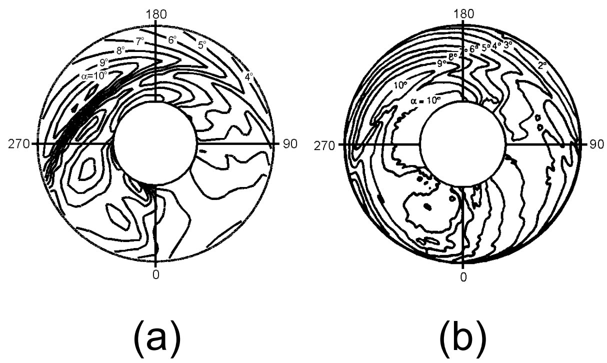

Figure 2: *Azimuthal variation of loading on a Puma main rotor in forward flight, expressed in terms of the angle of attack experienced by a single blade as it rotates once around the rotor hub.(a) RASCAL simulation with a vorticity transport based representation of the rotor wake. (b) A typical angle of attack variation as inferred from flight test measurements.*[1]

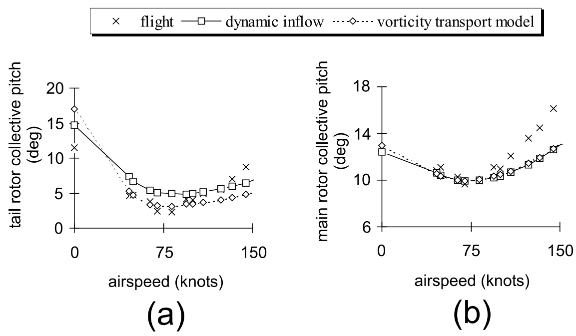

Figure 3: *Example correlations between RASCAL simulations and flight test data. (a) An example where the choice of wake model has significant effect on the results of the simulation. (b) An example where the choice of wake model has insignificant effect on the results of the simulation.*

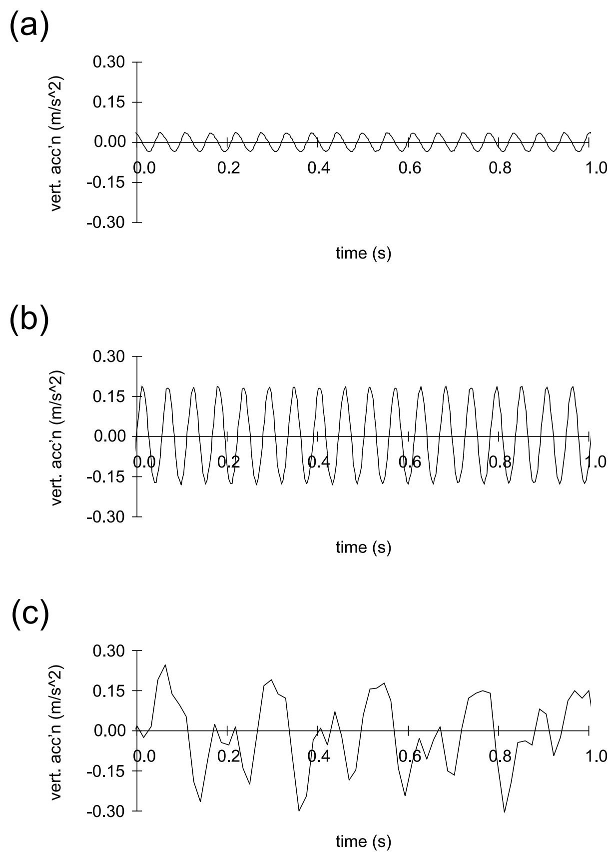

Figure 4: *Vibrational loading induced on a Puma fuselage in forward flight. (a) RASCAL simulation with a dynamic inflow representation of the rotor wake. (b) RASCAL simulation with a vorticity transport based representation of the rotor wake. (c) A typical variation as obtained from flight test measurements.*

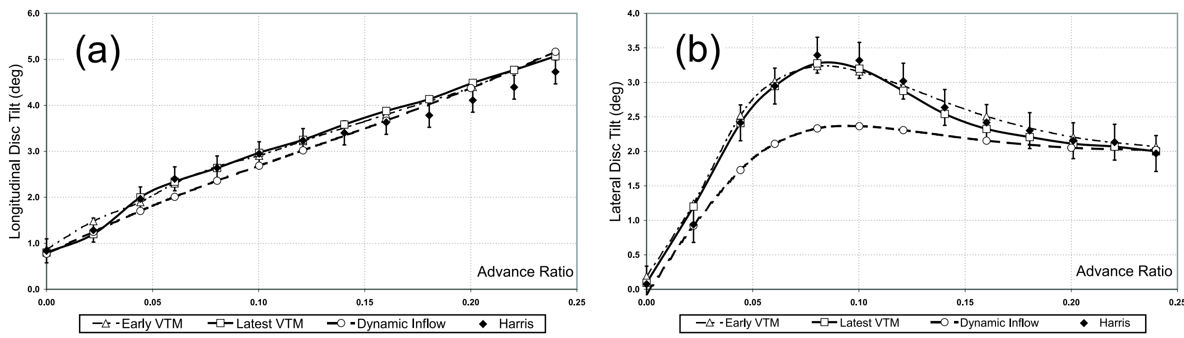

Figure 5: *Correlation between simulation and Harris' experimental data*[15] *for the disc tilt of an isolated rotor as a function of forward speed. (a) Longitudinal disc tilt (b) Lateral disc tilt.*

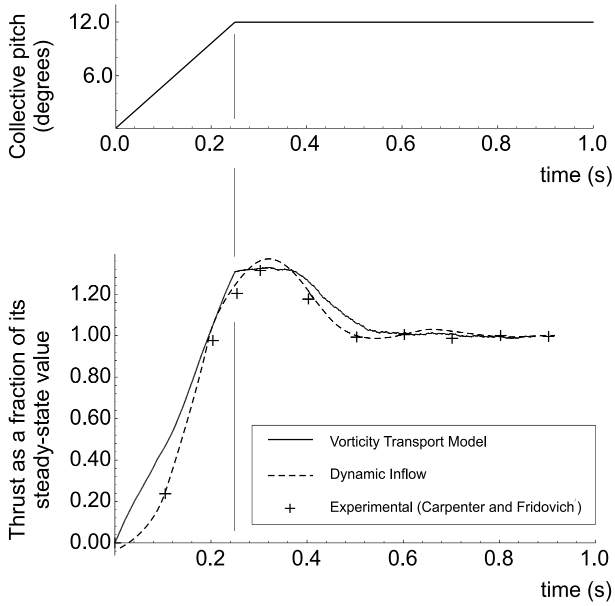

Figure 6: *Correlation between simulation and Carpenter and Fridovich's experimental data[16] for the response of an isolated rotor to a sudden control input.*

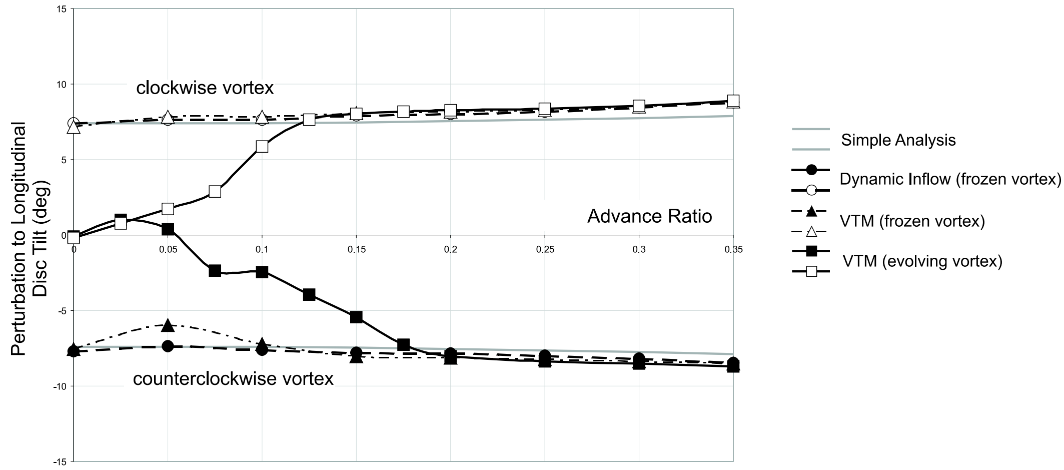

Figure 7: *Simulated response of a rotor to an encounter with a vortex constructed to represent part of the wake of a fixed-wing aircraft. Comparing the predictions of the various approaches shows how elimination of the frozen vortex assumption radically modifies the predicted response of the rotor, particularly at low forward speed.*