

Exciting vibrations: the role of testing in an era of supercomputers and uncertainties

D. J. Ewins 

Received: 7 July 2016 / Accepted: 1 November 2016 / Published online: 22 November 2016
© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract This paper revisits the traditional technology of structural dynamics with particular reference to the applications to critical structures for which structural integrity is a primary requirement. The concept of structural performance is developed with a view to emphasising the positive benefits of advanced structural dynamics capabilities, in particular in the prediction and verification of the safe working life of critical products: i.e. design and demonstration. It focusses on the two primary strands of this capability—analysis for design and test for demonstration—and explains how these tasks are hindered by uncertainties of different types—aleatoric imprecision, and epistemic incompleteness—which are incurred by the inevitable approximations and simplifications that are made in the interest of pragmatic cost-effectiveness. An approach to managing these uncertainties is proposed by exploiting the supporting roles that validation testing can offer the analysis-led design process, and that design models can provide for specification and interpretation of the test-led

verification demonstration. The key to this strategic approach is to ensure that an appropriate balance and integration of analysis and test activities is achieved. The approach is illustrated with specific examples which serve to highlight what are seen as the major challenges ahead in both design and demonstration. These include (1) the need to extend advanced modelling of components to the joints which connect them in every product, (2) the growing importance of including nonlinear characteristics, and the possibility of exploiting them, and (3) the need to ensure that expensive verification tests are adequately defined and executed. Future developments are anticipated to extend test-analysis integration activities into manufacture and the post-delivery service phase of the product's life by combining data collected for monitoring and diagnosis with the design models in order to provide advanced structural health management—the so-called digital twins concept.

Keywords Vibrations · Structural dynamics · Testing · Validation · Verification

This paper is based on the JP den Hartog Award Keynote Lecture with the same title delivered at the ASME IDETC Vibrations Conference in Boston, August 2015.

D. J. Ewins (✉)
Imperial College London, London, UK
e-mail: d.ewins@imperial.ac.uk

D. J. Ewins
University of Bristol, Bristol, UK

1 Introduction to structural dynamics analysis and test

1.1 The challenge

The vibration of structures has been a concern to engineers in many disciplines for a very long time, largely because of the associated damage and

disturbance experienced. A considerable understanding of the underlying physics has been acquired and a technology has been developed for anticipating and controlling the effects of the vibrations generated and experienced by machines, vehicles and structures of all types. The existence of a large body of literature, together with software for analysis and hardware for testing, to implement the technology is taken as read. The issue here, and the objective of this paper, is to consider how best to deploy this technology. The ultimate goal of all these activities is the ability to design our various products so that they experience predictable and acceptable vibration response levels throughout their entire service life. Of particular interest is to explore how to achieve an ideal balance between **analysis**—prediction—and **test**—measurement—of structural dynamic behaviour. In this paper, this is undertaken against the background of an unstated expectation that much testing will—sooner or later, to a greater or lesser extent—be superseded by inexorable advances in computing power. The perspective adopted here is from the testing side—as suggested in the title—and counters such prognosis by demonstrating the integral role played by advanced testing activities. The approach is strategic rather tactical with a definition of ‘strategy’ appended which might usefully be consulted before the main body of the text is read.

1.2 The main issues in structural dynamics are deleterious: failure, malfunction, disturbance

Although there are exceptions, the overwhelming majority of structural dynamics issues are deleterious. The vibrations experienced by our products generally have a negative effect, albeit of varying severity, which can be classified as: 1. Failure, 2. Malfunction or 3. Disturbance. Class 1 outcomes are irreversible, in the sense that once broken, the components in question have to be repaired or replaced. Class 2 and 3 outcomes are generally reversible in that they ‘cease’ if the vibration is eliminated but, even then, some irreversible damage is likely to have occurred. In practice, these outcomes must be considered in design (and test) for a range of in-service operating conditions, often themselves classified as: (1) normal (2) abnormal or (3) extreme conditions. It is rarely the case that vibration problems can be completely eliminated by suitable design, and so there will be a balancing of compromises

of vibration levels and the resulting ‘damage’ that might result by some form of integration over time, often resulting in assigning a finite safe working life for the components or products in question.

1.3 The subject systems: critical structures with failure and malfunction as primary issues

In this paper we shall focus our interest on a subgroup of structures, which we shall classify as ‘critical’. This subgroup refers to structures for which failure or malfunction represents a major threat to product integrity and/or of excessive secondary damage to other equipment or personnel (passengers, operatives, bystanders...). Such cases are found throughout the aerospace and high-performance power industries, as well as defence and transportation, and they are of particular interest because they demand the most advanced technology available. The various examples used to illustrate the main points of the discussion in this paper relate primarily to aero-engines, missiles, and rotorcraft.

1.4 The basic toolkit of analysis and test techniques

We have referred above to analysis (prediction) and test (observation) procedures that have been developed to assist us navigate the dynamics of our structures to ensure their safe and reliable lifetime operation. In fact, there are three primary skills that must be available to the structural dynamicist. These are: **theoretical modelling**, **numerical analysis** and **experimental measurement** and they are usually grouped as ‘analysis’ (or ‘simulation’) which is a combination of theoretical modelling and numerical analysis, and ‘test’ (a group of different types of measurement). The two parts of ‘analysis’ are quite different: the first part requires a thorough understanding of the underlying physics of the elements or components being designed, and the ability to use this to define a set of equations which describe the structure’s behaviour. The second part—numerical analysis—is concerned with providing accurate and efficient algorithms for solving these equations of motion under a wide range of user-specified operating conditions. It goes without saying that the most advanced numerical analysis tools are ineffective if supplied with deficient equations of motion, a situation

which can result from an imperfect or inadequate understanding and representation of the physics. These three basic skills can be set in an application context as illustrated in Fig. 1, showing **simulation**, **validation** and **identification** as techniques that can be applied using a combination of the three basic skills. For the remainder of this paper, we shall refer to **analysis** and **test** as the two fundamentally different approaches of prediction and observation.

2 The challenge of achieving structural performance: design and demonstrate

2.1 Critical structures: dual requirements of functional performance and structural performance

Most structures are designed to meet specific ‘**functional performance**’ targets—fuel consumption; power output; range etc. Although the structural dynamics of the designs may have some direct impact or bearing on the functional performance, this is usually secondary. However, the vibration levels that will be experienced in service can have a very significant effect on the reliability and effective working life of the product. It is convenient to express these effects as constituting the ‘**structural performance**’ of the product, comprising a set of specific metrics that should be met in the same way as those for the functional performance. There are several elements in the structural performance domain and these are shown schematically in Fig. 2. Most of the items here relate to specific mechanisms of deterioration of

the fabric of the component or product being described, and they include fatigue, wear, material degradation as well as other external factors. Structural dynamics plays a particularly significant role here in that it represents the ‘driver’ for most of these mechanisms of degradation of the structure itself. Without vibration of the structure, many of the deterioration mechanisms would not be activated.

So, in this context we can see our role as structural dynamicists in a more positive light. Our goal in managing the vibration characteristics of our structures is to quantify and to extend the working life of the products themselves. Our essential task is “to predict and to verify the life of the product”, and to use our specialist capabilities to ensure that the life is along as possible. This can be translated into ‘**design and demonstrate**’ which, in turn, can be considered as ‘**analysis and test**’.

2.2 Task of the structural dynamicist—design and demonstrate: analyse and test

This seems a straightforward distribution of tasks—**design**, using computer models and algorithms to find an optimum configuration from the functional performance perspective, and then **demonstrate**, by running performance tests on a full-size prototype structure. The structural performance specifications have a less immediate impact than those of the functional performance and can only really be demonstrated by endurance tests, perhaps accelerated, over a period of time which is always ahead of actual service duty.

2.3 Experience shows right-first-time to be rare

Experience with this design-first-then-test approach is that it is rarely successful on the first attempt, and several iterations may be necessary to achieve a satisfactory result. By ‘satisfactory’ we mean that the actual dynamic behaviour of the structure(s) concerned, as determined from the later test programme, closely matches that which is predicted by the model in the design phase. If this result is not achieved, and this is only discovered after the design phase has been completed, then the implications are very serious because the timescales involved in revising the design can be very long and unrealistic. What is required is confirmation that the model which is being assembled for the final design optimisation process (which itself

THE STRUCTURAL DYNAMICIST’S TOOLKIT

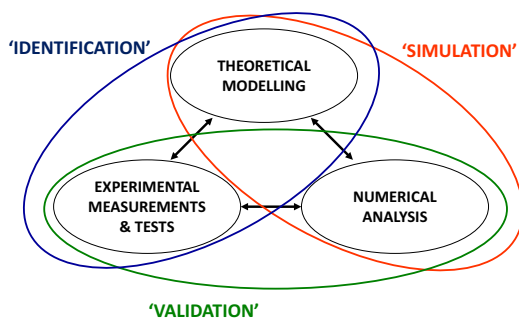
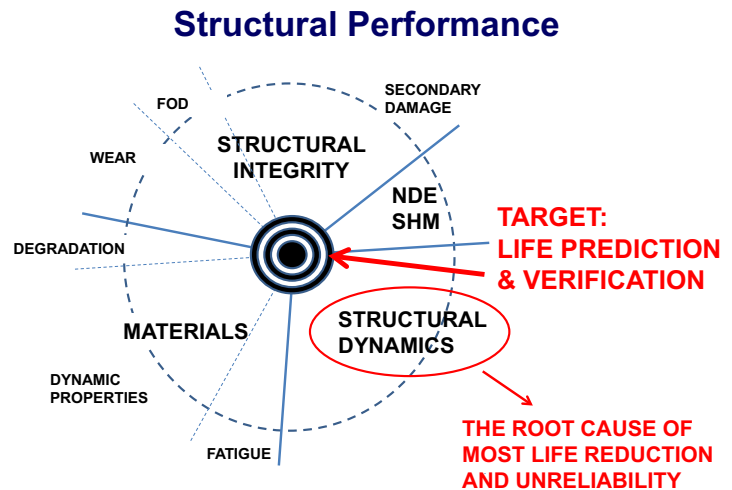


Fig. 1 Basic skills and technologies in structural dynamics

Fig. 2 The role of structural dynamics in structural performance



© D J EWINS 2015

can be very costly) is good enough for the task, and for this to be established before the optimisation is implemented. This calls for a validation procedure to be carried out in the model before it is used in the full design process. There are existing validation procedures available for this task which involve the prediction and measurement of closely-matched dynamic response characteristics. These are then correlated and used to identify errors on the model, which can be corrected using model updating techniques. In practice, it usually proves most efficient to carry out this model validation procedure at several intermediate stages: first, on individual components, to ensure the basic structure's model is adequate; next, on sub-assemblies, when the interfaces that have been introduced to connect components need their models to be checked; and so on, as the complete structure approaches full assembly.

Here we see the value in validation tests being undertaken to provide direct input to the modelling process by ensuring that a valid (good enough) model is created for the design process. Later, when we arrive at the demonstration activity—clearly, a test-driven process—it is often realised that the optimum test to achieve the desired demonstration may not be obvious. The requirement here is for a verification test in which the product either passes or fails. Clearly, this is often a very expensive test and so it is critical to succeed first time. This means, first, that the correct test is specified, and second, that the correct data are measured in order to ensure comprehensive confirmation of the result. To achieve these goals it is almost always essential to carry

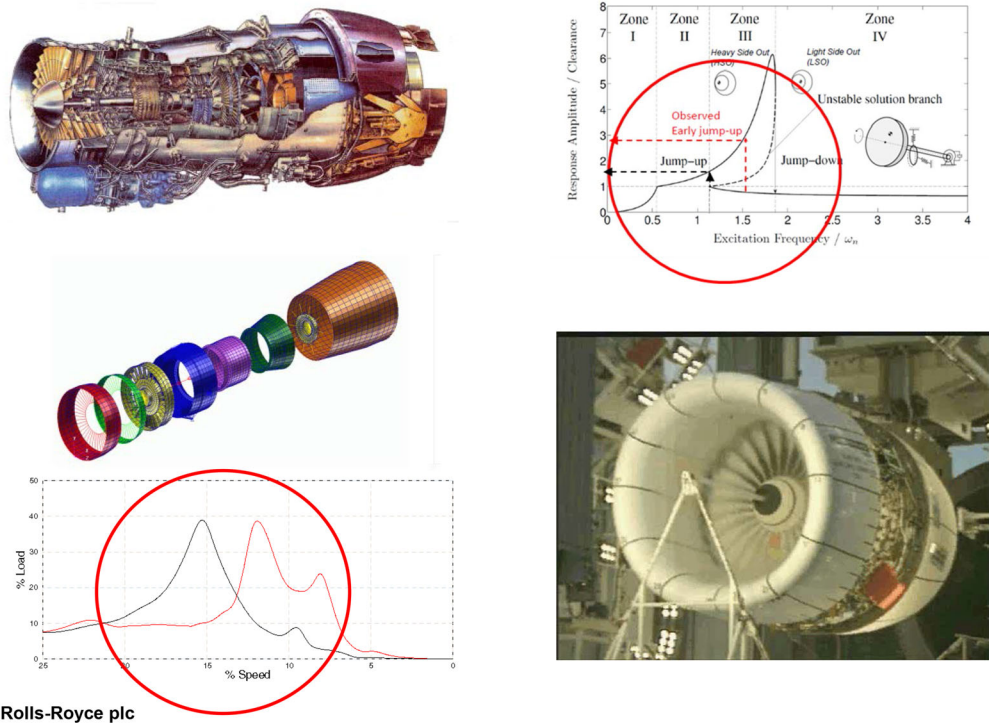
out a detailed numerical simulation of the proposed tests in order to be sure that the one test eventually carried out is definitive. So for the design and the demonstration procedures, we see that both analysis and test must be used in tandem to maximise the effectiveness of the whole process. Figure 3 illustrates typical examples of design and demonstration activities in the development of engineering products, both of which require iterations to achieve the required accuracy of result: design supplemented by validation tests and demonstration refined by analysis.

At this stage, it is appropriate to consider why neither of these processes—design or demonstration—can be expected to deliver the required results first time. An answer can be found in the inevitability of uncertainties being encountered at almost every stage of each process. We make simplifications, approximations, assumptions, selections,... throughout both analysis and test procedures and, as a result, the outputs of our endeavours will not be 100% accurate. Indeed, as engineers, we do not expect or demand 100% accuracy but we do—or should—have a clearly-quantified accuracy which our designs must achieve.

3 Managing uncertainties in analysis and test: validation and verification

3.1 Two types of uncertainty

Uncertainty quantification is currently a widely-discussed topic but it is not always fully appreciated



© 2015 Rolls-Royce plc

Fig. 3 Typical examples of design and demonstration activities

at the outset that there are two quite distinct types of **uncertainty**—**aleatoric** and **epistemic**—and these are quite different in their origins, significance and difficulty to resolve. In simple terms, aleatoric uncertainty is the most familiar and it relates to imprecision or a lack of knowledge of the precise numerical values of individual parameters, whether predicted or measured. Resolution of aleatoric uncertainty in an engineering context is, in effect, a matter of reducing the imprecision to what is decided to be an acceptable level. Typically, in a structural dynamics problem, this would be specification of a vibration response level to within an accuracy of, say 10 or 15%.

In contrast, epistemic uncertainty refers to the inadequacy or incompleteness of a set of parameters that are used to describe behaviour—again, both predicted or measured. This type of uncertainty is more difficult to resolve since it arises because some parameter(s) may be missing from the model or the measured data set. Their imprecision cannot be addressed until they have been identified and included and this can be a much greater challenge than reducing

the inaccuracy resulting from aleatoric uncertainties. One good example of epistemic uncertainty arises if we try to describe the behaviour of a structure which exhibits nonlinear effects by using a model which contains only linear characteristics. Such a model omits the higher coefficients that are necessary to describe—for example—a cubic stiffness effect (i.e. uses $f = kx$ instead of: $f = kx + \beta x^3$). No amount of adjustment of the coefficient k can compensate for the absence of coefficient β .

Similarly, in the measurement domain, there are always issues of imprecision of the acquired data (aleatoric uncertainty) but there is also a significant likelihood of incompleteness in a measured data set, for many perfectly good practical reasons. One common example of this can be seen in the measurement of mode shapes, or operating deflection shapes, when the vibration deflection amplitude is measured at each grid point in just one direction (e.g. along the z axis, which might be normal to the surface of the test structure). At each measurement point on the structure, there will actually be deflections in 3 translation and 3 rotation directions,

but if only one of these 6 is actually measured, then the ‘missing’ data, related to the other 5 directions, is usually recorded as zero by default. As a result, when an animated display of the measured deflection pattern is created, it falsely indicates zero motion in several directions at several points. Incompleteness of data such as this can cause severe problems in subsequent analysis or interpretation of measured data, and these are completely independent of any measurement accuracy concerns which, by definition, can only apply to data which has actually been measured.

Clearly, it is necessary to be aware of the possibility of both types of uncertainty, and to take measures to minimise the consequences of these. For aleatoric uncertainties, it is a matter of seeking more precise estimates of the parameters under investigation—for example, by making repeated measurements or specifying tighter manufacturing tolerances. It is much more difficult to deal with epistemic uncertainties, primarily because it may not be at all obvious which parameters are missing. But, for sure, epistemic uncertainties can rarely be addressed by making repeat estimates of the initial parameter set.

3.2 To manage uncertainties, use tests to validate models for design; use analysis to design verification tests

The central thesis of this paper is to advocate the use of a balanced integration of analysis and test activities in order to arrive at a satisfactorily cost-effective outcome of both design and demonstration of the structural performance of our products. For design, the primary task is to create a mathematical model which is capable of representing the dynamic response characteristics of the product when subjected to different types of loading, these being selected to cover the wide range of operating conditions it will experience during its service life. This task comprises essentially two stages. First, the sub-models of the many individual components that make up the complete structure (or machine or vehicle) need to be defined and checked for suitability. In this phase, conventional correlation and updating processes in model validation testing are widely used to refine preliminary models and, when appropriate, to update previously-estimated material property data to reflect the actual behaviour experienced in hardware products. Generally, though not exclusively, this phase of validation is well served by

commercially-available linear modal testing and analysis approaches. The second phase is more challenging as it involves the assembly of these components into the subsystems and then eventually to the complete assembled product. At each of these assembly stages, attention must be given to the details of the fixture or interfaces between components—‘boundary conditions’—as these are often not modelled as parts of the components. Frequently, when applying the previously-used model validation techniques to an assembly of 2 or more components, the degree of correlation between predicted and measured dynamic characteristics of the assembly is significantly worse than was found on any of the individual components separately. This fact highlights the first major challenge to our control of the structural performance of our products: namely, the need to include appropriate models of the joints (interfaces, connections,...) which are used to connect the components of those products.

One of the first realisations of the importance of including specific models for the joints and interfaces can arise when it is found that the model updating process—invoked to improve the accuracy of the initial model parameters—cannot arrive at a physically-acceptable set of parameters for the initial model. When the outcome of a model updating exercise declares that inertia or elasticity parameters need to be adjusted by 30 or 40% in order to achieve a ‘good’ fit between the model and the test data, this usually indicates that the model is inadequate or incomplete (and does not contain a sufficient number of variables). This can happen when there is a physical flexibility between two components at the point where they are connected, but for which the model supposes a rigid connection. This introduces an epistemic uncertainty which must be resolved before model updating can be applied. The process of ensuring that a model is capable of being updated has recently been re-enforced and termed **model upgrading** [6].

4 Design, and the need for an integration of analysis and test capabilities

4.1 First major challenge: taking account of joints

4.1.1 Current status

It has been recognised for several years that in many critical structures taking proper account of the

influence of the joints on an assembled structure's dynamics constitutes a major challenge to our ambitions to manage structural dynamics by analysis-led design techniques. The task of modelling such joints or interfaces is very difficult, and often infeasible or uneconomic. In fact, there is evidence emerging that many conventional joints have highly unrepeatable dynamic characteristics, which change with time in service and/or with dismantling and reassembly. To offset this, there is also experimental evidence that for a given assembly, on a given day, the dynamic characteristics of a typical engineering joint such as a bolted flange can be quite reasonably described by a relatively low-order effective stiffness and/or damping, such as shown in Fig. 4. However, determination of the values for this stiffness and/or damping is only possible by indirect experimental measurements.

If a simple model of the joint between two components is introduced, with arbitrary values for the associated coefficients, application of model correlation and updating based on test data will often return numerical values which serve to describe the dynamics of the assembly quite adequately. While this might seem to provide a viable route to constructing a model suitable for the design process, the drawback is that the model obtained in this way is not based on a description of the physical system and so is not 'predictable'. It is based entirely on observed behaviour and so is empirical and thus the only means of determining the numerical values for the joint model is by indirect measurement. There is an international research effort [1] seeking to develop methods for constructing such joint models based on physics, but this is proving to be a major challenge. Simply constructing highly detailed models of the interface surfaces, with 1000s of elements, does not seem to be an effective way of resolving this problem. It is clear, however, that joints present a significant epistemic uncertainty issue in structural dynamics, not least because they are widely found to be nonlinear and so the nature of the nonlinearity must be identified first in order to know how to proceed with a model validation exercise.

4.1.2 Prospects: next-generation joints

The prospects for this challenge are unclear. It is noted that there are many joints in a typical structural assembly (at least as many as there are components) and that the great majority of these are not modelled in

conventional design procedures. It is suggested here that this constitutes a major limitation in our current structural dynamic analysis capabilities. It is also suggested that further improvements in our computation capabilities will be compromised if this limitation is not resolved. This is because the uncertainties associated with the dynamics of the joints that are used in engineering structures today do not lie in our computation capabilities but in the ineffectiveness of our modelling capabilities, brought about by an inadequate understanding of the physics of these important and omnipresent features of real engineering structures. In fact, it is believed that the situation is worse than first appears. It is probably the case that inadequate modelling of joints cannot be solved simply by constructing bigger models. We have already seen that empirical data suggests a suitable joint model derived from measurements which is of relatively low order, thereby suggesting that the limitation in current modelling is not one of scale, but of inadequacy in our understanding of the underlying physics. Additional insight into this problem can also be gained from empirical test data which suggests that many joints exhibit a high degree of unrepeatability. This is observed both when re-measuring the same jointed structure at various times after its acquisition and—worse—when re-measuring the same structural assembly after dismantling and reassembly to notionally the same 'condition'. In both cases, the variability cannot be fully explained by uncertainties in dimensions or other basic properties, and so that leads to the conclusion that such joints have features that make them highly susceptible to variation in some parameters that we do not consider, or have chosen to ignore, and which are not included in our modelling efforts. Such features could include: non-flatness of jointed surfaces, micro-wear effects that result from cycles of vibration or changes the micro-level dimensions or conditions of the contacting surfaces. It is well known that jointed structures can have different static properties depending on the sequence in which the connecting bolts are tightened. Such a feature is not accommodated in conventional modelling procedures.

What is proposed here is that joint designs should be reviewed with a view to introducing new criteria such that dynamic properties can be predicted and controlled as well as the conventional static properties that ensure sealing and alignment (the primary function of most joints today). This will probably require a




Fig. 4 Inclusion of models for joints in connected structures

revision of the way joints are configured and designed, so that they can be modelled in respect of their dynamic properties as well as their static characteristics.

Such an approach may have significant benefits to functional performance as it could lead more efficient joints which, in turn, means joints that use less material than is currently the case. This would be highly attractive for applications in the aerospace industries where weight is an extremely expensive commodity. It is likely that many current designs will be conservative, and thus heavy, in order to ensure their robustness to damaging dynamic features of dynamic operating environments. Reducing these uncertainties will almost certainly be accompanied by reductions in weight and thus significant economic benefits in long term service.

4.2 Second major challenge: increasing influence of nonlinear effects

4.2.1 Current developments in dynamic testing of nonlinear engineering structures

The focus on joints in recent years has brought the question of nonlinearity to the forefront of discussion.

Much structural dynamics as applied to industrial applications today is effectively based on the presumption of linearity and assumes that the structures are sufficiently close to linear that the resulting response predictions are within the target range of accuracy. However, as both analysis and test become more accurate, and as designs grow ever less conservative, the incidence of nonlinear effects being clearly significant enough to influence the structural performance has risen noticeably. As a result, it is now necessary to consider their influence as matter of routine, rather than exception, when dealing with critical structures in high-performance application areas. Reflecting this trend, there has been a growth of interest in the engineering aspects of nonlinear structural dynamics recently with special issues of two journals, including some 15–20 papers, being published in 2015 [2] and 2017 [3]. Also, in the UK, a major research programme concerned with the structural dynamics aspects of Engineering Nonlinearity, has been under way since 2012 [4].

The first level of assessment of nonlinearity in a structure's dynamics can conveniently be made in the routine process of modal testing as performed for model validation. One of the standard checks of the

quality or integrity of measured data used for model validation is, in effect, a test of the linearity of the test structure. Most model validation exercises are carried out at relatively low levels of vibration for a number of reasons, including the desire to avoid damage from accidental overtesting and the complications that can arise in data analysis when nonlinearities are present. However, once the primary validation has been completed and a validated **underlying linear model** (ULM) has been developed, it is usually of some interest to explore how the structure's dynamics change as the vibration excitation and response levels are increased towards those which might be expected to be encountered in service—both in normal and abnormal operating conditions: see Fig. 5.

4.2.2 *Prospects: dealing with nonlinearity in engineering structural dynamics*

These checks are usually made today to demonstrate that any nonlinearity in the test structure is not significant enough to compromise the linear-based modelling that has been undertaken. However, we must now be looking ahead to the next level of modelling and validation, to situations where nonlinearities cannot be ignored. One recent publication has set out a proposed procedure for managing this situation [5] and another provides a first example application [6]. The proposed approach builds on current (linear) methodology by seeking to use test data to enhance, or upgrade, the underlying linear model by identifying which parts of the structure are exhibiting non-linear characteristics. It is suggested that in many practical applications, the sources of nonlinear behaviour are likely to be focussed in a relatively small number of localised regions—such as (some of) the joints. The essential approach in such cases is to seek to upgrade the preliminary model by identifying those regions/elements which are contributing the most significant non-linear effects. Once these features have been located and characterised so that the additional spatial parameters to describe them have been identified, they can then be updated by an extension of the conventional (linear) updating methodology (although this technology is still under development). Figure 6 shows a flow diagram of the proposed methodology.

It should be noted at this stage that there will inevitably be a wide range of the degree of

nonlinearity found in engineering structures. At one end of this range there will be many structures which are largely linear but which have discrete localised nonlinearities, such as joints. These nonlinearities are not necessarily 'weak' but they are confined to a very small percentage of the elements in the overall model and the task here is to be able to identify these few elements and to ensure they are appropriately configured to describe the (sub)components behaviour. Often these effects are such that they are only really significant at the higher levels of excitation and response that are encountered in the more extreme operating conditions. Normal operating conditions may well be adequately represented by the underlying linear model. At the other end of the range there will be structures which contain components that have a primary nonlinearity, widely distributed throughout the structure and effective across the operating range of the product. These cases will demand a more extensive and perhaps individual treatment than the earlier type, which can possibly be accommodated by an incremental extension of the underlying linear model, and will almost certainly require a more customised approach than we are considering here.

In many practical engineering applications, we are seeking an incremental extension to our model which makes it capable of delivering predictions of the structure's response to typical in-service loading to our prescribed level of accuracy. In this respect, our primary objective is to identify and quantify those elements in the structure which have a non-linear, rather than linear, characteristic. With this information we can construct a design model which can then be used for prediction of the structure's behaviour under many different excitation conditions and it is assumed that there (will) exist numerical analysis algorithms that can be used to predict the forced response of models which include such nonlinear elements. This approach differs somewhat from other activities in nonlinear dynamics where advanced techniques of non-linear normal modes, backbone curves and other descriptors of inherent nonlinear dynamics properties are widely reported. The difference is largely one of emphasis as here we are primarily interested in constructing a valid spatial model and less so in the complex response characteristics that will be encountered when that model is used for design purposes. Here, once again, we are seeking to integrate our analysis and our test capabilities in the most effective

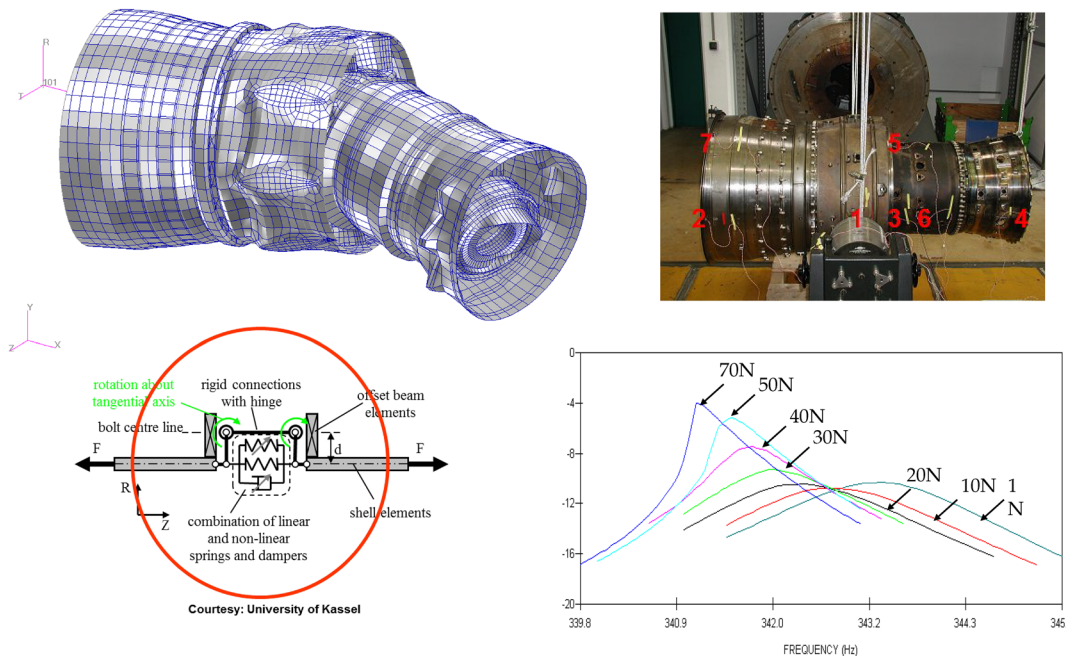


Fig. 5 Evidence of nonlinear behaviour in jointed structure

way in order to achieve our goal of being able to construct models to design for structural performance of the most advanced structures. We may well seek to do this without making recourse to some of the more complex nonlinear characteristics, but by relying on traditional response function measurements under suitably controlled conditions to validate our models.

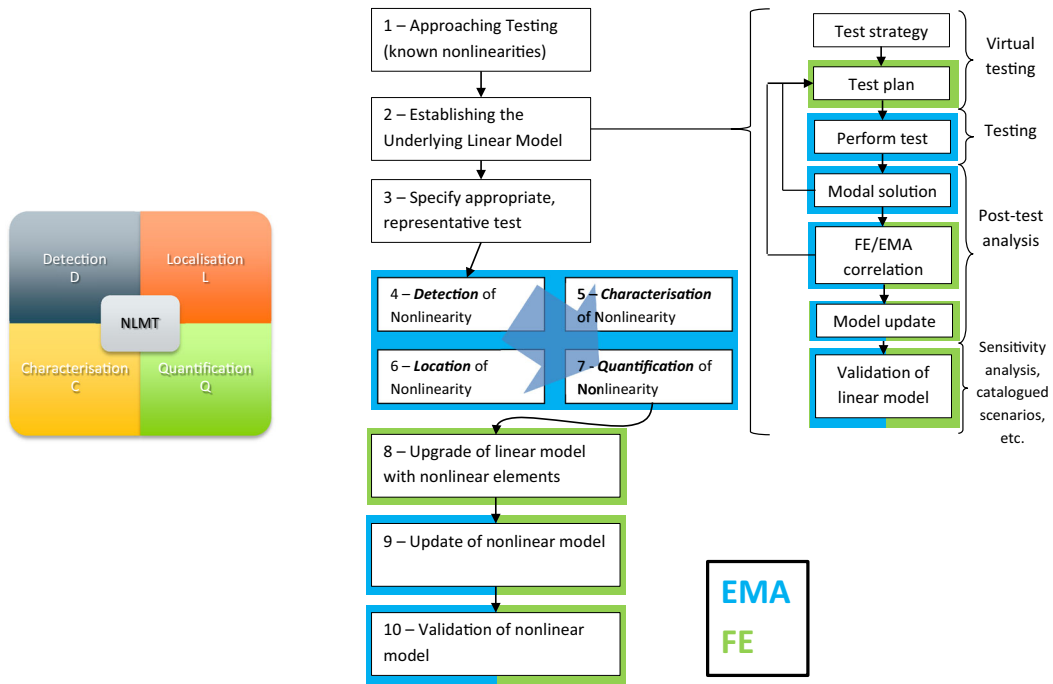
5 Demonstration, and the need for an integration of test and analysis

5.1 The role of verification tests

Once the model has been validated, the product design can be optimised to achieve both functional and structural performance targets. There then follows a practical demonstration of the overall performance of the product, in order for the customer to be shown that the finished product really does deliver what has been agreed in the contractual specifications. In many cases, there may also be a requirement to demonstrate the structural performance, especially in respect of safety issues, not only to the customer but also to relevant certifying authorities. These demonstrations can only be performed by physical tests, referred to here as

a verification tests, the outcome of which is usually pass or fail.

At this stage, the stakes are high, because not only must the product pass the test, but the test itself must be fully representative of the operating conditions under which the product will be in service. In many cases, the responsibility for defining what the verification test(s) should be will rest with the manufacturers themselves, not least because they will have a better perspective than anyone of the boundaries of the performance envelopes of their product and these need be demonstrated in a physical test. The requirement is not just that ‘the product has not failed after 10,000 h in service’ but rather that ‘it is anticipated that it will fail (which may mean ‘fail to deliver the full performance level’, rather than ‘break’) between 12,000 and 15,000 h’, and to demonstrate in tests that this is a reliable prediction. This means that the verification test must be carefully designed so as to demonstrate the credibility of the design itself. This, in turn, requires an extensive input to the testing programme from the analysis capability that has been used to design the product in the first place. This is where, in order to minimise the uncertainties that might compromise the authority of the verification test, another close integration of test and analysis



© D J EWINS 2015

Fig. 6 Proposed schedule for model validation of nonlinear structures

capabilities must be enacted. This can be illustrated by two different industrial applications, below.

5.2 Qualification tests of stores subjected to long-term dynamic environments before deployment

There are many examples in the aerospace and defence industries where a product has to endure significant and often extreme vibration environments for long periods of time before it is deployed on its own mission—at which time its structural integrity is of paramount concern. Clearly, there are many defence applications where this is the case, but so also are the various space projects where satellites have to endure extreme dynamic environments just to reach their place of deployment.

It is common practice to determine or to assess or prescribe the dynamic environment that such stores will experience in service or in transit, and then to demonstrate their ability to survive these environments, undamaged, by conducting suitable endurance or qualification tests. In simple terms, the service environment can often be measured on existing host vehicles and then a specification made for a test to

reproduce these environments in the laboratory where the behaviour of the product can be closely monitored. It is also possible in this format to carry out accelerated tests so that a safe in-service life can be demonstrated ahead of actual service experience.

For many years, it has been normal practice to define the in-service environment based on a set of measured spectra of vibration levels in different directions at different locations on the host vehicle, close to the points of attachments of the product. Then, these vibration levels are reproduced in the laboratory using shakers to replace the physical excitation that occurs in service (which might be aerodynamic in origin). The product is then subjected to long-duration endurance tests to demonstrate its robustness, and thus survivability in that environment. In many respects, this procedure involves a number of considerable assumptions and simplifications that can have a bearing on the validity of the tests. These assumptions are often not detailed but they involve the interdependencies of the different vibration levels which define ‘the service environment’. Simply, the originating excitation forces which are applied to the host vehicle are not taken fully into account in the simulated endurance test, nor is the different structural interface

which the product ‘sees’ (1) in service and (2) under test. Recent studies of this widespread practice have led to the formulation of a more representative approach to this class of endurance test, and they have done so by a carefully managed integration of test and analysis, as illustrated below (and documented in [7–9]).

In summary, a new methodology has been developed and demonstrated on a small-scale model of a missile, shown in Fig. 7a. This was first installed in a wind tunnel and subjected to aerodynamic excitation to simulate the operating environment of a missile carried under the wing of a host aircraft. The vibration levels at the attachment points, along with other critical locations on the missile body (control locations), were recorded for use as a definition of the environment to be reproduced in the subsequent endurance test, and examples are shown in Fig. 7c. The model was then installed on 2 shakers as shown in Fig. 7b and these were controlled so as to replicate the spectra of displacements at the 2 control points. The results of this test are shown in Fig. 7c where it can be seen that the two control spectra are very closely replicated under the test operation. Next, measurements were compared at two other points on the missile, which were not included in the test specification, and here it can be seen that the actual test vibration levels are very different to those which had been recorded in the in-service wind tunnel measurements Fig. 7d. The discrepancies amount to both over-testing and under-testing by orders of magnitude, and are not unrepresentative of experience in such testing in industry. It is realised that although the vibration levels at the control points are well replicated from the in-service data, this is because the controllers driving the shakers are set to achieve just this condition. However, the structural interactions between the missile and its support structure in service, and that when attached to the shakers, are not the same, and so there is no reason to expect vibration levels away from the control locations to be the same in the wind tunnel measurements and the on-shaker tests.

What needs to be done is to simulate not only the vibration response levels, but the complete structural configuration, and interactions between missile and host vehicle. In turn, this requires treating the missile as a 3D structure, and not just to consider single-axis vibration, as in this first qualification test. This can all

be done by incorporating a mathematical model to describe the structural interfaces of the actual structures into the test setup, and by exciting the structure in x , y and z directions simultaneously, as shown in Fig. 8a. When tested this way, the verification test succeeded in reproducing both the control response levels (as expected, and as found in the previous test) but it also enabled the responses at other points on the missile to be controlled to the specification (this was not possible in previous setups)—see Fig. 8b.

By this means, using a combination of direct testing together with an analytical compensation for the missing parts of the complete structural assembly, a much more effective and realistic verification test of the product was made possible. Curiously, the power requirements for the latter enhanced test were found to be considerably lower than those necessary for the first, traditional, type of test. It is probable that the avoidance of massive overtesting of the test structure observed in the original test is linked to the need for much lower power levels in the test cell.

5.3 Verification of an extreme event situation

In many aerospace applications, there will be a number of extreme event situations for which verification of the structure’s capacity to survive is mandatory, at least to the extent of demonstrating that it does not compromise the integrity of the host vehicle, even if the structure is itself no longer functional. A classical example of this type of verification of structural integrity is the fan-blade-off (FBO) test which is a certification requirement for aero engines to ensure that the aircraft is not compromised following an unanticipated fan blade detachment from a running engine. The essential FBO test is to run the engine up to full speed, and then detach a fan blade using an explosive bolt in the root. The engine casing is required to contain the fragments of blades and other debris thus created, so that they do not puncture the pressurised fuselage of the aircraft. But it must also then survive the ensuing run-down period during which the fan rotor decelerates from the full speed at initiation to a much lower steady speed rotation, referred to as the ‘windmilling’. This rotation involves a significant out of balance disturbance being applied to the engine body, and thus to the aircraft, as a result of the missing blade or blades. While the first phase of containment is more a strength of materials issue, the

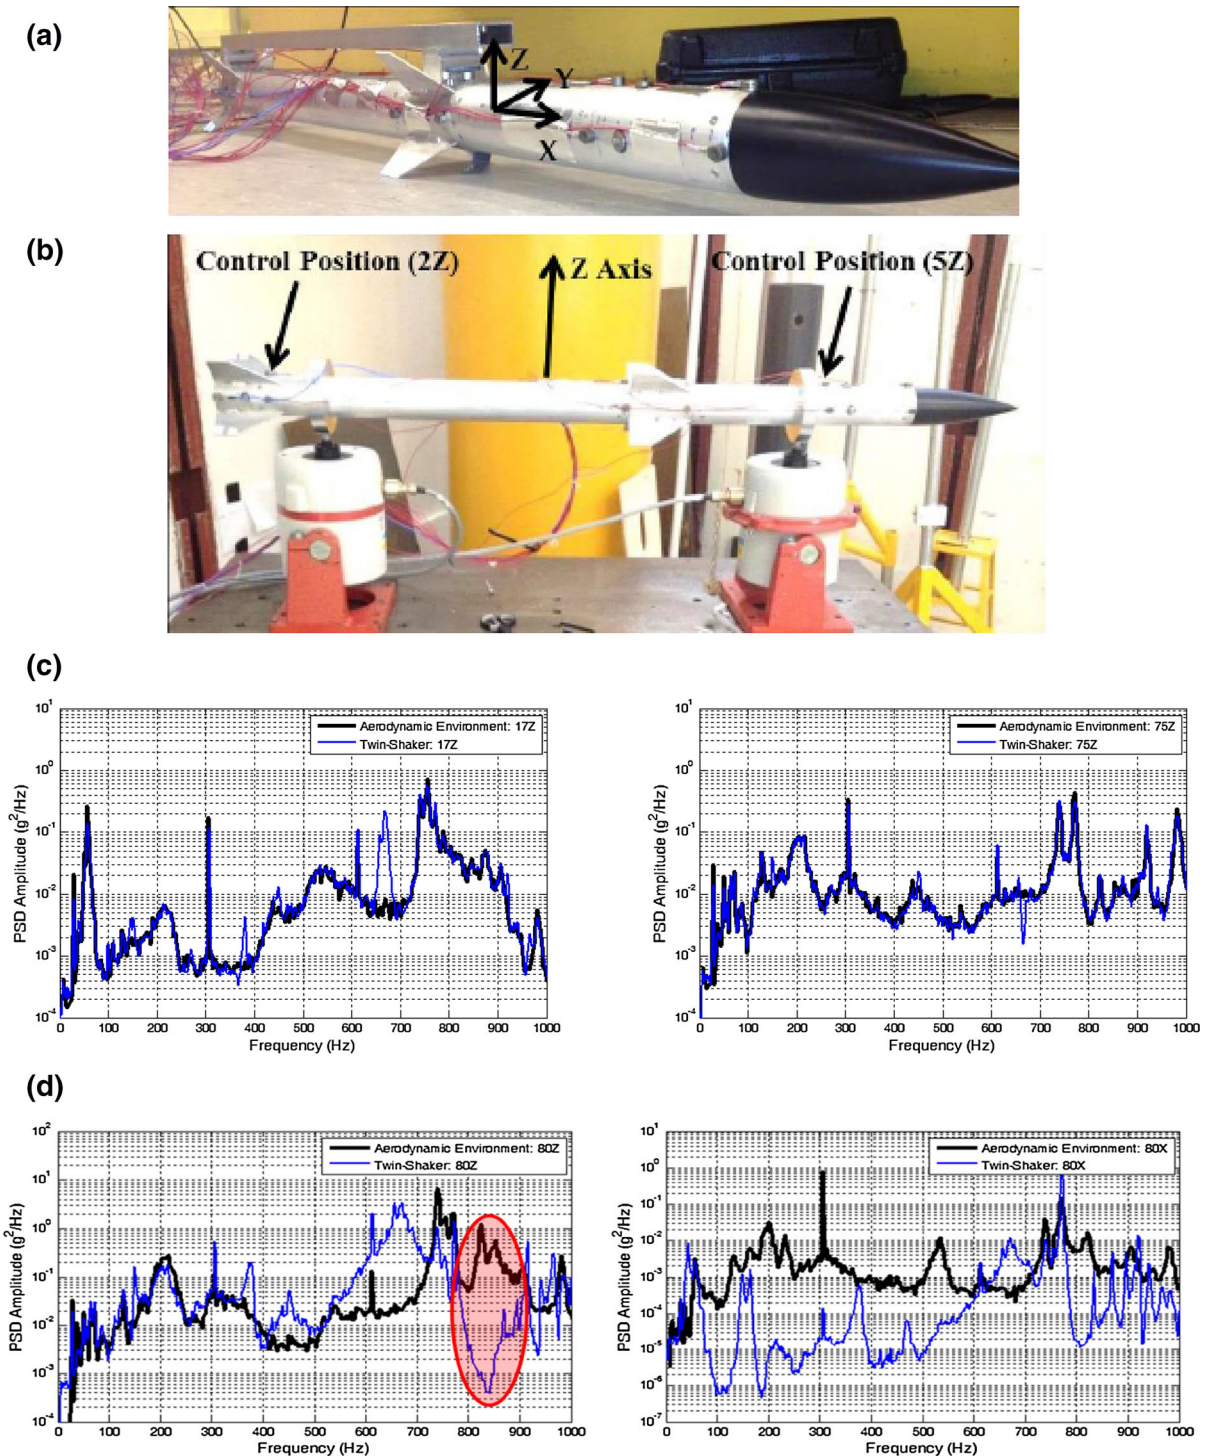
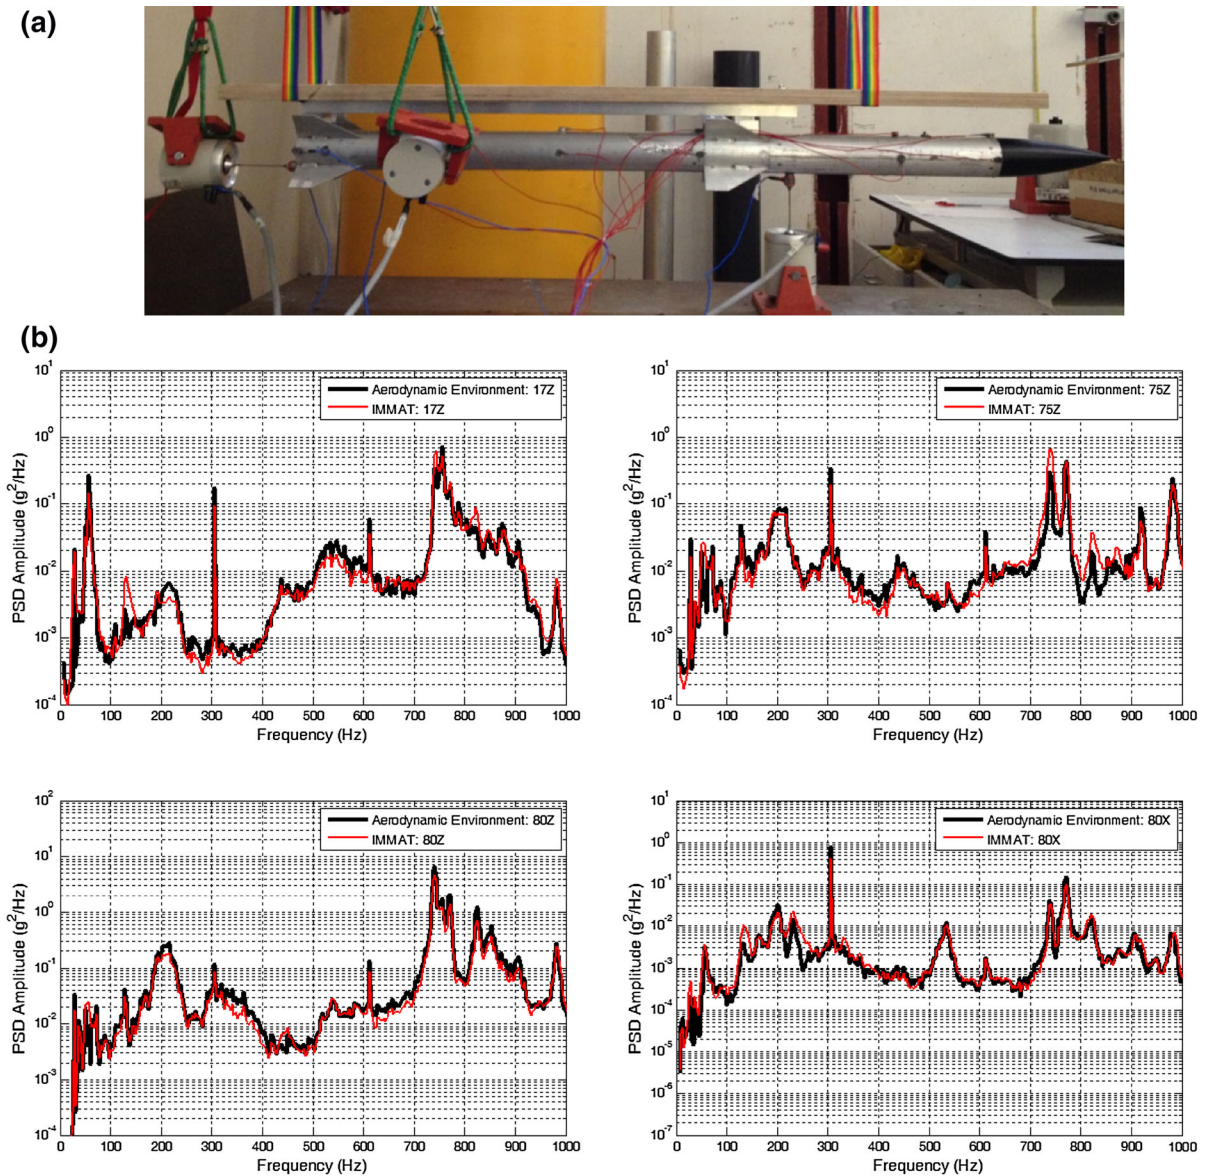


Fig. 7 **a** Small-scale missile model used for verification test. **b** Traditional verification test using twin-shaker mono-axis test configuration. **c** Matching control displacement spectra in

traditional verification test. **d** Failure to match additional response spectra in traditional verification test



© British Crown Owned Copyright 2015/AWE

Fig. 8 **a** Qualification test using advanced test/analysis methodology. **b** Qualification testing using advanced test/analysis methodology

second phase—which will probably continue for the duration of the flight back to a landing site—involves very high level vibration of the whole aircraft, at frequencies which, even though they may not be dangerous for the airframe integrity, are almost certainly extremely discomforting for the on-board crew and passengers. Clearly it is necessary to demonstrate that the product will survive such an

extreme event, and this can only be done convincingly by a physical test.

The FBO test is a good example of a test which must be ‘right-first-time’, by which is meant that the actual test performed must not only be successfully passed, but it must also be accepted as having demonstrated the worst case version of the extreme event that might realistically be encountered in

service. In practice, meeting this requirement of demonstrating the worst case can only be done by using a mathematical model (the one used for final design) to establish exactly what conditions would constitute the worst case. Trying to identify this purely by testing would be unrealistic in both time and cost, and so here again an appropriate integration of test and analysis is the only viable way of reducing the uncertainties surrounding what would be a worst case to minimal proportions. A specific example is worth noting here, and this is summarised from a recent ASME paper, [10]. That reports refers to a specific FBO test case in which the maximum vibration response level experienced in a post-FBO rundown test was observed to be somewhat higher than anticipated—a phenomenon illustrated by a simplified model in Fig. 9. In that case, the maximum response level was expected to occur in a jump-up phenomenon as the rotor decelerated through a major resonance of the rotor system. The jump-up occurred in practice at a higher speed than originally anticipated, with the result that the response level was considerably higher than had been expected. Detailed study of the test setup using a simplified but representative model revealed that the jump-up might occur at different stages of the rundown depending on the deceleration rate, as opposed to just the actual rotation speed as had been assumed. Here, again, integration of test and analysis succeeded in removing residual uncertainties from the verification process such that the test itself was deemed valid, and the verification successfully passed.

6 Summing up and future prospects

In this paper we have set out a philosophy and strategy for addressing the many structural dynamics issues that are encountered in designing and using machines, vehicles and structures which are subjected to dynamic loading in service. We have focused on technologies most applicable to the critical structures which are the most aggressively affected. Running through the discussion we see two parallel and complementary primary capabilities—**analysis** and **test**—and an overriding interest in attaining the right balance between these two.

If we look at the complete life cycle for a typical critical structure (Fig. 10), we see that there are four

distinct stages from concept to decommission and that at each stage there are both analysis-led and test-led procedures involved in achieving and maintaining the structural performance.

In this paper, we have focussed our attention here to the first two core activities of designing and delivering the product, although it can be seen that there are two other areas—**manufacture** and **service**—that are also highly relevant, and should be similarly addressed in another report.

As we review the various stages, we repeatedly encounter two complementary aspects, echoing the duality of analysis and test. The first example is the need to consider both **functional performance** and **structural performance**—the first representing the immediate capabilities (speed, thrust, power,...) and the second representing the longer-term issues based on the duration for which the functional performance can be sustained—both as a safe working life but also maintaining the structural features that have a direct influence on operating performance. In this context, the usual negative view of the significance of dynamics and vibration phenomena can be turned into a positive one of determining the useful life and reliability and, as such, quantifying their economic value, alongside that of the primary functional performance metrics.

The next major distinction is between the basic technology capabilities of **design** and **demonstrate**: being able to design something to have competitive functionality, on the one hand, and being able to demonstrate this performance convincingly to critical audience of customers and authorities on the other.

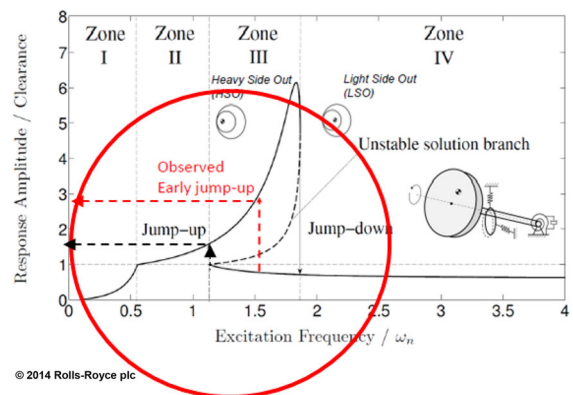
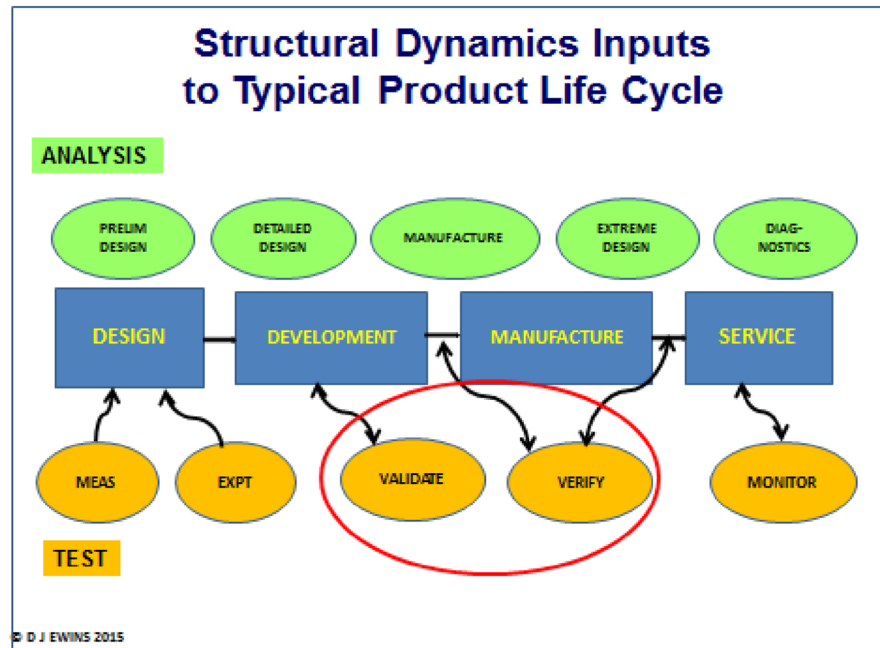


Fig. 9 Details of rundown jump-up at different speeds of rotation

Fig. 10 Analysis and test procedures through a product life cycle



These both draw on the two fundamental skill sets of **test** and **analysis**, or what we can measure and what we can predict before any physical hardware is available. It is found the two capabilities of design and demonstration are both most effectively carried out using a combination of both test and analysis, design being analysis-led but supported by test, and demonstration being test-led but supported by analysis. This integration and balance of test and analysis is the primary focus of the paper.

The challenges to successful design and demonstration are represented by **uncertainties**—**aleatoric**, of imprecision and variability inevitable in an imperfect world, but largely manageable—and **epistemic**, which are more profound in that they represent an ignorance or inadequacy which is displayed here by our omission of certain features and our unawareness of their importance. Identifying and then correcting those omissions is one of the engineer's greatest skills. In one of the inner procedures used in this subject we find **model updating** and **model upgrading** addressing, respectively the aleatoric and epistemic uncertainties in our design models.

Next, the tactics proposed as the means of successfully navigating through the uncertainties which confront us in the form of **validation** and **verification**: systematic methods to identify the uncertainties and to

adjust our models and tests to take them into proper account. These two processes both involve an appropriate integration of analysis skills and techniques and corresponding experimental ones. It is the integrated balance of these two approaches that provides us with a methodology for managing structural dynamics.

Most of this paper has been concerned with the first 2 of the 4 major activities in Fig. 10—**design** and **development**. After this point, the project moves into **manufacture** and **service** and these two phases will also have some significant contributions to the overall success of the product. It is inevitable that some new uncertainties will be introduced in the manufacture stage, the most obvious of which will be the existence of some scatter in the dimensions and other properties in the manufacture of a batch of nominally-identical units. These variations will, in turn, lead to scatter in the structural performance parameters which have been carefully evaluated and demonstrated in the prototype verification testing phase. In that phase, we shall have demonstrated that our simulation tools are capable of predicting the required structural performance characteristics to within a target accuracy—say, $X\%$. If we now introduce a new element of scatter in the various model parameters due to limitations in manufacture, this will result in an additional uncertainty on the accuracy of the predictions. It may be

convenient to talk of these two effects: (1) the quality or accuracy of the model (say, X%) and (2) the confidence that we can have in those predictions when taking account of the scatter of nominal parameters (a confidence of, say, Y%). The quality of the model determines X, while the quality of the manufacture plus the sensitivity or robustness of the inherent design determine the (<100%) confidence in the simulation output when applied to a fleet of nominally-identical products. It can be noted that the robustness of the design is itself a characteristic of the design, and can be quantified from the models used for design.

What has not been included in this review is the important post-delivery in-service activities of **monitoring** and **diagnostics**. Throughout the service life of such structures—typically 30 years or more—it is essential to maintain the structural performance in order for the product to continue to deliver the functional performance. While it is routine today to monitor key parameters throughout the working life of such a structure, these data are most often used to archive past experience and to analyse this from a statistical perspective. It is less common that such data are used together with the design model to diagnose the source of individual discrepancies, or faults, and to specify a remedial action. It is considered that with the emergence of advanced design models capable of delivering the structural performance as described above, there is now a major opportunity to develop a next generation of monitoring with diagnosis and prognosis for a powerful structural health management technology.

Recently, the concept of ‘digital twins’ has been introduced, referring in effect to the construction of two complementary models of the product in question, one from **design** and the other from **service**. The first, which is the primary subject of this paper, is an analysis of the behaviour of the product based on our understanding of the underlying physics. The second is based on a potentially vast collection of measured data from its actual behaviour in service and so is an empirical model. As always, one twin is the senior and here it must be the physics-based model, but this can be refined and perfected by constant reference to the empirical model.

Acknowledgements The author wishes to acknowledge with gratitude the major contributions to this paper which have been provided by a generation of colleagues, students and sponsors at

Imperial College London, the University of Bristol and the University of Oxford. The most recent of these are included in the specific references but there are many others who prepared the ground for what is presented here. The author is particularly grateful to Rolls-Royce and AWE for their longstanding sponsorship of the research behind this article and also to Sandia National Labs and others who have had a significant input to the more recent applications.

Compliance with ethical standards

Conflict of interest The author declares that he has no conflict of interest.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix

A note about strategies

The overall context of the material in of this paper is strategic. It addresses the subject at a high level, ‘descending’ to specific examples to illustrate various issues that are highlighted. Much of what we do as practitioners is tactical—developing and applying the tools of our trade in pursuit of building and operating excellent products. This paper seeks to step back from these details to review the subject from a more distant perspective—less of ‘how’? to solve specific problems and more of ‘which’ problems to address? and ‘why’?

It is worth noting that a strategy has 4 elements, although the word is often used for only one of these. The four stages are:

1. *Objective* What exactly do we want to achieve? It is necessary to define a clear and usually quantitative target for our task
2. *Current position* How far are we from achieving our objective? Define the current position with respect to the objective.
3. *Options* What are the possible ingredients of a solution to the task? What methods might be tried or need to be developed? Where can we find additional ideas?
4. *Plan* With a comprehensive set of prospects (from 3), draw up a plan of action, with timings and costing, to tackle the task. The plan can be

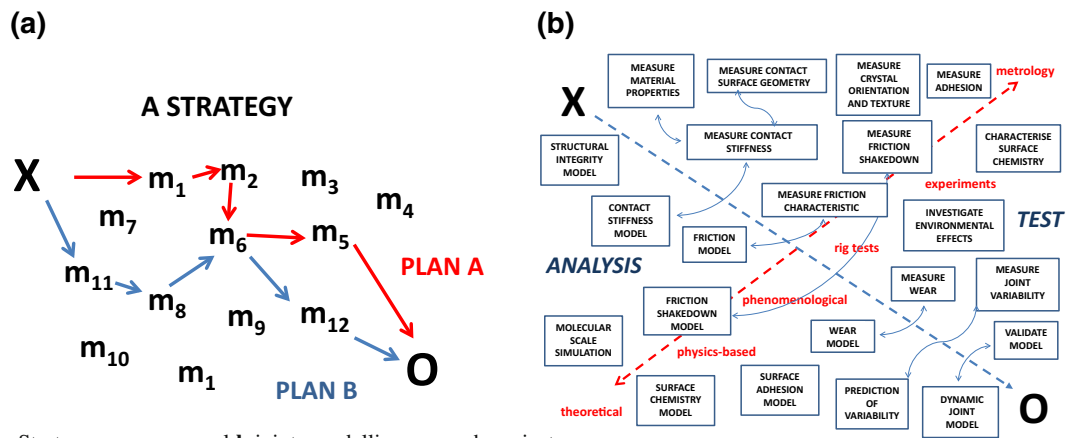


Fig. 11 Strategy map **a** general **b** joints modelling research project

changed in the light of experience in working through it.

Often, the term strategy is used for what is, in reality, just the Plan (step 4), frequently without a proper definition of the Objective (1), or a realistic assessment of the starting point (2) or a comprehensive set of options (3).

This interpretation of Strategy can be illustrated by the diagram in Fig. 11a together with an appropriate example for one of the key topics of this paper: the Joints Modelling Challenge in Fig. 11b. **O** is the Objective; **X** the Current Position; m_1, m_2, \dots possible methods; **A, B** alternative plans

References

1. ASME Joints Modelling Research Committee. <http://asmercmjs.weebly.com/>; https://community.asme.org/research_committee_mechanics_jointed_structures/default.aspx
2. Butlin T, Woodhouse J, Champneys A (eds) (2015) A field guide to nonlinearity in structural dynamics. Philos Trans R Soc A 373(Special Issue)
3. Kerschen G (ed) (in press) Mechanical Systems & Signal Processing. Special Issue: Recent advances in nonlinear system identification. vol 84 Part B
4. EPSRC Programme (2012) Grant “Engineering Non Linearity” Grant No. EP/K003836/1
5. Ewins DJ, Weekes B, delli Carri A (2015) Modal testing for model validation of structures with discrete nonlinearities. Philos Trans R Soc A 373:20140410. doi:10.1098/rsta.2014.0410
6. delli Carri A, Weekes B, di Maio D, Ewins DJ (in press) Extending modal testing technology for model validation of engineering structures with sparse nonlinearities: a first case study. Mech Syst Signal Process. Vol 84 Part B, pp 97–115: doi:10.1016/j.ymssp.2016.04.012
7. Daborn PM, Roberts C, Ewins D, Ind P (2014) Next Generation Random Vibration Tests. In: Topics in modal analysis II, vol 8: proceedings of the 32nd IMAC, 2014, conference proceedings of the society for experimental mechanics series. Doi:10.1007/978-3-319-04774-4_37
8. Daborn PM, Ind PR, Ewins DJ (2014) Enhanced ground-based vibration testing for aerodynamic environments. Mech Syst Signal Process 49:165–180
9. Roberts C, Ewins D (2016) Multi-axis vibration testing of an aerodynamically excited structures. J Vib Control. doi:10.1177/1077546316642064
10. Zilli A, Williams R, Ewins D (2015) Nonlinear dynamics of a simplified model of an overhung rotor subjected to intermittent annular rubs. ASME J Eng Gas Turbines Power 137(6):065001