Nonlinear Dynamics (2006) 43: 127–136

Nonlinear Dynamics of Mechanical Systems Under Complicated Forcing

ALLAN MCROBIE
Engineering Department, Cambridge University, Trumpington St, Cambridge CB2 1PZ, U.K.;
(e-mail: fam@eng.cam.ac.uk; fax: +44-1223-332662)

(Received: 7 April 2004; accepted: 9 February 2005)

Abstract. This paper considers the problem of simple linear mechanical oscillators forced by complicated forcing functions, such as those induced by fluids or humans. Such problems arise in civil engineering in the design of bridges, grandstands and towers. The paper questions whether low dimensional models have much to contribute to the understanding of such problems, given the complicated possibilities of high-dimensional forcing.

Key words: human-structure interaction, fluid-structure interaction

1. Introduction

It is a great pleasure to contribute to this volume on Modern Trends in honour of the work of Michael Thompson, having worked in his Centre at University College London from 1989 to 1994. At the end of this period, I married his daughter, Helen, and, feeling it appropriate to insert a little geographical and professional distance, left. Professionally, this separation can be loosely summarised as my moving from the left- to the right-hand side of the equations of motion. In the standard descriptions of mechanical systems, the left-hand side is the demesne of the structure and the right-hand side is where the external forces are written.

In the early 1990s, the work at the Centre for Nonlinear Dynamics and its Applications at UCL was largely focused on the effects of nonlinearities in the structural behaviour. My original interest in this arose during the design of a large power station chimney on the cyclone-prone Queensland coast. The design was dominated by considerations of dynamics, from both wind and possible earthquake excitations. For both excitations, the existing design practices assume that the material behaviour is largely linear, and yet it is well known that the moment–curvature relationship for reinforced concrete sections can be highly nonlinear, particularly as the structure approaches the ultimate limit state. Shortly after completing this design, I arrived at UCL intent upon addressing the question of what possibilities can arise in forced mechanical oscillations if the structural properties are nonlinear. Michael was at that time working on the problem of ship capsize which he had by then idealised and described in his paper ‘Chaotic behaviour triggering the escape from a potential well’ [1]. The archetype under study was ‘the escape equation’

$$\ddot{x} + \beta\dot{x} + x - x^2 = A\sin(\omega t) \tag{1}$$

Conversations with Michael soon convinced me that the chimney problem bore strong similarities, and could, in the first instance, be seen as an escape problem, although a certain suspension of disbelief seemed to be required. Firstly, the chimney was not a single degree-of-freedom system, and the wind forces, particularly in the along-wind buffeting case, were far from the perfect eternal sinusoids assumed on the right-hand side of the escape equation.

Nevertheless, the two problems, chimney collapse and ship capsize, did share some central features. One was the realisation that ultimately a designer probably cared little if the chimney/ship was performing a very interesting period-two subharmonic orbit. All the designer really wanted to know was whether or not the thing would collapse or capsize. It was here that one of Michael's insights shed much-needed light. Much of the work in the newly-forming Centre at the time was looking at the capsize problem, idealised as the escape equation. The detailed structure of sub- and super-harmonic periodic orbits, all entwined in extraordinarily beautiful, convoluted homoclinic tangles, was being uncovered daily by the cell-to-cell mapping and invariant manifold tracking algorithms that ran all night on IBM PS2's. Michael's insight was that for all their filigree beauty, these structures were little more than the skeleton through which the transient orbits flowed. Capsize is inherently a transient problem. The underlying message was that if you can understand the periodic orbit and invariant manifold structure, you can better understand the transient behaviour which is the real object of interest. Since transients may only last for a few cycles of forcing before escaping from the well, the need for perfect sinusoids that exist from here to forever is removed.

By carefully building a picture of the transient possibilities under the action of a handful of forcing cycles, one could begin to 'ring-fence' the conditions under which collapse/capsize was likely. Parameter regimes were identified [2] in which large proportions of the transient orbits subsequently escaped. The phenomenon was known as 'basin erosion', and it occurred beyond a parameter locus that Michael called 'the Dover Cliff'. Subsequent students and co-workers then incorporated layers of added reality to this underlying picture. The interaction of heave was added to the roll motions. Real G–Z curves for actual ships were used. Collaborators performed experiments in wave tanks and this volume contains papers from some of these.

For the chimney dynamics problem, however, it was not clear that the greatly increased understanding of the possibilities of single-degree-of-freedom nonlinear oscillators was having quite as much useful impact. Some of the differences between the ship and the chimney problem could not be easily assimilated. For example, the collapse of the chimney is typically abrupt. There is a discontinuity at the point of failure. This removes the smoothness of the potential barrier over which transient orbits are escaping. However, all the invariant manifolds of the unstable hilltop saddle which so neatly form the skeleton of the transient motions depend upon that smoothness for their existence. An even bigger difficulty is presented when the excitations are considered in more detail. Another major difficulty is that for many wind-sensitive structures the fluid forces depend upon the structural motion. This is further complicated by the fact that the forces can depend not only on the *present* state of structural motion, but on its whole *past* history: vortices may have been created in the past which are still around to influence the present motions.

On leaving UCL in 1994, my research focus changed to look at the forcing mechanisms. To a fluid dynamicist, a fluid–structure interaction problem largely involves solution of the Navier–Stokes equations over a domain with a moving boundary. To a structural engineer, the same problem is usually formulated as a mechanical oscillator excited by an external force f. It was evident that this letter f was an overly-brief designation of some highly complicated physical processes.

Much of the nonlinear dynamics research that was conducted at UCL in the early 1990s can be characterised as computer exploration of the low-dimensional phase spaces of simple nonlinear mechanical oscillators. Although significant progress was made in that arena, it was always clear that it would prove difficult to extend those techniques to problems of substantially higher dimension. Hopes were expressed that perhaps some high dimensional problems may have their important dynamics confined to some low dimensional submanifold of the full phase space. The literature contained a variety of ideas, ranging from nonlinear Galerkin methods and inertial manifold approximation [3, 4], Karhunen–Loeve decomposition [5], and local adaptive Galerkin bases [6].

One research direction at UCL did look at the application of inertial manifold techniques, techniques that had originally been proposed in connection with fluid turbulence, to the hopefully simpler case of shell and plate vibrations [7, 8]. The essential idea was that the response of the structure could be split into low and high frequency components, with an assumption that the high frequency response was slaved to the low frequency behaviour. The low frequency response could thus be tracked with some numerical efficiency and the high frequency response could then be superimposed at a later stage when required. Algorithms were constructed which achieved superior numerical speeds than tracking the full dynamics. However, it was not clear that the gains in processing time in any way compensated for the substantially increased programming time required to set up the algorithm. As is often the case, increases in computer processor speed can largely obviate the need for complicated problem-specific algorithms which give only marginal increases in speed yet require substantial preparatory formulation. In purely structural problems, computation time is rarely large compared to the time required for model construction. The main advantage of structural finite element models is the ease with which models of complicated domain shapes can be created. Problem-specific algorithms seemed to be stepping in the opposite direction.

This paper moves on to consider two examples of engineering dynamics that were encountered by the author after leaving UCL. In both cases there is strongly nonlinear behaviour that originates not from the structural geometry or the material response but from the complications in the applied forcing. The first concerns the phenomenon of synchronous lateral excitation as observed by pedestrians walking on the London Millennium Bridge during its opening days in June 2000. The second concerns the fluid–structure interaction of wind-sensitive structures, particularly long-span bridges, and the numerous aeroelastic instabilities that dominate their design.

2. Pedestrian-Induced Lateral Excitation of Footbridges Caused by Synchronous Walking

The London Millennium Bridge is a multiple-span suspension bridge. The dynamic analysis of such structures is routinely performed via finite element modelling. Such suspension structures can exhibit substantial geometrically nonlinear behaviour which requires some care in assessing. An appropriate analysis technique involves the construction of a geometrically nonlinear static analysis which tensions the cables under the imposed selfweight and some static-average live load. The equilibrium configuration may be located by a standard path-following technique (arc-length or Riks methods, perhaps) by almost any of the many modern nonlinear finite element packages. The small-deflection dynamic behaviour about this equilibrium may then be assessed in the first instance by an appropriate linearisation of the geometical nonlinearities that include tension-stiffening of the cables and pendulum effects. This results in a linear system of the form

$$m\ddot{x} + c\dot{x} + kx = f \tag{2}$$

where x is a vector of nodal displacements (measured relative to the static equilibrium), m, c and k are respectively mass, damping and stiffness matrices and f is a vector of nodal loads. In particular, k is the total tangent stiffness matrix, the word ‘tangent’ signifying that it is a linearisation about the equilibrium configuration of the nonlinear static problem.

The creation of such a model is the first level of approximation, whereby the infinite-dimensional continuum model has been reduced in dimensionality by the finite element discretisation to give an approximate model with perhaps several hundred or even thousand degrees of freedom, as represented by the vector x.

The second level of dimension reduction occurs after extracting the dynamic mode shapes, these being the eigenvectors of the undamped, unforced problem

$$m\ddot{x} + kx = 0 \tag{3}$$

If x has a high dimension, $N \approx 10^3$ say, the model is simplified by keeping only a small number of modes, $n \approx 10$, say, these being the modes with lowest eigenvalue, the low frequency modes. The value of n to choose for the truncation should typically be such that all modes are retained which have frequencies in the range of the dominant dynamic forces. Standard linear theory then leads to a set of n equations

$$M\ddot{q} + C\dot{q} + Kq = F \tag{4}$$

where q are the modal displacements, M, C and K are the mode-generalised mass, damping and stiffness matrices and F is vector containing the mode-generalisation of the applied external forces. Not only is the number of equations to solve now much smaller, but under usual conditions and assumptions, orthogonality considerations lead to M, C and K being diagonal. There are fewer equations, and they are uncoupled and may be solved individually.

A number of complexities can arrive even at this stage. For example, such suspended structures will tend to have modes which involve vertical, lateral, longitudinal and torsional motions and combinations of these. Nonlinear dynamicists familiar with the myriad possibilities of autoparametric and parametric resonance in the spring pendulum [9] will recognise that such possibilities exist here. If appropriate modal frequencies are sufficiently commensurate then the vertical oscillatory motions predicted by the finite element eigenvalue analysis, although being possible free vibration solutions to the equations, may be dynamically unstable. Originally vertical oscillations may therefore lose stability and become predominantly horizontal oscillations at a different frequency. Certainly the possibility exists that a bridge supporting a walking crowd which is exerting large vertical forces may subsequently exhibit large lateral oscillations via this parametric energy-exchange mechanism.

For the most suspension bridges, however, a more immediate nonlinear process should be considered. All the discussion of the introduction focused on the possibilities that arise due to geometric or even material nonlinearities. The linearitisations of these effects lie on the left-hand side of Equation (3). On the right-hand side are the mode-generalised forces F. These forces are caused by humans walking. As a general rule, engineers are trained to deal competently with the possibilities on the structural left-hand side. They are trained less to consider the possibilities of the right-hand side which represents the possible forces exerted by a crowd. A fundamental error would be to assume that $F = A\sin(\omega t)$ and to think that this had covered all possibilities, for human behaviour can be substantially more subtle and unpredictable. For example, F may depend upon x and $\dot{x}$: if the structure is oscillating then people may walk differently and create different forces. The form of F may be rather intangible: it may be gender-specific, and it may depend upon whether the people are drunk, or excited, or dancing (see Figure 1).

Even here, dynamical systems theory can provide useful indications of the possibilities. For example, the work of Strogatz on the synchronisation of systems of coupled oscillators has shed much light on the

Figure 1. Morris dancers on the London Millennium Bridge. Bridge designers should be aware that their structures may be subject to significant dynamic loading scenarios that may not have been envisaged at the design stage. Dancing is a particularly severe load case for grandstands and concert hall auditoria. There is less agreement amongst designers about what the dynamic loading possibilities are from humans than from fluids. [Photo: A. Blagborough]

mechanisms by which such systems can undergo the transition from disorder to synchrony (see [11], for example). The diversity of spheres wherein this transition occurs, including mechanical, electrical, astronomical and biological processes, indicates the prevalence of the effect. Its occurrence in the presence of even weak coupling between oscillators also indicates how difficult the phenomenon will be to "ring-fence". On a bridge there are numerous interactions, between people themselves, between the people and the bridge and between the people and the wider environment. Full-scale physical experiments by the author at Cambridge [10] indicated a number of complex feedback mechanisms whereby pedestrians walking on a laterally flexible surface responded to subtle stimuli. In one experiment, the tactile knowledge of the presence of a hand-rail proved to be sufficient to change the gait of a pedestrian, even though the mechanical forces that were transmitted from the hand-rail to the walker were practically non-existent. In other experiments, adjacent pedestrians who were not walking in synchrony interacted strongly by banging shoulders, with a transition to synchrony removing such undesirable social interaction.

The sensible approach to progress on this matter must involve experiments with real, as opposed to virtual, humans and if possible with real rather than laboratory structures. In the case of the London Millennium Bridge, the actual structure was in existence and experiments with walking crowds could be undertaken on the structure itself. Such experiments were subsequently performed by Arup, the engineers who had designed the bridge [12].

Academic engineers can, on occasion, have a tendency to emphasise mathematical and computational modelling. For that structure, the specific problem was the uncertainty about the possibilities of the right-hand side and their interactions with the left-hand side and fundamental progress was made by full scale experimentation rather than by theoretical analysis. One imagines that future designers, whose structures by necessity do not exist at the design stage, can draw upon the experience and results of those tests. Some designers are even building post-construction walking tests into their project timetables, to verify that the actual structure performs adequately under various pedestrian scenarios. Such designers should however note that the specific scenarios of their own or even of the Millennium Bridge tests (which involved rather orderly walking by sober, well-behaved individuals) may not cover the range of circumstances that their own structure may encounter (see Figure 1). The design against further occurrences of the phenomenon of human–structure interaction becomes not merely a change of focus from the physics of the left- to that of the right-hand side but instead requires a higher level activity: an assessment of the risks of the various scenarios occurring. A framework for undertaking such assessments has been proposed in McRobie et al. [10].

There is a strong parallel between the human–structure and the fluid–structure interaction problem that forms the basis of the next section. Interestingly, one proposal [10] for the codification of footbridge design draws heavily upon this similarity and proposes that all experimental results on pedestrian–footbridge interaction be plotted in terms of dimensionless groups that are largely analogous to the Scruton and Strouhal numbers encountered in wind-engineering.

3. Aeroelastic Instabilities of Flexible Structures

The design of flexible structures such as long-span bridges and tall towers is dominated by considerations of aeroelastic stability. Numerous phenomena need to be considered. These range from buffeting due to wind gustiness, through to the aeroelastic possibilities of flutter, galloping, vortex-induced lock-in, and various other divergent-amplitude possibilities.

Analysis of such structures has traditionally proceeded along the lines already described for the human–structure case in the previous section in that the structural dynamics may be represented by the first few modes of the linearisation of a geometrically nonlinear finite element model, as per Equation (4). Again, the subtleties of the structural aspects of the left-hand side pale in comparison to the complexities of the Navier–Stokes equations, whose moving-boundary integrals give rise to the right-hand side forcing terms.

For each of the aeroelastic phenomena, various low-dimensional models have been proposed. For example, for vortex-shedding, there is the van der Pol oscillator model, as typified by the Hartlen–Currie lift-oscillator model [13]. The rationale behind the proposal of such a model appears to have less to do with projections of the physics from a high-dimensional scheme to a low-dimensional approximation, than to do with the fact that vortex-induced oscillations are observed to be self-limiting in magnitude, as were the electrical oscillations in the circuits studied by van der Pol [14]. Such models thus involve little more than the selection of the appropriate coefficients such that the two phenomena (the van der Pol cycles and the vortex-induced resonance) coincide. Attempts have been made to extend such empirically-based “lift-oscillator” models by inclusion of physical reasoning about the associated fluid behaviour. Examples are the “wake-oscillators” models that follow from the insights of Birkhoff [15] whereby the behaviour of a mass of fluid in the wake is tracked, its behaviour depending upon the oscillation history of the structure. All such models were largely attempts to replicate the experimental results of Feng [16] whose spring-supported cylinders exhibited limit cycles, lock-in and hysteresis. The hysteresis that was observed at very low levels of structural damping has since been observed by few others.

In the study of flutter, there is a historical denouement which encompasses the empirical formulae of Selberg [17], the classical aerodynamic theory of Theodorsen [18] for flat-plates, and the Scanlan's introduction of empirical flutter derivatives for bluff bodies [19]. The latter is the most general, but it has no *ab initio* predictive power as the flutter derivatives cannot be determined *a priori* but must be measured in wind tunnel tests. It is thus not so much a first-principles theory as a means of transferring the results of wind-tunnel tests to full-scale structures.

All such models of vortex-induced oscillation and flutter suffer from a lack of predictive power. The models fit the behaviour over some limited range of validity, but little can be said about what behaviour is to be expected if the shape of the cylinder (or bridge deck) is even slightly different. Here, as was the case for human–structure excitation, there is thus no substitute for physical experimentation, and preferably at full-scale. This however, can present numerous difficulties. As ever, the structures are usually one-of-a-kind designs, and do not exist at full-scale until they have been built. At the design stage, the designer thus does not have access to a full-scale prototype with which to experiment. Physical tests are thus necessarily at model scale, and two problems arise. Firstly, for all the advances in modelling, it is not always possible to remove all scale effects to ensure that model and full-scale tests have a guaranteed one-to-one correspondence. Secondly, for the case of the determination of flutter derivatives, very many model tests are required and these can be both expensive and time-consuming, the latter being particularly problematic if design lead times are tight.

One possible way to progress in this problem is by means of computational fluid dynamics (CFD), and specifically by means of models capable of simulating the multiphysics fluid–structure interaction case. The main obstacle is computation time. The fluid domain needs to be highly resolved in the vicinity of the boundary layer, where vorticity is created. Complex physical processes such as boundary layer separation and reattachment need to be properly resolved on very fine scales if the larger scale dynamics is to be properly replicated. Progress has been made in this arena, with the traditional CFD approach of finite volume formulations, with coupled fluid–structure finite element models, and with grid-free Discrete Vortex formulations. The latter have proved to be particularly instructive in the bridge design case. Larsen's exposition [20] of the interaction between vortex production and structural response gave the first real physically based explanation of the collapse of the first Tacoma Narrows Bridge in 1941.

Figure 2 shows a sample output from a discrete vortex code written by the author. The simulation concerns flow around a square cylinder, and involves around 10^5 discrete vortices. The algorithm is based on that of Morgenthal [21]. Essentially, the vorticity at the body surface is discretised to a series of piecewise linear approximations over short elements (panels) on the body boundary. The velocity induced at the centre of each panel by the vorticity on each of the other panels can be computed by means of the Biot–Savart Law. Imposition of a no-flow boundary condition at the body surface leads to a system of equations which can be solved for the panel vorticity distribution. At each time step, 'bound' vorticity is then released into the adjacent flow as 'free' discrete vortices. The velocity induced at any panel may then be calculated from the effects of the (new, unknown) bound vorticity and the free vortices. Again, the resulting system of equations can be solved for the bound vorticity necessary to ensure the no-flow boundary condition. At each further time step, each free vortex is convected to a new location by computing the velocity induced at its core by all the other free and bound vortices. The vortex–vortex interactions in this convection step involves $O(N^2)$ calculations, where N is the number of vortices. As N becomes large, the procedure becomes computationally prohibitive. Following Morgenthal's approach, numerical performance is improved by projecting all vorticity to a rectilinear grid over the whole fluid–structure domain. The vorticity equations are then solved efficiently (albeit approximately) using Fast Fourier Transforms. In order to maintain the high resolution over short distances comparable to vortex spacing rather than grid size, nearby vortex–vortex interactions are computed exactly using the

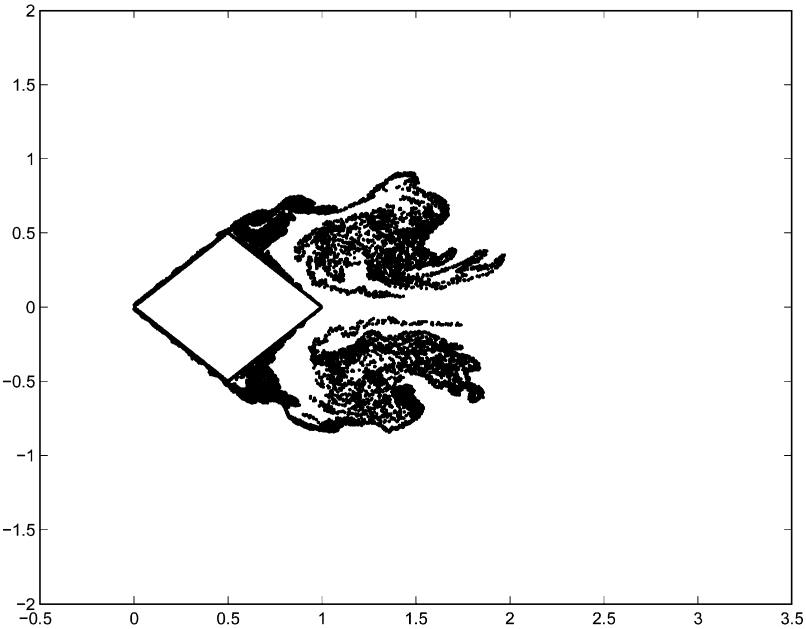

Figure 2. An example simulation of the flow past a square cylinder computed using a high-resolution discrete vortex method. The model involves around 10^5 discrete vortices.

Biot–Savart Law and their contributions via the grid-projection are subtracted exactly. The diffusional part of the Navier–Stokes equations is included by operator splitting. Treatment of the convective part has already been described. The diffusional part is then achieved by applying a suitable random walk to each free vortex after the convection operation. As vortices move further from the body, the resolution can be reduced by merging vortices, or by projecting the vorticity to the grid and computing its subsequent convection on the grid.

Experience with such models has shown that a large number of vortices is required, particularly in the vicinity of the body, if the numerical simulations are to be capable of matching the important features observed in physical experiments.

In both the mesh-based finite-volume formulation and the discrete vortex method, in order to properly predict the effect of a small change in the structural shape on the resulting aeroelastic performance, many thousands of degrees of freedom are required. In the two modelling approaches, these degrees of freedom are, respectively, the nodal values in the very dense meshes over the fluid domain in the vicinity of the structure or the very many discrete vortices that need to be produced and tracked. The obvious question arises as to whether it is possible to reduce the dimensionality of the modelling over and above that which may be achieved by judicious coarsening of the mesh or merging of vortices in regions well away from the viscous boundary layer. On seeing the large scale coherent vortex streets downwind of a structure in the results of a discrete vortex model, numerous dynamicists have commented to the author that some form of principal component analysis should be able to reduce the number of degrees of freedom from the tens of thousands down to something more manageable. The author's response is always to challenge the dynamicist to produce such a low-dimensional model without first having access to the high-dimensional model. A specific challenge would be to ask a low-dimensional model to predict, *ab initio*, the changes to vortex excitation and flutter stability that are caused by the addition of small guide vanes beneath a bridge deck, as occurred on the Great Belt East Bridge in Denmark.

The furthest progress in this direction appears to be the work of Cohen et al. [22] who use a proper orthogonal decomposition of the results of a high-dimensional CFD model of vortex-shedding from a cylinder to create a low-order model which forms the basis of a control algorithm. The massive reduction in computational time achieved by use of the low-order model makes it possible to implement a model-reference control system. By changing the question from "what exactly happens?" to "how can we best control the system?", the dimension reduction is useful. However, one may still question the validity of the predictions of that model and the robustness of the controller outside the parameter regime of the CFD runs from whose results the model has been extracted.

It therefore appears that the proposed low-dimensional models, which would require some non-negligible intellectual input to create, would be capable of little more than interpolating between the results of high-dimensional models. In terms of understanding the full dynamical possibilities, this would be a very minor accomplishment. Since they are at best interpolations, low order models alone would not be able to 'ring-fence' the possible design-critical instabilities. It therefore appears to the author that a prudent combination of physical experimentation and high-dimensional numerical experimentation still presents the only way forward for progress on the design of aerodynamically bluff, flexible structures. Unlike for the case of the human–structure interaction, the fluid–structure interaction at least admits the possibility of progress by modelling since the equations of the physics of a fluid are generally agreed upon.

4. Conclusions

The paper has surveyed recent advances in nonlinear dynamical systems theory. It has looked at two examples from civil engineering dynamics, a human–structure and a fluid–structure interaction case. In both cases, the structural behaviour was largely linear and all complexity occurred in the external forcing. The paper concludes that the high dimensionality of the forcing systems presents a challenge that cannot yet be overcome by anything other than an appropriate combination of high-dimensional simulation and physical experimentation.

In both of the examples presented, there exists an even higher level issue to be addressed. Once the complicated physics and/or biology has been understood, the risk of the occurrence of the envisaged scenarios needs to be addressed. That, however, is wholly another question.

References

1. Thompson, J. M. T., 'Chaotic phenomena triggering the escape from a potential well', *Proceedings of the Royal Society of London* **A421**, 1989, 195–225.
2. Thompson, J. M. T. and Soliman, M. S., 'Fractal control boundaries and thoer relevance to safe engineering design', *Proceedings of the Royal Society of London* **A428**, 1990, 1–13.
3. Marion, M. and Temam, R., 'Nonlinear Galerkin methods', *SIAM Journal of Numerical Analysis* **26**(5), 1989, 1139.
4. Foias, C., Jolly, M. S., Kevrekidis, I. G., Sell, G. R., and Titi, E. S., 'On the computation of inertial manifolds', *Physics Letters A* **131**(7/8), 433.
5. Holmes, P., Lumley, J. L., and Berkooz, G., *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, Cambridge, UK, 1996.
6. Broomhead, D. S., Indik, R., Newell, A. C., and Rand, D. A., 'Local adaptive Galerkin bases for large-dimensional dynamical systems', *Nonlinearity* **4**, 1991, 159–197.
7. Foale, S., McRobie, F. A., and Thompson, J. M. T., 'Numerical dimension-reduction methods for nonlinear shell vibrations', *Journal of Sound and Vibration* **215**, 1998, 527–545.

8. Laing, C. R., McRobie, F. A., and Thompson, J. M. T., 'The post-processed Galerkin method applied to nonlinear shell vibrations', *Dynamic and Stability System* **14**, 1999, 2.
9. Rusbridge, M. G., 'Motion of the sprung pendulum', *American Journal of Physics* **46**, 1978, 1120–1123.
10. McRobie, F. A., Morgenthal, G., Lasenby, J., and Ringer, M., 'Section model tests on human-structure lock-in', *Proceedings of the Institute of Civil Engineering and Bridge Engineering* **156**(BE2), 2003, 71–80.
11. Strogatz, S., 'From Kuramoto to Crawford: Exploring the onset of synchronization in systems of coupled oscillators', *Physica D* **143**, 2000, 1–20.
12. Dallard, P., Fitzpatrick, A. J., Flint, A., Le Bourva, S., Ridsdill Smith, R. M., and Willford, M., 'The London Millennium footbridge', *The Structural Engineer* **79**(22), 2001, 17–33.
13. Hartlen, R. T. and Currie, I. G., 'Lift-oscillator model of vortex-induced vibration', *Journal Engineering and Mechanical Divison* **96**(EM5), 1970, 577–591.
14. Van der Pol, B., 'A theory of the amplitude of free and forced triode oscillation', *Radio Review* **1**, 1920, 701–716.
15. Birkhoff, G., 'Formation of vortex streets', *Journal of Applied Physics* **24**, 1953, 98–103.
16. Feng, C. C., *The Measurement of Vortex-Induced Effects in Flow Past Stationary and Oscillating Circular and D-Section Cylinders*, M.Sc. Thesis, University of British Columbia, Vancouver, BC, 1968.
17. Selberg, A., 'Oscialltin and aerodynamic stability of suspension bridges', *Scandanavica Ci13, Acta Polytechnica*, 1961.
18. Theodorsen, T., 'General theory of aerodynamic instability and the mechanism of flutter', TR 496, NACA, 1935.
19. Scanlan, R. H. and Tomko, J. J., 'Airfoil and bridge deck flutter derivatives', *ASCE Journal of Engineering and Mechanics* **97**, 1971, 1717–1737.
20. Larsen, A., 'Aerodynamics of the Tacoma Narrows Bridge – 60 years later', *Structural Engineering International* **10**, 2000, 4.
21. Morgenthal, G., *Aerodynamic Analysis of Structures Using High-Resolution Vortex Particle Methods*, Ph.D. Thesis, Cambridge University, Cambridge, 2002.
22. Cohen, K., Seigel, S., and McLaughlin, T., 'Sensor placement based on proper orthogonal decompostion modeling of a cylinder wake', in *33rd AIAA Fluid Dynamics Conference*, June 23–26, 2003, Orlando, FL.