



When a Mechanical Model Goes Nonlinear: Unexpected Responses to Low-Periodic Shaking

Author(s): Lisa D. Humphreys and P. Joseph McKenna

Source: *The American Mathematical Monthly*, Vol. 112, No. 10 (Dec., 2005), pp. 861-875

Published by: [Mathematical Association of America](#)

Stable URL: <http://www.jstor.org/stable/30037627>

Accessed: 08/12/2010 13:34

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=maa>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Mathematical Association of America is collaborating with JSTOR to digitize, preserve and extend access to *The American Mathematical Monthly*.

<http://www.jstor.org>

When a Mechanical Model Goes Nonlinear: Unexpected Responses to Low-Periodic Shaking

Lisa D. Humphreys and P. Joseph McKenna

1. INTRODUCTION. This paper had its origin in a curious discovery by the first author in research performed with an undergraduate student. As described in [9], the following odd fact was noticed: *when a mechanical model of a suspension bridge (linear near equilibrium but allowed to slacken at large distance in one direction) is shaken with a low-frequency periodic force, several different periodic responses can result, many with high-frequency components.*

This result ran counter to many earlier results in the literature [1], [4], [5], [6], [7]. There, the main feature investigated was multiple solutions for small-amplitude forcing terms, where the frequency of the forcing term was relatively high, usually in an interval below the linearized resonant frequency. So the conventional wisdom was, if you want to explore for “interesting stuff,” investigate forcing terms with frequency in the vicinity of the resonant frequency.

Not only were these new solutions unexpected, they were also somewhat counterintuitive. One’s first guess would be that reducing the restoring force toward equilibrium should actually cause an even lower frequency response, rather than introducing new high frequencies. It also was a little unnerving, in that the long-term response of the mechanical model to forcing was quite unpredictable and would be highly dependent on initial conditions as soon as the model was forced out of the linear range.

Several periodic responses were observed by the relatively naive method of feeding in initial conditions, solving an initial-value problem for a large time, and observing the eventual response. Of course, there is no theorem that says this must result in convergence to a periodic solution, but it happened a gratifying number of times.

The natural question is: How does one go from the simple near-equilibrium linear behavior, where there is one solution that matches the forcing, to the manifoldly more complex situation where the amplitude of the forcing pushes the oscillator into the nonlinear region? In other words, what is the global structure of the solution set? A full picture of the solution set involves finding not just stable solutions but the unstable ones as well. This means we need to find ways of identifying periodic solutions other than just solving initial-value problems for long times and hoping for convergence.

A word of warning before we begin: some readers may think that just because we have a nonlinear equation, we should expect complicated behavior. This is not true. Most mechanical systems oscillate into a nonlinear range to some extent. And most do not exhibit unpredictable behavior. Indeed, most suspension bridges function predictably, even in earthquakes. The long-term goal of our analysis is to identify the characteristics of those systems that behave unpredictably and to find ways to eliminate these behaviors.

The methods we use include Newton's method, continuation methods, and the method of steepest descent.¹ All have limitations, and none worked one hundred percent of the time. However, a combination of methods, and a lot of persistence, have finally given us a picture of astonishing complexity in a very simple model.

In section 2, we describe the model. In section 3, we briefly outline the (mostly low-tech) methods that we use. Our hope is that this paper can stimulate similar exploration in other undergraduate settings. In section 4, we describe our findings on what the solution set looks like, including period-doubling and period-tripling solutions in astounding multiplicity. In section 5, we present conclusions and suggest areas for further research.

2. THE DEVELOPMENT OF THE MODEL. The basic idea of the model that we investigate was developed in [6], [11], and [12]. A nice exposition can be found in [2]. The model assumes that cables resist expansion but not compression. Thus, if we simplify a suspension bridge to a single oscillating particle, we expect to find two different restoring forces acting on it. The first is a unilateral force from the cables that hold the bridge up (but not down), and the second is a linear force that resists deflection from equilibrium in both directions. This second force is due to the inherent stiffness of the roadbed. Some suspension bridges are, like the George Washington Bridge, very stiff, whereas others, such as the original Tacoma Narrows Bridge, have little resistance to vertical deflection. The weight of the bridge pushes the bridge down and extends the cables to an equilibrium around which the forces from the cables remain linear. However, if the upward deflections become too large the cables slacken, and this part of the restoring force should disappear.

These considerations lead to the following model. Consider a mass attached to a vertical spring and a cable providing additional support. While the spring causes a restoring force in both the upward and downward directions, the cable only resists expansion. Let $y(t)$ denote the downward displacement of the mass at time t , where $y = 0$ denotes the position before elongation of the spring by the addition of the mass (see Figure 1).




Figure 1. Two different types of forces hold up the mass: a linear spring that resists displacement in either direction and a cable-like force that resists displacement only in the downward direction.

There are three main forces acting on the mass: a restoring force from the cable, a linear restoring force from the spring, and gravity. Without the cable, the restoring

¹Both authors have spent a considerable portion of their careers on numerical methods of finding “non-obvious” periodic solutions of nonlinear ordinary and partial differential equations, using the aforementioned methods. The version of steepest descent that we use here is apparently new, remarkably effective, and very easy to implement.

force in the linear model in the spring is given by $k_1 y$. A cable has a restoring force of $k_2 y^+$, where $y^+ = \max\{y, 0\}$, because it resists deflections in the downward direction (expansion) but not in the upward direction (compression). For example, in the George Washington Bridge, k_1 would be large and k_2 relatively small. In the more flexible Tacoma Narrows Bridge, k_1 would be small and k_2 relatively large. Thus the combined restoring force is given by $ay^+ - by^-$, where $a > b$ and $y^- = -\min\{y, 0\}$. We are interested in the response of this system to periodic forcing, and we take $f(t) = \lambda \sin \mu t$ as our generic periodic force. A small amount of damping leads us to the following model:

$$y'' + 0.01y' + ay^+ - by^- = 10 + \lambda \sin \mu t. \quad (1)$$

For the situation of [9], with $\mu = .17$, we expect complicated behavior if λ is large but simple linear behavior when λ is small. The model relies on the cables to hold up the mass, so we take the linear spring constant to be relatively small ($k_1 = 1$), while the cable constant k_2 is taken to be the relatively large value $k_2 = 16$. We take $a = 17$ and $b = 1$. In the absence of forcing, the mass is at equilibrium at $y = 10/17$, and the equation remains linear until the oscillation magnitude approaches $10/17$ and $y(t)$ becomes negative.

The most significant property of the model as we have chosen it is that, with the large difference between a and b , the equation is “very” nonlinear. For example, if we had chosen $a = 17$ and $b = 7$, none of the strange behavior that we shall later discover in the presence of low-frequency forcing would exist. So our question becomes: What do the solutions of the boundary value problem (1) with boundary conditions given by $y(0) = y(2\pi/\mu)$, $y'(0) = y'(2\pi/\mu)$ look like as we vary λ from small to large amplitudes?

When describing the structure of the set of solutions of a nonlinear equation that includes a parameter, we will use the language of bifurcation theory. We start with a simple example. Consider the algebraic equation

$$F(u, \lambda) \equiv u((u - 1)^2 - \lambda) = 0.$$

There are a variety of ways of describing the solutions of this as λ varies. One could observe that for $\lambda < 0$, $u = 0$ is the only solution, whereas for $\lambda > 0$ there are the three solutions $u = 0$ and $u = 1 \pm \sqrt{\lambda}$. The easiest way to convey this information is to view the picture in Figure 2. In a sense, this is typical of most nonlinear equations with an




Figure 2. The solution set for $F(u, \lambda) \equiv u((u - 1)^2 - \lambda) = 0$. Note how the solutions lie on curves that occasionally turn round and intersect but are “mostly” parametrisable in λ .

embedded parameter. “Most of the time” the solutions are on curves parametrizable by λ . Sometimes these curves turn around, as at $(0, 1)$, and we would need another parameter such as arclength to describe them. Another thing that can go wrong is that curves can collide with each other, as at $(1, 0)$. This is called a *bifurcation point*.

Another instructive aspect of Figure 2 is that it suggests a way to find these solutions if one were extremely analytically challenged. Having established some point on the curve (u_0, λ_0) , one could then use this value of u_0 as a good initial guess for the equation $F(u, \lambda_1) = 0$ for nearby $\lambda_1 = \lambda_0 + \delta\lambda$. Then one could employ Newton’s method $u_{n+1} = u_n - (D_u F(u_n, \lambda_1))^{-1} F(u_n, \lambda_1)$ to converge to the nearby point of the curve (u_1, λ_1) . In this way, by iterating, one could hope to trace out the various pieces of these curves numerically. This is the intuitive idea behind one of our main tools, *continuation methods*.

This simple idea would break down in two places, namely at $(0, 1)$, the turning point, where continuing to increment λ in the negative direction would enter a regime where there is no nearby solution and Newton’s method, deprived of a good initial guess, would probably fail. It would also break down at $(1, 0)$, the bifurcation point, where $D_u F(u_n, \lambda_1)$ goes to zero. Other techniques would be needed at these points.

The same intuitive idea holds when trying to describe periodic solutions of equation (1) with the embedded parameter λ . If, for fixed λ , we have a solution, then typically, as we vary λ , we expect nearby solutions to depend continuously on λ . In other words, “most of the time,” when we vary λ continuously, the solutions trace out a curve in function space. These curves cannot simply die out at any point but are always continuable [10]. As in our first simple example, they may collide at bifurcation points or fail to be parametrizable in λ at turning points. We call these curves of solutions *bifurcation curves*. Of course, they are no longer nice simple curves in the plane but now reside in some infinite-dimensional function space.

To visualize these curves of solutions in infinite dimensional space, we project onto two dimensions and plot λ versus the amplitude of the corresponding periodic solution of (1). In a slight abuse of terminology, we also call these two-dimensional projections bifurcation curves. An example is shown in Figure 3. As mentioned earlier in the discussion of the model, small λ means the solution stays in the linear range. This is seen in Figure 3 in the region where λ is smaller than about 9.9. Since the equation is linear in this range, the bifurcation curve is linearly dependent on λ . However, as increasing λ pushes the solutions into the nonlinear range, the curve bends backwards and forwards. In particular, when $\lambda \approx 10.5$ there are at least three solutions. Describing the complete set of solutions for the entire range of λ is rather ambitious, as the set of curves may be extraordinarily complicated and one can never be sure that there are no other curves of solutions somewhere out there in unexplored areas of initial conditions. Nonetheless, this is what we attempt in this paper.

3. NUMERICAL TECHNIQUES. We begin by reviewing some basic material from ordinary differential equations. For any pair (c, d) the initial-value problem

$$y'' + 0.01y' + ay^+ - by^- = 10 + \lambda \sin \mu t, \quad y(0) = c, \quad y'(0) = d \quad (2)$$

has a unique solution on $[0, T]$ that is continuously and differentially dependent on c and d . Finding a T -periodic solution of (2) is equivalent to finding a pair (c^*, d^*) so that when the initial-value problem is solved with $c = c^*$ and $d = d^*$, the resulting solution satisfies $y(0) = y(T)$ and $y'(0) = y'(T)$. Strictly speaking, we should write $y(T, c, d)$, since the definition of $y(T)$ depends on the initial values (c, d) . As remarked earlier, $y(T, c, d)$ is differentially dependent on c and d .

This problem was attacked in several ways. We began with techniques whose underlying concepts are accessible to most undergraduates and readily available in standard mathematical software. We modified these algorithms to fit our specific equation and then used more complicated methods to finish the analysis.

Since we are searching numerically for periodic solutions, we have to be content with approximating T-periodic solutions, which are obtained by running individual algorithms until *the quantity* $(y(0) - y(T))^2 + (y'(0) - y'(T))^2$ is less than a prescribed tolerance, which we refer to as the *error*.

Initial-value problem methods. Before discussing the first method, we need to remind the reader that in regard to periodic solutions of nonlinear problems there are two fundamentally different types. We call a solution *stable* if initial values close to the initial data of the solution yield convergence to that solution in large time. In the case of *unstable* solutions, one can start off arbitrarily close to a solution and eventually end up far away from it.

Our goal is to find periodic solutions of (1). We begin by treating the boundary value problem as an initial-value problem and observing long-term behavior. Almost all standard software products have built in features to do this and are easily accessed. We wrote our own fourth order Runge-Kutta solver, varied the initial conditions, solved the initial-value problem over a large fixed time, and looked to see what resulted. This was a way of finding many different stable periodic solutions. What makes this technique work is that many stable solutions exist and, if one ever comes close to one, one expects to converge to it in large time. Of course, this method has no chance of finding unstable solutions. Another drawback is that the limiting solution needs further refinement. Typically we observed error on the order of 10^{-8} . By using other techniques, we were able to do much better.

Newton's method. The power and speed of Newton's method for finding zeros of a function of one variable is well known. Indeed, it is taught in calculus classes. This procedure works equally well in finding zeros of vector-valued functions. If we define the function

$$F\left(\begin{bmatrix} c \\ d \end{bmatrix}\right) = \begin{bmatrix} c \\ d \end{bmatrix} - \begin{bmatrix} y(T, c, d) \\ y'(T, c, d) \end{bmatrix}, \quad (3)$$

then F is a map from R^2 to R^2 and zeros of F correspond to T-periodic solutions of (1). As in [9], we implemented a two-dimensional Newton solver. We searched for the initial conditions that would put us on a periodic solution.

Here we briefly outline the Newton solver. Using vector notation and letting c and d denote the initial position and initial velocity, respectively, let $G\begin{bmatrix} c \\ d \end{bmatrix}$ denote the position and velocity of a solution to the equation at $T = 2\pi/\mu$. Thus, finding periodic solutions of (1) is equivalent to finding zeros of the function F defined by

$$F\begin{bmatrix} c \\ d \end{bmatrix} = \begin{bmatrix} c \\ d \end{bmatrix} - G\begin{bmatrix} c \\ d \end{bmatrix}. \quad (4)$$

So to implement Newton's method, we follow the familiar iterative scheme

$$z_{n+1} = z_n - (DF)^{-1}(z_n)F(z_n), \quad (5)$$

which in our notation becomes

$$\begin{bmatrix} c_{n+1} \\ d_{n+1} \end{bmatrix} = \begin{bmatrix} c_n \\ d_n \end{bmatrix} - \begin{bmatrix} 1 - \frac{\partial G_1}{\partial c} & -\frac{\partial G_1}{\partial d} \\ -\frac{\partial G_2}{\partial c} & 1 - \frac{\partial G_2}{\partial d} \end{bmatrix}^{-1} F \begin{bmatrix} c_n \\ d_n \end{bmatrix}. \quad (6)$$

There are two different ways to calculate the partial derivatives to be used in this approach. One simple method is to use a central difference scheme. This involves evaluating terms like

$$(G(c + h, d) - G(c - h, d))/(2h)$$

for suitably small h . In other words, to calculate each partial derivative, one must solve the initial-value problem (1) over the interval $[0, T]$ twice. This is the method that was used in [9] and is probably best when working with undergraduates.

Another more efficient (and more complicated) way to calculate the partial derivatives required in the derivative matrix of equation (6) is to differentiate (2) with respect to the initial conditions c and d . This results in two additional differential equations. Thus, to implement one iteration of the Newton method, we simultaneously solve the following coupled system of three equations:

$$y'' + 0.01y' + ay^+ - by^- = 10 + \lambda \sin \mu t, \quad (7)$$

$$\left(\frac{\partial y}{\partial c}\right)'' + (0.01)\left(\frac{\partial y}{\partial c}\right)' + 17 \begin{Bmatrix} \frac{\partial y}{\partial c} & y > 0 \\ 0 & y \leq 0 \end{Bmatrix} - \begin{Bmatrix} 0 & y > 0 \\ -\frac{\partial y}{\partial c} & y \leq 0 \end{Bmatrix} = 0, \quad (8)$$

$$\left(\frac{\partial y}{\partial d}\right)'' + (0.01)\left(\frac{\partial y}{\partial d}\right)' + 17 \begin{Bmatrix} \frac{\partial y}{\partial d} & y > 0 \\ 0 & y \leq 0 \end{Bmatrix} - \begin{Bmatrix} 0 & y > 0 \\ -\frac{\partial y}{\partial d} & y \leq 0 \end{Bmatrix} = 0, \quad (9)$$

with the initial conditions

$$\begin{aligned} y(0) = c, \quad y'(0) = d, \quad \left(\frac{\partial y}{\partial c}\right)(0) = 1, \\ \left(\frac{\partial y}{\partial c}\right)'(0) = 0, \quad \left(\frac{\partial y}{\partial d}\right)(0) = 0, \quad \left(\frac{\partial y}{\partial d}\right)'(0) = 1, \end{aligned} \quad (10)$$

over an interval of length $(0, 2\pi/\mu)$. Having solved the initial-value problem (7)–(9), we note that the term $\partial G_1/\partial c$ is given by $\partial y/\partial c(2\pi/\mu)$ and $\partial G_2/\partial c$ is given by $(\partial y/\partial c)'(2\pi/\mu)$ and similarly for partial derivatives with respect to d . (A word of warning: it is not at all clear that these derivatives really exist! Certainly we end up solving a system in which many of the terms are discontinuous, a fact that should give us a certain discomfort with the quality of the solutions of the initial-value problem. As long as the solutions are not identically zero on an open interval, one can justify this approach [16].) To find appropriate initial conditions that yield the periodic solution we iteratively compute $z_{n+1} = z_n - (F')^{-1}[z_n]F[z_n]$ until our error is sufficiently small.

Any application of Newton's method relies heavily on a good first guess. If we use this approach to refine a solution found by solving the initial-value problem, as in the previous subsection, then we have an excellent starting point. After such a refinement we typically found solutions with error on the order of 10^{-20} .

Naturally, Newton's method doesn't work if we haven't made a good guess at the start. If, for example, we wanted to find a periodic solution at a fixed λ without any hint as to what we were looking for, Newton's method usually wouldn't be successful. We needed a third, more global method.

Steepest descent. Thus far, we have discussed finding stable solutions through initial-value methods and refining those solutions with Newton's method. We could have used the Newton approach to find unstable solutions as well, but we lacked the good initial guesses that this method requires. We found that the method of steepest descent was an effective way of getting around this difficulty.

This method is based on a simple and familiar idea from multivariable calculus, namely, the directional derivative. To find a minimum of a scalar function f of two variables, c and d , one evaluates the gradient ∇f at some initial point and, assuming that this vector is not zero, then moves in the opposite direction for a small distance. So long as the distance is small, this should lower the value of the function. Iterate! When you can't reduce the value of the function any further, you have presumably found a critical point, which is generically a minimum.

As far as we know, this method has not previously been used to find periodic solutions of nonlinear ordinary differential equations. The idea is to find a zero of the error function given by

$$E(c, d) = (c - y(2\pi/\mu, c, d))^2 + (d - y'(2\pi/\mu, c, d))^2. \quad (11)$$

Notice that we have emphasized the differentiable dependence of $y(2\pi/\mu)$ on the initial condition pair (c, d) . As remarked, one way to find zeros of this function (i.e., the places where it takes its minimum value) is to start at an arbitrary point (c, d) and move in the direction opposite the gradient. Naturally, this involves computing the partial derivatives $\partial E/\partial c$ and $\partial E/\partial d$. One way to do this would be to evaluate finite differences of the form

$$(E(c + h, d) - E(c - h, d))/(2h)$$

for some small h . Calculating this one finite difference involves solving the initial-value problem two more times on the interval $[0, 2\pi/\mu]$ and twice more to find $\partial E/\partial d$.

Let $T = 2\pi/\mu$. Another way to calculate the gradient is to take the partial derivatives of equation (11) with respect to c and d . This gives

$$\begin{aligned} \frac{\partial E}{\partial c} &= 2(c - y(T)) \left(1 - \frac{\partial y}{\partial c}(T) \right) + 2(d - y'(T)) \left(- \left(\frac{\partial y}{\partial c} \right)'(T) \right), \\ \frac{\partial E}{\partial d} &= 2(c - y(T)) \left(- \frac{\partial y}{\partial d}(T) \right) + 2(d - y'(T)) \left(1 - \left(\frac{\partial y}{\partial d} \right)'(T) \right). \end{aligned}$$

The various partial derivatives like $\partial y/\partial c(T)$ are exactly the ones we have already discussed in Newton's method.

We then iterate $z_{n+1} = z_n - h \nabla E$, where h is taken to be sufficiently small. This procedure is relatively easy to implement since the calculations for all the partial derivatives have already been computed in the Newton solver.

Steepest descent had, for us, the advantage of finding solutions, including unstable ones, without having a good guess. This method proved unreasonably effective, especially in finding solutions that were not connected to any obvious branches of solutions [8].

A priori, there was no reason why this process should not converge to some local minimum of the function E that was not zero, (and therefore not a periodic solution), but most times it did in fact converge to zero, giving us a periodic solution. One curious result of our calculations was the observation that the only local minima to which the process converged were zeros of the scalar function $E(c, d)$. One cannot help wondering whether this is a theorem, or whether there are other nonzero local minima.

Continuation algorithms. To get a sense of the complete picture of the solution set, it is not very informative to say simply that for a given λ there are, say, ten solutions, and then just show their pictures. We would like to exhibit a more complete picture of how these solutions get distributed as we vary the external forcing λ and to gain some sense of how certain solutions are related to others. As mentioned earlier, mapping out as many of these bifurcation curves as we can will give us a more complete picture. This is the job of the continuation algorithm. It will generate the bifurcation curves shown in Figures 3, 4, and so on.

Although the implementation can be very technical, the basic idea is intuitive. Suppose we have found a solution for a given value of λ by some other method. As mentioned, we expect this solution to lie in a continuum of solutions that exist for nearby λ . These nearby solutions should be close to the one we have already found. Therefore, our initial data for the known solution should be a very good guess for nearby solutions. Since a reasonable initial guess is exactly what is required for Newton's method, we should be able to find the nearby solutions and then trace out an arc of solutions by inching along the curve in solution space in tiny increments of λ . One good place to start this process is when λ is very small, since $y = 0$ is a natural initial guess.

To visualize the resulting curves of solutions, we plot λ against the amplitude A of the corresponding solution. This simple intuitive idea breaks down as we approach turning points where we may not be able to increase λ to find another solution. For example, it would work well on the lowest arc of Figure 3 but will fail when $\lambda = 10$ as we reach the first turning point, at which the derivative matrix becomes singular. The bifurcation curve can no longer be parametrized in terms of λ .

There are various ways to fix this problem, some quite difficult. One is to parametrize the curve in terms of *arclength*. This allows us to sometimes pass round the turning points. This method is quite complicated and beyond the scope of this paper. Details can be found in [3] or [10].

Another simple fix, if we suspect that a turning point is near and that there may be no solutions for increased λ , is to hunt around using steepest descent or Newton's method till we find another point on the other side of the turning point. We used both these methods.

Having found another solution on the other side of the turning point, we can then start up the simple continuation procedure again. Steepest descent could be employed to search for another solution at a fixed λ nearby. One could piece together several branches from this simple continuation and get an idea of the appearance of the solution space. Alternatively, we could inch along the branch using steepest descent to find solutions as we incremented λ .

The bottom line. None of the methods that we have described was perfectly reliable. Often the matrix in Newton's method would become singular as we moved along a branch. Typically, we would then search with the steepest descent method until we found another point on the branch, after which we would try to work backwards from there with continuation. Each of the global pictures shown in this paper represents many different pieces patched together.

4. THE COMPLETE PICTURE. We begin the bifurcation curve of the main branch starting at $\lambda = 0$. Recall that the horizontal axis records λ , while the vertical axis represents the amplitude of the resulting solution. Prior to about $\lambda = 9.9$ the linear solution exists and appears to be the only one, at least until we are close to nonlinear amplitudes. This can be seen in Figure 3. As we approach the region where the solution becomes nonlinear, things change rapidly.




Figure 3. The bifurcation curve starting at $\lambda = 0$. The graph plots λ versus the amplitude of the corresponding solution. Note that for λ near zero, the graph is linear and remains so until λ approaches 10, where we enter the nonlinear regime.

As we climb up the branch of the bifurcation curve that starts at zero, we begin with the linear solution. Then as we move into the nonlinear range, the solution immediately begins to exhibit a slight wiggle. As we progress up the branch, we can clearly see a high frequency response superimposed on the outline of the linear solution. This high frequency response increases as we move up the path in Figures 3 and 4. Figure 6 shows solutions as we go up this branch.

Of course, there is no reason to limit oneself to continuation from the zero solution. As mentioned in [9], several solutions were found at $\lambda = 13.7$. We started the continuation algorithm on these solutions and discovered several isolated branches. In Figure 5, we superimpose these branches on the primary branch from Figures 3 and 4.

Period-doubling bifurcations. As we tested for stability along the main branch shown in Figures 3 and 4, we discovered an anomaly. We expected and observed a change of stability as we rounded turning points. We also discovered that along certain arcs away from turning points, we could likewise detect a change of stability. This led us to suspect some other form of bifurcation, and after searching in the neighborhood of these points, we found solutions whose periods were twice the period of the forcing term. This led to new branches, which came off the main branch and later rejoined it. These are shown superimposed on the period-one branches in Figure 8.




Figure 4. A better view of the main branch from Figure 3. When λ is 9.9, notice how we get five different solutions for the same forcing term.




Figure 5. The complete structure of the various solution paths, including three isolated branches. The main branch is shown with a dotted line and the three isolated branches are shown with solid lines. Of course, the isolated branches only appear to intersect the main branch, an artifact of projecting an infinite-dimensional space into R^2 . Around $\lambda = 13$, there are at least seven solutions for the same forcing term.

Period-three solutions. On a hunch, we looked for solutions whose period was three times that of the forcing term. We did not expect to find any because of the multitude of period-one solutions, which of course are also of period three. However, steepest descent did find some. Having found these, we were able to use continuation to find more. They did not seem to connect up to the main branch but existed off on their own




Figure 6. Four solutions for various λ in (9.8, 14) as we move up the bifurcation curve of Figures 3 and 4 that starts at 0. For each solution we plot $(t, y(t))$ over two periods. As we climb the path, a pronounced high frequency wiggle becomes more noticeable. As we round a turning point, a stable arc turns unstable, and vice versa. In the top left, we can see the solution making only a tiny incursion into the half-plane $y \leq 0$, and already there is a pronounced high-frequency oscillation.




Figure 7. Two of the period-two solutions. For each solution we plot $(t, y(t))$ over four periods. Note on the left how the period-two solution goes through one period of small vibration followed by a second period of more violent vibration.

in function space. Two samples of these are illustrated in Figure 9, and the continua we found are in Figure 10.

5. CONCLUSIONS AND OPEN PROBLEMS. Our main conclusion is surprising: as soon as the supporting cables lose tension under very low-frequency forcing, even for a tiny percentage of the period, high-frequency oscillation becomes an inevitable




Figure 8. The main branch of the bifurcation curve displayed with the period-two-branches. Compare with Figure 3. Again, the main branch is shown with a dotted line, while the new period-two curves are shown with solid lines. There is one isolated figure-eight continuum of double-period solutions and three that start off on a lower arc and rejoin the main branch further up the continuum. At least two of the additional branches appear to be quite similar.




Figure 9. Two of the period-three solutions. For each solution we plot $(t, y(t))$ over two periods (or six periods of the forcing term). The solution on the left is particularly counterintuitive, with its alternating large and small shaking.

part of the picture. As the amplitude is increased further into the nonlinear range, the solutions become less and less intuitive, and the high-frequency shaking becomes more widespread and violent.

This may have some implications for vibrations in cable-stayed bridges, where unexpected high-frequency oscillations have been observed. Ships encounter large-amplitude, low-frequency forcing from waves and obey Archimedes' law, which becomes nonlinear when the ship lifts out of water. They occasionally exhibit puzzling high-frequency oscillations [15].

We emphasize that the ideas underlying most of what we have said are accessible to good undergraduates. Some variations on this theme would make for interesting re-

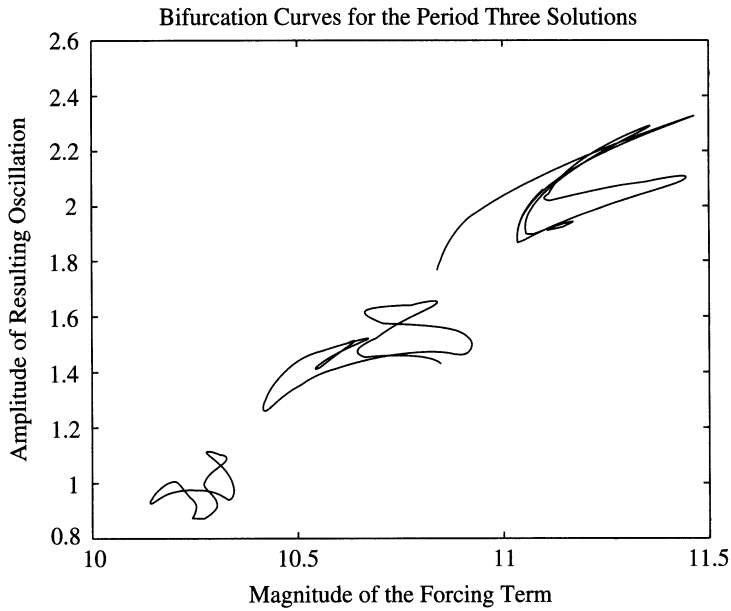


Figure 10. The continua of period-three solutions. They don't seem to connect to any other lower-period solutions.




Figure 11. All solutions of period one, two, or three! Challenge to the reader: find the λ with the most periodic responses.

search projects. A certain amount of patience is required, since we found that continuing along arcs of a solution branch was a highly unpredictable process, with Newton's method becoming singular for no apparent reasons. We then had to hunt around for another point on the branch and continuously increment our way backwards. Modulo that cautionary note, this problem gives students fairly easy access to the sheer fun, unpredictability, and occasional frustrations of research in applied mathematics.

We conclude with some open questions:

1. What is the structure of the solution set as one varies the frequency of the forcing term from the extremely low values used here to the higher values used in earlier studies, say, around 4? (Varying μ a little gives substantially similar pictures. Higher up, near $\mu = 4$, the pictures become simple S-curves.)
2. Can we find the basins of attraction of the various solutions for a given value of λ ? In other words, can we find a pattern in the sets of initial conditions that eventually lead to the different solutions? A start in this direction was made in [9].
3. Can we find more examples of cascades of period doubling as was found in [4]? In other words, do the period-two solutions give period-four solutions? Or eight?
4. Are there other more exotic solutions such as the one in Figures 6 and 8 that alternately go quiet and shake violently? Maybe there is a period four solution that is relatively quiet for three periods but shakes violently for the fourth.
5. Are there any solutions of period other than one, two, or three? (We searched a little but didn't find any.)
6. Is there any chaos in this system with periodic forcing?
7. What happens if one has a two-mass system or two degrees of freedom? (See, for example, [12], [13], [14], and [5]). There are multiple solutions, but the general structure of the solution set is still a mystery in these systems.

REFERENCES

1. N. Ben-Gal and K. S. Moore, Bifurcation and stability properties of periodic solutions to two nonlinear spring-mass systems, *Nonlinear Anal.* **61** (2005) 1015–1030.
2. P. Blanchard, R. L. Devaney, and G. R. Hall, *Differential Equations*, 2nd ed., Brooks/Cole, New York, 2002.
3. Y. S. Choi, K. C. Jen, and P. J. McKenna, The structure of the solution set for periodic oscillations in a suspension bridge model, *IMA J. Appl. Math.* **47** (1991) 283–306.
4. S. H. Doole and S. J. Hogan, A piecewise linear suspension bridge model: Nonlinear dynamics and orbit continuation, *Dyn. and Stab. Syst.* **11** (1996) 19–47.
5. ———, Non-linear dynamics of the extended Lazer-McKenna bridge oscillation model, *Dyn. and Stab. Syst.* **15** (2000) 43–58.
6. J. Glover, A. C. Lazer, and P. J. McKenna, Existence and stability of large-scale nonlinear oscillations in suspension bridges, *Z. Angew. Math. Phys.* **40** (1989) 171–200.
7. L. Humphreys, Numerical mountain pass solutions of a suspension bridge equation, *Nonlinear Anal.* **28** (1997) 1811–1826.
8. L. D. Humphreys and P. J. McKenna, Using a gradient vector to find multiple periodic oscillations in suspension bridge models, *College Math. J.* **36** (2005) 16–26.
9. L. Humphreys and R. Shammass, Finding unpredictable behavior in a simple ordinary differential equation, *College Math. J.* **31** (2000) 338–346.
10. H. B. Keller, *Lectures on Numerical Methods in Bifurcation Problems*, Springer-Verlag, Berlin, 1987.
11. A. C. Lazer and P. J. McKenna, Large scale oscillatory behaviour in loaded asymmetric systems, *Ann. Institut H. Poincaré Anal. Non Linéaire* **4** (1987) 243–274.
12. ———, Large-amplitude periodic oscillations in suspension bridges: some new connections with non-linear analysis, *SIAM Rev.* **32** (1990) 537–578.
13. P. J. McKenna, Large torsional oscillations in suspension bridges revisited: Fixing an old approximation, this MONTHLY **106** (1999) 1–18.
14. P. J. McKenna and Cillian ÓTuama, Large torsional oscillations in suspension bridges visited again: Vertical forcing creates torsional response, this MONTHLY **108** (2001) 738–745.
15. H. E. Saunders, *Hydrodynamics in Ship Design*, Soc. Naval Arch. Press, New York, 1965.
16. S. Solimini, Some remarks on the number of solutions of some nonlinear elliptic equations, *Ann. Institut H. Poincaré Anal. Non Linéaire* **2** (1985) 143–156.

LISA D. HUMPHREYS earned her Ph.D from the University of Connecticut in 1994. She has published papers in various areas of nonlinear partial differential equations, including mountain-pass techniques, continuation methods, and research involving undergraduates. In her spare time, she enjoys her two young children as well as reading, gardening, and weightlifting.

Department of Mathematics and Computer Science, Rhode Island College, Providence, RI 02908
lhumphreys@ric.edu

P. JOSEPH MCKENNA earned his B.Sc. from University College Dublin and his Ph.D. from the University of Michigan in 1976, under the supervision of Lamberto Cesari. He received the Lester Ford Award in 2000 and the 2004 Northeast Section Award for Distinguished College or University teaching of Mathematics. According to ISIhighlycited.com, he is one of the most highly cited researchers working in the mathematical sciences. His spare time is devoted to cooking, theater, opera, and detective novels.

Department of Mathematics, University of Connecticut, Storrs, CT 06269
mckenna@math.uconn.edu

