

Projective Methods for Stiff Differential Equations: *problems with gaps in their eigenvalue spectrum*

C. W. Gear*
Ioannis G. Kevrekidis†

February 4, 2002

Abstract

We show that there exist classes of explicit numerical integration methods that can handle very stiff problems if the eigenvalues are separated into two clusters, one containing the “stiff”, or fast components, and one containing the slow components. These methods have large average step sizes relative to the fast components. Conventional implicit methods involve the solution of non-linear equations at each step, which for large problems requires significant communication between processors on a multiprocessor machine. For such problems the methods proposed here have significant potential for speed improvement.

Keywords Integration, stiff, explicit, stability

1 Introduction

In this paper we consider explicit numerical methods for problems with a large gap between the time constants of the fast components (which are assumed to be damped out after a short time) and the time constants of the slow, *active* components (those still present in the solution). The fast, damped components arise from eigenvalues with large negative real parts, while the active components can arise from driving terms in a non-autonomous system or from eigenvalues of smaller magnitude. It will be convenient to think of the problems as having two sets of eigenvalues: one with very negative real parts and a second set that are close to the origin. The *gap* is, loosely, the distance between these sets.

In the absence of fast driving terms, the fast components in the true solution corresponding to the large eigenvalues will be rapidly damped. If we had a very good numerical integrator for the fast region, the numerical solution would have the same properties.

*NEC Research Institute, retired

†Princeton University, Dept of Chemical Engineering. Supported in part by AFOSR (Dynamics and Control program, Dr. Mark Jacobs)

Such an “inner” integrator could be an explicit one, using step sizes of the same order of magnitude as the fast time constants. It could, however, be any method provided that it is accurate and damping. For example, it might involve the simulation of the underlying system at a different level of modeling detail: it could be a Lattice-Boltzmann simulation or even a Molecular Dynamics simulation of a flow, a more fine-grained model than the Navier-Stokes description of the same problem. Indeed, a major motivation for this approach, which will be pursued in subsequent publications, is the desire to be able to exploit short-term simulation results of existing “fine” microscopic model codes and to derive from them results over long time scales. This will be done by a “coarse-motivated” post-processing through the “outer” integrator.

Once beyond the region where the fast components are active, we would like to use step sizes in a numerical integration commensurate with the slow components. Conventional wisdom tells us that we have to use implicit methods in this case, and then to solve the resulting coupled nonlinear equations using some approximation to the system Jacobian or of the subspace corresponding to the large eigenvalues (e.g. by a Krylov technique through an approximation of the dominant subspace of the Jacobian). The methods we propose do not require any approximation to the Jacobian or any representation of its dominant subspace.

We are concerned with the integration of the initial value ODE

$$y' = f(y), \quad y(0) = y_0 \tag{1}$$

where y and f are n -dimensional vectors and n is typically large.

We wish to consider problems in which it is not practical to use implicit methods. This could occur because we already have a large code that implements an explicit integration and it is not practical to re-engineer that code; it could occur because of the size of the problem, or because nonlinearities make the solution of the system of nonlinear equations impossible when large steps are used. (Although we will talk in terms of ordinary differential equations, the primary application is undoubtedly to problems arising from the semi-discretization of partial differential equations.)

Intuitively, the idea is very simple: we take a small number of steps of an *inner integrator* at a time scale corresponding to the fast time constants until those components are heavily damped. Then we perform a (polynomial) projection (or extrapolation) forward over a long step commensurate with the slow time constants from the results of the inner integration. The inner integrator can be an explicit method since it is using a small step (although it does not have to be). For now, the reader may find it convenient to think of the inner integrator as a “perfect integrator” - one that exactly reproduces the behavior of the differential equation over very small temporal step sizes. Of course, such an integrator does not exist except for those problems whose integral can be expressed explicitly. In hybrid (micro- / macro-) versions of the codes we propose, the “inner” integrator could be the microscopic model of the problem, while the “outer” extrapolation procedure would be a more conventional one. We will see later that the existence of a perfect integrator is not a requirement for the method, nor would its existence be of any great value compared to existing numerical methods.

The extrapolation can be viewed as an *outer integrator*. In this view, the steps of the inner integrator serve to damp the fast components and develop the numerical solution of the slow components over a small interval. That (numerical) solution is then numerically differentiated to get the derivatives needed for the outer integrator. Performing numerical differentiation avoids some of the error amplification inherent in an evaluation of $f(y)$ due to the large eigenvalues of the Jacobian $\partial f/\partial y$.

When the outer integrator is simply an extrapolation from the small-step smooth solution, it is essentially an explicit method. We will see that the extrapolation can also be viewed as the predictor to be followed by the outer integrator analog of an implicit method. This implicit method *can* generally be solved by a predictor-corrector functional iteration, even at the large outer step size.

The reader might think that these should be called “extrapolation methods”, but that name has already been used for methods which estimate the error terms by extrapolation [1] to $h = 0$, an unrelated class of methods. Hence we call the proposed methods *Projective Integration Methods*.

These methods have some similarities with methods in the literature. We will mention a few recent papers containing ample references to earlier ones. Many of the related literature methods generate the *dominant subspace* of the Jacobian (the subspace spanned by the eigenvectors corresponding to the large eigenvalues) and then solve for the fast components in that (usually lower-dimensional) subspace. They often form the subspace from the Krylov sequence (see [4], [5], and [2]), or they may use Chebychev polynomials to get a lower dimensional approximation to the solution operator in the fast subspace ([7]).

A related set of papers ([6], [8], and [10]) addresses the issue of greatly extending the region of stability along the negative real axis by using a set of steps of varying sizes so that the region is suitable for parabolic equations. The set of step sizes determine the polynomial approximation to the exponential operator and are chosen to make this approximation less than one in magnitude in the desired regions. Reference [6] indicates how the regions can take a number of shapes. Such methods could be used to generate the shape of stability region we address. However, the methods we will use are quite different and rely only on the basic stability properties of the explicit integrators we will use, not on a special collection of step sizes.

We assume that the stability of the differential system (1) and the numerical method can be approximated by studying the stability of a local linearization, and that the eigenvalues of the Jacobian $J = \partial f / \partial y$ are clustered into two groups G_0 and G_1 . We assume that there is a large separation, or gap, (relative to the size of the clusters) between the clusters (see Figure 1), that G_0 is in the neighborhood of the origin, corresponding to the slow components and that G_1 is well into the left-half plane. (If there are driving terms, the time constants of the driving terms should be viewed as if they were additional eigenvalues in the G_0 cluster in this discussion since we will be concerned with methods that are accurate for the components in G_0 .)

In the methods proposed below we are happy to use conventional integrators in the region where the fast components are active since a small step size is needed in this region. When the problem becomes *stiff*, i.e. when the fast components become negligible at the accuracy desired, the conventional inner integrator damps the fast components, and thus avoids any direct representation of the dominant subspace. Taking advantage of this we can achieve an explicit outer integrator that is stable even at large step sizes. The slightly surprising result is that the stability region for the outer integrator is essentially the same as that for the related conventional explicit integrator for just the slow components.

In the next section we will consider the simplest realization of the proposed method and analyze its stability. Then we will examine a number of extensions of the method and show that the behavior of the simplest method is not a special case, but typical of many possible realizations.

Finally we will apply the method to some examples.

2 The Projective Forward Euler Method

One outer integration step of the *Projective Forward Euler* (PFE) method integrates over $k + 1 + M$ steps of size h from t_n to $t_{n+k+1+M}$ in the following manner:

1. Using a suitable *inner integrator* integrate for k steps from t_n to t_{n+k} . It does not matter in our discussion what method is used for this inner integrator except that it is stable and of at least first order. (We will also assume that it is linear as described below.)
2. Perform one more inner integration to compute y_{n+k+1} from y_{n+k} .
3. Now perform an extrapolation over M steps using y_{n+k+1} and y_{n+k} to compute $y_{n+k+1+M}$ as

$$y_{n+k+1+M} = (M + 1)y_{n+k+1} - My_{n+k} \quad (2)$$

We call this outer procedure a “Forward Euler” method because we could also write it as

$$y_{n+k+1+M} = y_{n+k+1} + (Mh)y'_{n+k+1}$$

where the derivative approximation is given by

$$y'_{n+k+1} = (y_{n+k+1} - y_{n+k})/h$$

The k steps of the inner integrator must damp the fast components sufficiently to offset the growth of the same components in the extrapolation step. In fact, each application of the inner integrator multiplicatively reduces the fast components, so the error reduction scales with a power of k whereas the growth in the extrapolation is linear in M . This will be seen in the simple analysis below.

We will assume that the inner integrator is a *linear* method. This is defined to be a method that commutes with a linear transformation of the dependent variables in the equations. That is, applying the numerical method to the differential equation (1) to get y_n and then computing $z_n = Qy_n$ for some nonsingular, constant transformation Q gives the same results (within roundoff error) as applying the method to the transformed differential equation $z' = Qf(Q^{-1}z)$. Most numerical integration methods are linear. When a linear method is applied to the linear, constant coefficient equation $y' = Ay$ it is equivalent to applying it separately to the set of scalar equations $y' = \lambda y$ where the λ are the eigenvalues of A . Hence, we can do linear stability analysis of linear methods by analyzing their effect on the scalar *test* equation

$$y' = \lambda y \quad (3)$$

for each of the eigenvalues, λ , of the problem.

One step of the inner integrator applied to eq. (3) over step size h starting from $y = y_n$ will give

$$y_{n+1} = \rho(h\lambda)y_n$$

where ρ is the *amplification* of the method. For the perfect integrator, $\rho(h\lambda) = \exp(h\lambda)$.

We now consider the error propagation of PFE for a linear equation. Suppose that the error at t_n in an eigencomponent corresponding to eigenvalue λ is ϵ_n . (We must

consider the behavior for each λ in G_0 and in G_1 .) After k inner integrations steps, the error is “amplified” (normally, that will be a decrease) to

$$\epsilon_{n+k} = \rho^k \epsilon_n.$$

(We write ρ to mean $\rho(h\lambda)$.) When the extrapolation (2) from step k and $k + 1$ to $k + 1 + M$ is performed, that error will be amplified to

$$\epsilon_{n+k+1+M} = (M + 1)\epsilon_{n+k+1} - M\epsilon_{n+k} = [(M + 1)\rho - M]\rho^k \epsilon_n$$

Hence, the error amplification in the compound step, which we will denote by $\sigma(h\lambda)$, is given by

$$\epsilon_{n+k+1+M} = \sigma(h\lambda)\epsilon_n$$

where

$$\sigma(h\lambda) = [(M + 1)\rho - M]\rho^k \tag{4}$$

The method is *absolutely stable* if $|\sigma(h\lambda)| \leq 1$. Absolute stability depends on the value of $h\lambda$. The *region of absolute stability* (hereafter called the stability region) in the $h\lambda$ -plane is the set of $h\lambda$ for which Eq. (4) gives $|\sigma(h\lambda)| \leq 1$. We can find this region by plotting the locus of all $h\lambda$ for which $|\sigma| = 1$. This locus will divide the $h\lambda$ -plane into two or more regions. By continuity, if any point in a region is stable, then all are.

This region depends on the form of the inner integrator. If the inner integrator is perfect, we get the stability regions shown in Figures 2, 3, and 4 for PFE methods with $k = 2$ and $M = 5, 7$ and 9 . We refer to this class of methods as Pk - M methods. Note that only the strip for imaginary values in $[-i\pi, +i\pi]$ is shown. This strip is repeated at a spacing of $2i\pi$. We see that, in the first two figures, the plane is divided into two regions. Since the point at $-\infty$ gives $\rho = 0$ for the perfect integrator, the region to the left of the boundary is the stability region. In the Figure 4 the stability region has split into multiple parts - the interior of the closed contour on the right-side of the graph which is repeated every $2i\pi$, and the semi-infinite part on the left. Now there is a gap in the $h\lambda$ values on the negative real axis for which the method is stable.

It is impossible to find even an approximation to a perfect integrator for large values of $h\lambda$ except for trivial problems, so the apparent infinite stability region in this $h\lambda$ -plane is somewhat misleading.

We can, instead, plot these stability regions for any specific inner integrator. For example, figures 5, 6, and 7 show the stability regions for the same three methods (P2-5, P2-7, and P2-9) if the inner integrator is the FE method. (For reasons which will be apparent shortly, the regions are displaced a distance one to the right so that the origin in the $h\lambda$ -plane appears at the point (0,1).) Notice now that the regions (which are the interior(s) of the curves shown) are finite because the FE method used as the inner integrator is unstable for large $h\lambda$.

The stability regions in the $h\lambda$ -plane are determined by two factors - the form of the outer integrator and the form of the inner integrator. In analyzing specific methods it will be necessary to consider the impact of the inner integrator and its effect on the stability regions. Fortunately we can consider the impact of the outer integrator independently of that of the inner integrator by considering the stability in the ρ -plane using eq. (4). It is the set of values of (inner) ρ for which the (outer) $\sigma \leq 1$. Because FE gives $\rho(h\lambda) = 1 + h\lambda$ it is simply a shift by 1 of the stability plot of the method using FE as the inner integrator, which we have shown in Figures 5 to 7. To find the (overall) stability region for any specific inner integrator it is sufficient to map back from

the ρ -plane to the $h\lambda$ -plane using the particular form of $\rho(h\lambda)$ for the inner integrator. Thus Figure 2 is simply the logarithm of Figure 5.

We notice that as M gets larger, the stability region breaks into two. It is easy to show that no matter what the value of k , for sufficiently large M the stability region will separate in this way as it has in Figure 7. In fact, the regions split when M is around $3.6k$ for large k . (Asymptotically for large k , the split occurs when $M > \gamma k$ where γ is the real root of $\gamma = \exp(1 + 1/\gamma)$. See Section 4.)

However, as long as we have a gap in the spectrum and can arrange for cluster G_1 to be in the left-hand region and cluster G_0 to be in the right-hand region, the method will be stable. (In fact, we can not expect all eigenvalues in cluster G_0 to be inside the right-hand region since there can be small positive eigenvalues corresponding to slowly growing components. All that matters is that the method is accurate for those components.)

M will be chosen so that the “effective step size” of the outer integrator, Mh , is commensurate with the slow components, that is, $Mh = \Delta/\lambda_0$ where λ_0 is an eigenvalue in the slow cluster G_0 and Δ is chosen for error control, typically in the range $0.001 < \Delta < 0.2$. If there is a large gap in the spectrum then M can be large (and considerable savings in integration achieved).

We show in Section 4 that, however large M , we can always choose k so that the method is stable, provided that the step size of the inner integrator can be chosen so that $\rho < 1$ for all eigenvalues in G_1 . The result follows from the lemma in Section 4 that says that as M becomes large the stability regions of the PFE method in the ρ -plane approach two disks, one centered at the origin with radius $1/M^{1/k}$ and the other centered at $1 - 1/M$ with radius $1/M$. Figure 7 illustrates this, although M is still relatively small and thus the regions are far from circular (this is an asymptotic result).

The effective outer integrator step size is Mh . Hence it is interesting to consider stability in the $Mh\lambda$ -plane and ask where the eigenvalues in G_0 must lie. This map depends on the form of the inner integrator. If the inner integrator is the FE method and we map the second (right-hand) disk into the $Mh\lambda$ plane we get a circle centered at -1 and radius 1. This is precisely the stability region for FE at step size Mh .

Interestingly, this result is also true asymptotically for *any* inner integrator that is at least first order. This occurs because first order implies that $\rho(h\lambda) = 1 + h\lambda + O(h\lambda)^2$. If $Mh\lambda$ is fixed as M becomes large, this implies $\rho = 1 + (Mh\lambda)/M + (Mh\lambda)^2 O(M^{-2})$. Hence, the map of the stability region from the ρ -plane to the $Mh\lambda$ -plane in any fixed regions of the $Mh\lambda$ -plane asymptotically approaches the stability region of the FE method.

The radius of the first (left-hand) disk shrinks as M increases. We must choose a k large enough to shrink the ρ^k values corresponding to the eigenvalues in cluster G_1 into this disk. As long as the maximum ρ corresponding to any of these eigenvalues is less than one, this is possible and we have

$$k \geq k_1 = -\log(M)/\log(\rho_{max}) \quad (5)$$

where ρ_{max} is the maximum value of $|\rho(h\lambda)|$ for eigenvalues λ in the set G_1 .

2.1 Efficiency of the PFE Method

We define the efficiency improvement of the method as

$$E = M/(k + 1) \quad (6)$$

because, with $k + 1$ inner integrations we are actually advancing an additional distance of M steps. Thus, if E has a value of 0 there is no advantage, while if E were 9 we would have a speed up of $E + 1$, or 10.

If we know something about the inner integrator, it is possible to estimate E based on the location of the eigenvalue clusters. We will assume that G_1 is a cluster centered on real negative λ_1 with radius $\lambda_r = D|\lambda_1|$ as shown in Figure 1. Here, D is the relative radius of the cluster G_1 . If we use FE for the inner integrator with

$$h = -1/\lambda_1 \quad (7)$$

we can see that

$$|\rho(h\lambda)| \leq h\lambda_r = \lambda_r/|\lambda_1| < 1$$

Note that λ_1 and λ_r are values representing our imperfect *knowledge* of the eigenvalue locations and are used to estimate the maximum value of ρ . This estimate is calculated from these assuming both constant eigenvalues and linearity. Neither of these will, in general, be true. Hence in practice λ_r must be taken as somewhat larger than would be given from a direct eigenvalue calculation.

Let us suppose that the *time constant* of the slow components is given by $1/\lambda_0$. λ_0 could be the largest (in magnitude) of the eigenvalues in cluster G_0 , or it could represent the speed of driving components that have to be tracked accurately. The outer integrator must use a step commensurate with $1/|\lambda_0|$, but for accuracy it will have to be smaller than that, say $\Delta/|\lambda_0|$, where Δ is dependent on the integration accuracy required. (Typical values of Δ for 2 to 4 digits of accuracy are 0.2 to 0.01 for FE.) Hence we expect the outer “step size” to be given by

$$Mh = \Delta/|\lambda_0|$$

or, using (7)

$$M = \Delta|\lambda_1/\lambda_0| \quad (8)$$

Substituting this in eq. (5) we have an approximate value for k given by

$$k = -\frac{\log(\Delta|\frac{\lambda_1}{\lambda_0}|)}{\log(|\frac{\lambda_r}{\lambda_1}|)}$$

Since the time scale is arbitrary, it is convenient to normalize and use the descriptors $S = |\lambda_1/\lambda_0|$ and $D = |\lambda_r/\lambda_1|$. Here, S is the “stiffness,” defined as the ratio of the typical fast component (that is, center of the eigenvalue set) to the slow component of interest, while D is the “dispersion” of the stiff eigenvalues - that is, the extent to which they are spread out. A dispersion of one means that there is no gap and the proposed method is of no value.

With these parameters we can express the efficiency improvement as

$$E = S\Delta \frac{\log(D)}{\log D + \log(S\Delta)} \quad (9)$$

Note that from eq. (8) $S\Delta = M$. Eq. (9) shows that the speed improvement depends on the accuracy required - through Δ ; the relative speeds of the fast and slow components - through S ; and the amount to which the fast components are dispersed - through D . The values of E as a function of the relative radius, D , for several values of $M = S\Delta$ are shown in Figure 8.

2.2 The Inner Integrator

Generally we are not concerned with the exact nature of the inner integrator, only that it is at least first-order accurate for the slow components and damping for the fast components. If we use an explicit method for the inner integrator, $\rho(h\lambda)$ will be a polynomial in $h\lambda$. For example, if the forward Euler method is used, ρ is the linear polynomial $1 + h\lambda$ so that ρ^k is simply $(1 + h\lambda)^k$ if a fixed step size is used for the inner integrator. By using the approaches of [6], [8] and [10] and varying the step size over the first k inner steps, say by using the sequence $h_j = \alpha_j h, j = 1, \dots, k$, we can achieve an amplification factor over the first k inner steps of

$$\rho_k(\mu) = \prod_j (1 + \alpha_j \mu)$$

where $\mu = h\lambda$. The α_j can be chosen to give this polynomial other desirable properties - for example the minimax polynomial over the cluster G_1 . (If all of the eigenvalues in G_1 are negative real, the minimax polynomial is approximated by the Chebyshev polynomial of degree k over the line segment containing the eigenvalues.)

3 Higher-order and Implicit Outer Methods

We are not usually content with a first order integration method so we naturally ask about higher-order methods. Here we are referring to the “outer integrator” since the inner integrator method can be chosen independently of the outer one. What is important for the inner integrator is its ability to damp the fast components quickly. Since its step size is small compared to the outer step, accuracy is almost certainly assured for the slow components in the inner steps.

The Pk - M Method described above is an example of an *explicit method*. Since the stability regions of implicit methods are superior to those of explicit methods, we naturally ask about the existence of “implicit” outer methods. We put “implicit” in quotes because we are deliberately trying to avoid solving large implicit systems that involve a Jacobian. We will show that they can be solved by a simple predictor-corrector iteration using the “power of the power.”

3.1 Higher-order Explicit Methods

In the Pk - M Method, the inner integrator computed two important pieces of information from the value y_n for the outer integrator: approximations to y_{n+k} and to y'_{n+k} . The approximation to y_{n+k} provides damping of the fast components by k applications of the inner integrator while the approximation to the derivative was used in the outer integration formula. The obvious extension of this first-order method is to ask the inner integrator to provide approximations to additional derivatives at t_{n+k} and use these in a Taylor’s series to compute y_{n+k+M} . Each additional derivative will require one more inner integration step.

If we do q additional inner integrations steps after the initial k such steps, we can estimate the first q derivatives at t_{n+k} . An alternative view is that we do a q -th order extrapolation forward a distance Mh from the values $y_{n+k}, y_{n+k+1}, \dots, y_{n+k+q}$ to form $y_{n+k+q+M}$. We will call this a Pk - q - M method (so that the previous method is a Pk -1- M

method). This method has properties similar to the Pk - M Method: for small M there is a single stability region (see Figures 9 and 10) and for large M its stability region consists of two regions in the ρ -plane, a region around the origin in the ρ -plane that is approximately a disk of radius

$$r = \left[\frac{q!}{M^q} \right]^{1/k} \quad (10)$$

and a region containing $\rho = 1$ on its boundary that approximately maps to the stability region of the q -th order Taylor's Series method in the $Mh\lambda$ -plane. This region is the set of μ such that

$$|1 + \mu + \mu^2/2! + \mu^3/3! + \cdots + \mu^q/q!| \leq 1$$

in the μ -plane. See Figure 11.

Note from eq. (10) that we now need to choose a k such that

$$\rho_{max}^k < \frac{q!}{M^q} \quad (11)$$

or

$$k \geq k_q = -q \log(M) / \log(\rho_{max}) + \log(q!) / \log(\rho_{max})$$

Comparing this with eq. (5) we see that

$$k_1 > \frac{k_2}{2} > \frac{k_3}{3} > \cdots > \frac{k_q}{q}$$

so that the number of inner integration steps, $k + q$, increases more slowly than q , the order.

3.2 Implicit Methods

The Pk - q - M is closely related to q -th order Taylor Series Methods. The implicit extension of a Taylor's Series Method uses derivatives from both ends of the interval in the approximation. The simplest of these is the *trapezoidal rule* (called the *Crank-Nicolson method* in the PDE community). It uses the first derivative from each end. The general form of the implicit Taylor's Series Method is

$$y_{n+1} = y_n + h\beta_{01}y'_n + \cdots + h^{q_0}\beta_{0q_0}y_n^{(q_0)} + h\beta_{11}y'_{n+1} + \cdots + h^{q_1}\beta_{1q_1}y_{n+1}^{(q_1)} \quad (12)$$

Viewing the inner integrator as a device that can approximate derivatives we can implement a projective analog of this method which we will call the Pk - q_0, q_1 - M method. It uses q_0 additional steps following the calculation of y_{n+k} to estimate the first q_0 derivatives at t_{n+k+q_0} ; it also uses a similar estimate of the first q_1 derivatives at t_{n+k+q_0+M+k} . The latter is accomplished by using y_{n+k+q_0+M} to first integrate to t_{n+k+q_0+M+k} and then integrating a further q_1 inner steps and using the last $q_1 + 1$ values to estimate the q_1 derivatives). It then uses these derivative estimates in the analog of eq. (12).

We will look at the Pk -1-1- M method in a little more detail since it illustrates the general issues. To reduce the length of subscripts, we will write N for $n + k + 1 + M$. The implicit "corrector" formula is

$$y_N = y_{n+k+1} + \alpha M(y_{n+k+1} - y_{n+k}) + (1 - \alpha)M(y_{N+k+1} - y_{N+k}) \quad (13)$$

This method is of first order for any α and second order for one specific choice of α . The method is implemented with the predictor-corrector iteration given below, where the word “integrate” means perform one inner integration step.

1. Integrate from t_n for $k + 1$ steps.
2. Predict $y_N^{(0)}$ using eq. (2).
3. Integrate for $k + 1$ steps from $y_N^{(m)}$ to form $y_{N+i}^{(m)}$ with $m = 0$.
4. Correct y_N using:

$$y_N^{(m+1)} = y_{n+k+1} + \alpha M(y_{n+k+1} - y_{n+k}) + (1 - \alpha)M(y_{N+k+1}^{(m)} - y_{N+k}^{(m)}) \quad (14)$$

where α is a parameter that can be chosen to vary the stability regions and accuracy.

5. Repeat steps 3 and 4 for $m = 1, 2, \dots$ until the corrector has converged.

As long as steps 3 and 4 are iterated until the iteration error is small, it is sufficient to study the stability of the corrector formula (13) only.

We are using the equivalent of a functional iteration to solve the implicit corrector equation (13). In conventional methods, functional iteration cannot be used when the problem is stiff (i.e., $h\lambda$ is large). However, in this case, the k inner integration steps save the day. Writing $\delta^{(m)}$ for the difference $y^{(m)} - y$ we get by subtracting eq. (13) from eq. (14):

$$\delta_N^{(m+1)} = (1 - \alpha)M(\delta_{N+k+1}^{(m)} - \delta_{N+k}^{(m)}) \quad (15)$$

For the linear, constant coefficient case we can treat each eigencomponent separately and we have

$$\delta_{N+k+i}^{(m)} = \rho^{k+i} \delta_N^{(m)}.$$

Hence eq. (15) becomes

$$\delta_N^{(m+1)} = (1 - \alpha)M\rho^k(\rho - 1)\delta_N^{(m)} \quad (16)$$

This converges if

$$|(1 - \alpha)M\rho^k(\rho - 1)| < 1$$

which, for large M , means that either

$$|\rho| < \left| \frac{1}{(1 - \alpha)M} \right|^{1/k} + O(1/M)^{2/k} \quad (17)$$

or

$$\rho = 1 + \eta \quad (18)$$

where

$$|\eta| < \left| \frac{1}{(1 - \alpha)M} \right| + O\left(\frac{1}{M^2}\right)$$

Criterion (17) applies to the large eigenvalues. For these, ρ must be inside the unit circle and k must be chosen sufficiently large to satisfy (17). Criterion (18) applies to the

slow components with small eigenvalues. If the inner integrator is at least first order, $\eta = h\lambda + O(h\lambda)^2$. Hence criterion (18) requires that

$$|(1 - \alpha)Mh\lambda| < 1 - O(h\lambda)^2 \quad (19)$$

Note that once again, the gap requirement plays a role. Eigenvalues must either be in cluster G_1 where the inner method is such that ρ is inside a disk of radius of order $1/M^{1/k}$ or in cluster G_0 close to the origin where criterion (19) applies to $Mh\lambda$. Mh is, of course, the effective outer integrator step size.

In the related one-step method

$$y_{n+1} = y_n + h(\alpha y'_n + (1 - \alpha)y'_{n+1})$$

$\alpha = 1$ corresponds to the FE method, $\alpha = 0$ corresponds to the backward Euler method, while $\alpha = 0.5$ yields the second-order trapezoidal rule. In the Pk-1-1- M method $\alpha = 1$ corresponds to the PFE method, while with $\alpha = 0$ we have a method akin to the Backward Euler Method. Choosing $\alpha = 0.5$ looks like the trapezoidal method, but it is not second order. Second order is achieved with a slightly different value of α because the two derivative estimates (formed by differences) are only second order accurate at the midpoints of their difference intervals, namely $t_{n+k+1/2}$ and $t_{N+k+1/2}$ respectively. These have to be combined to get a second order accurate estimate of the derivative at the midpoint of the outer integration interval, namely $t_{n+k+1+M/2}$. This leads to

$$\alpha_2 = \frac{M + 2k + 1}{2(M + k + 1)}$$

as the value which yields second-order accuracy. (As M becomes large, this becomes 0.5.)

The stability regions for P2-1-1- M methods are illustrated in Figures 12 to 15. In Figure 12 we show the stability region for $M = 5$ and $\alpha = 0$ (or the Projective Backward Euler Method). The method is stable everywhere except in the interior of the three disk-like regions. The rightmost region corresponds to the region of instability of the standard Backward Euler Method (which is a circle centered at 1 and radius 1 in the $Mh\lambda$ -plane). The leftmost two regions are approximately circular, and asymptotically for large M they become disks centered at the k -th roots of $-1/M$ and radius $1/(kM^{(k+2)/k})$. (The reason we have used such a small M for Figure 12 is that for M of any practical size, the leftmost regions are so small that they can not be seen when plotted on the same scale as needed for the rightmost region. The next few terms in the asymptotic expansion are significant in the estimate of the centers for such a small M - the next two terms change the center of the upper left disk by $-0.1 - 0.06i$.)

As α increases towards the value for which the method is second order, the size of the right-hand region of instability spreads further into the right half plane, and the small disks of instability on the left increase. Figure 13 shows the P2-1-1-100 Method (note the much larger M) for $\alpha = 0.49$. Again, the region of stability is everywhere except the interior of the regions plotted. As α increases a little more, the right-hand region grows very rapidly, then briefly becomes a large region in the left-half plane enclosing the other regions. The stability region is now the interior of the region whose boundary passes through $\rho = 1$ with the exception of the interior disks, which by now have become distorted. It is shown in Figure 14 for $M = 5$. As α continues to increase, the larger

region, now in the left half plane, shrinks and its boundaries coalesce with those of the disk of instability. Figure 15 shows the regions just coalescing in the case $M = 5$. Figure 16 shows the P2-1-1-100 method for the α that gives a second order method. Note that the boundary comes within about $1/M^{1/2}$ of the origin, corresponding to the location of the centers of the disks of instability. As α increases further towards 1, the stability region shrinks, pinches off, and becomes as described for PFE. (See Figure 7.)

It appears at first sight as if the second-order method whose stability is shown in Figure 16 should be very good for problems with all real eigenvalues since its stability region includes the whole of the interval $[0,1]$ on the ρ -axis. (The inner integrator could presumably be chosen to keep ρ real and positive in that case, e.g., by choosing h in FE to be the inverse of the most negative eigenvalue.) However, this is an implicit method, and we saw in eq. (15) that the condition for predictor-correct convergence also restricts ρ . If one has to resort to solving the implicit corrector by some form of Newton iteration, nothing has been gained by the proposed method.

The main point to be noted from this discussion is that there will be instabilities about $1/M^{1/k}$ from the origin. Even if one has very good knowledge of the location of the eigenvalues (and the values of $\rho(h\lambda)$ for those eigenvalues) *convergence of the functional corrector iteration will still restrict the application of the methods to problems with gaps in their spectrum.*

4 Analysis of General Methods

This section will sketch proofs of results stated in previous sections and discuss some of the methods in more generality. It can be skipped by readers not interested in the general theory.

4.1 Stability Region Gap

In Section (2) we noted that in the Pk - M method, an M larger than about $3.6k$ would cause the stability region to break into two pieces. This can be seen by considering eq. (4) and asking when the locus of ρ can cross the real axis as σ traverses the unit circle. Clearly, there can only be a real root ρ of eq. (4) when σ is real, namely, $+1$ or -1 . Graphically it is easy to see that when $\sigma = +1$ there is a real root at $\rho = +1$, and one more real root at $\rho \approx -1/M^{1/k}$ (if k is odd) or no other real roots (if k is even). When $\sigma = -1$ we can rewrite eq. (4) as

$$1/\rho^k = M - (M + 1)\rho \tag{20}$$

The graph of the left-hand side is independent of M and composed of two pieces, one piece convex downwards in the positive half plane and a second piece in the negative half plane that is convex upwards and negative if k is odd or is convex downwards and positive if k is even. As M increases from zero, the graph of the right-hand side, which is a straight line passing through $(1,-1)$ and $(0,M)$ slowly tilts up until it intersects the graph of the left-hand side in the right-half plane at $M = M_{crit}$. For positive $M < M_{crit}$ there are no positive real roots for ρ . For $M > M_{crit}$ there are two positive real roots. There is one negative real root for any positive M if k is even, none otherwise. Hence the real axis in the ρ -plane is cut in only two places by the boundary of the stability region

for $M < M_{crit}$ as shown in Figures 5 and 6. For $M > M_{crit}$ the stability region boundary intersects the real axis in four places and we have the situation shown in Figure 7. The critical value of M is at the point that eq. (20) has repeated real roots.

We can compute this value by asking that the right-hand side of (20) be tangential to the left-hand side, or that eq. (20) and

$$k/\rho^{k+1} = M + 1 \quad (21)$$

be satisfied simultaneously. Eliminating ρ we get

$$\left[\frac{(M+1)}{M} \right]^k = \frac{M}{k+1} \left[\frac{k}{k+1} \right]^k$$

Setting $M = \gamma k$ we get

$$\left[1 + \frac{1}{\gamma k} \right]^k = \gamma \left[1 - \frac{1}{k+1} \right]^{k+1}$$

As $k \rightarrow \infty$ then we see that γ tends to the solution of

$$\gamma = \exp(1 + 1/\gamma)$$

The behavior of the stability regions as M gets large is given by the following lemma:

Lemma 1 For large M the roots $\rho_j, j = 0, \dots, k$ of Eq (4) when $\sigma = \exp(i\theta), \theta \in (0, 2\pi)$ are given by

$$\rho_0 = 1 - 1/M + \exp(i\theta)/M + O(1/M^2)$$

and

$$\rho_j = \exp(i(\theta/k + 2\pi j/k))/M^{1/k} + O(1/M^{2/k}), j = 1, \dots, k$$

The proof follows by simple asymptotic expansions.

4.2 Stability Regions for Higher-order Explicit Methods

From the q -th order extrapolation formula using forward differences, the error amplification for the Pk - q - M method can be seen to be

$$\sigma = [1 + (M+q)(\rho-1) + \dots + (M+q)(M+q-1)\dots(M+1)(\rho-1)^q/q!]\rho^k \quad (22)$$

The stability region is the set of ρ for which σ in the unit disk, and can be found once again by forming the locus of the values of ρ when $\sigma = \exp(i\theta), \theta \in [0, 2\pi)$. The regions are shown for $q = 4, k = 4$ and $M = 6, 8,$ and 10 in Figures 9 through 11. Once again we note that as M increases, the stability region “pinches off” and breaks into two regions.

As before, we can analyze the behavior asymptotically in large M . By series expansion we can demonstrate that the $k+q$ roots of eq. (22) are:

$$\rho_i = \left[\frac{\sigma q!}{M^q} \right]^{1/k} + O(M^{-(q+1)/k}), i = 1, \dots, k \quad (23)$$

(the k -th root yields k different values), and

$$\rho_j = 1 + \gamma_j/M + O(M^{-2}), j = k + 1, \dots, k + q \quad (24)$$

where the γ_j are the q roots of

$$\sigma = 1 + \gamma + \gamma^2/2! + \dots + \gamma^q/q!. \quad (25)$$

The locus of γ is precisely the boundary of the stability region of the q -th order Taylor Series Method, so that the scaling of γ by $1/M$ in eq. (24) maps the usual stability plot into the scaled version in the ρ plane, as stated earlier. Figures 17 and 18 show the two sections of the stability region for the P4-4-100 Method. Figure 17 shows the right-hand section that passes through $\rho = 1$ mapped into the $(\rho - 1)M$ -plane, corresponding approximately to the $Mh\lambda$ -plane for the outer integrator at step size Mh . Comparing this with the stability region for the classical fourth-order explicit Runge-Kutta Method (see, for example, [3] Section 2.6.1) we see that they are very close. As M tends to infinity they become the same. Figure 18 shows the left-hand section around $\rho = 0$ scaled by $M/(4!)^{1/4}$ (see eq. (23)). As can be seen, it is close to the unit disk.

4.3 Implicit Methods

The general form of the implicit Taylor's series method is given in eq. (12). The analogous projective method, the P k - q_0 - q_1 - M method will be

$$y_N = [1 + h\beta_{01}\nabla + \dots + h^{q_0}\beta_{0q_0}\nabla^{q_0}]y_{n+k+q_0} + [h\beta_{11}\Delta + \dots + h^{q_1}\beta_{1q_1}\Delta^{q_1}]y_{N+k} \quad (26)$$

where Δ is the forward difference operator ($\Delta y_n = y_{n+1} - y_n$) and ∇ is the backward difference operator. These provide approximations to the derivatives. (We use the backward difference operator on the left hand end to emphasize that we will project forward from the furthest forward point computed in the inner integration, namely y_{n+k+q_0}). The coefficients β_{ij} are chosen to achieve the desired order (the order is defined to be the maximum integer such that the method is exact for all problems whose solution is a polynomial of degree no greater than the order).

The coefficients β_{ij} have to be chosen to achieve some order of accuracy or other behavior. The maximum order of accuracy that can be achieved is $q_0 + q_1$. If less is used (as, for example, in the P- k -1-1- M method discussed earlier) then the degrees of freedom can be used to try to achieve other properties. Given the large number of possibilities, we will say little about the general methods, but will analyze one particular case - $q_0 = 0$ which is what might be called the Backward Taylor's Series (BTS) method. If the β_{1j} are chosen to achieve order $q = q_1$ then the coefficients are uniquely determined as functions of k and M .

We can construct the coefficients as follows:

The q -th degree polynomial that passes through y_{N+k+i} for $i = 0, \dots, q$ can be written as

$$p(t) = [1 + (t - t_{N+k})\Delta + \dots + \frac{(t - t_{N+k}) \dots (t - t_{N+k+q-1})}{q!} \Delta^q] y_{N+k}$$

(We have taken the step size h to be 1 in this equation to reduce the number of characters in the equation.) This polynomial passes through y_{N+k} . However, we wanted to construct

a polynomial that passes through y_{n+k+q} and has the same differences as $p(t)$ at the points t_{N+k+i} . (We are projecting forward from y_{n+k+q} because we assume that in the previous outer step we computed $k+q$ steps forward from y_n in order to perform the previous BTS step.) Hence we add the constant $y_{n+k+q} - p(t_{n+k+q})$ to $p(t)$ to get the desired “interpolating polynomial.” Then we compute the approximation to y_N from it, getting

$$y_N = y_{n+k+q} + p(t_N) - p(t_{n+k+q}) \quad (27)$$

(It may seem strange that the “interpolating polynomial” does not actually interpolate at the points t_{N+k+i} . We could achieve that by using a polynomial of one degree higher. However, if we do that, we appear to lose convergence of the predictor-corrector iteration in many cases.)

We can study the stability of BTS for large M by noting that if we keep the dominant terms in M in each difference in eq. (27) we get:

$$y_N = y_{n+k+q} - \sum_{i=1}^q (-M)^i \Delta^i / i! y_{N+k}$$

For the linear test equation with an inner integrator having an amplification ρ we find that the amplification from y_n to y_N is given by

$$\sigma = \rho^{k+q} / [1 + \rho^k \sum_{i=1}^q [M(1-\rho)]^i / i!] \quad (28)$$

We can analyze the stability region of this for large M as before by considering the locus of ρ corresponding to σ on the unit circle. In the region of $\rho = 1$ we consider $\rho = 1 + \gamma/M$ where $\gamma = h\lambda$. We find that the region is bounded by γ such that

$$1 - \gamma + \gamma^2/2 \cdots + (-\gamma)^q/q! = \exp(i\theta)$$

Comparing this with eq. (25) we see that it bounds a region that is the mirror image in the imaginary axis of the stability region for the explicit q -th order method. The remaining k roots of eq. (28) when $\sigma = \exp(-i\theta)$ can be shown to be small, nearly circular regions centered at the k -th roots of $-1/(-M)^q$ by asymptotic analysis. The BTS method is stable outside of these regions.

5 Examples

Example 1: The Brusselator with Rapidly Replenished Source

The simple Brusselator models a chemically reacting system with two time-varying concentrations, X and Y , two “source materials”, and two final products. The differential equations are:

$$\begin{aligned} X' &= A - (B+1)X + X^2Y \\ Y' &= BX - X^2Y \end{aligned}$$

Table 1: Results at $t = 10$ for Example 1 with $k = 4$

M	X	Y	B
10	0.48766	2.7234	2.9999
20	0.48794	2.7217	2.9999
40	0.48851	2.7181	2.9999
80	0.48970	2.7108	2.9999
160	0.49220	2.6960	2.9999
320	0.49777	2.6659	2.9999
640	0.51098	2.6037	2.9998
1280	0.55843	2.4536	2.9998
2560	0.48792	4.4590	2.9999

where the concentrations, A and B , of the source materials are usually assumed to be constant (and thus parameters). One might expect the A and B concentrations to be depleted locally as they are used in the reaction, although in the presence of a large supply reservoir, they will be replenished by diffusion and/or convection. We will add to this model a term for the depletion of the concentration B , followed by a replenishment at an exponential rate as might be expected with diffusion from a reservoir. The additional equation is

$$B' = (B_0 - B)/\epsilon - BX$$

Thus, the concentration B is reduced through its reaction with X but restored to its "base" level B_0 with a time constant of ϵ . This introduces a fast term into the reaction, or stiffness.

For values $A = 1, B_0 = 3$ the system has an unstable equilibrium at $X = A, Y = B_0/A, B = B_0$ but from any other starting point it tends to a limit cycle in which X ranges from about 0 to 4 and Y ranges from about 1 to 5. (See [9] for example.)

The system was integrated over the interval $[0,10]$ using the P- k -1- M method starting from $X = A + 0.1, Y = B_0/A + 0.1, B = B_0$. ϵ was 10^{-4} . The inner integrator was the FE method with $h = \epsilon$. For $k = 4$ we got the results at $t = 10$ shown in Table 1. In order to reach the end point exactly, the final outer step was modified as follows: If $k + 1$ inner steps of size h would pass the end point, the inner step size was reduced so that $k + 1$ steps got there exactly and no outer step was taken. Otherwise, M for the last outer step was reduced to that needed to reach the end point exactly. In a production code, one would interpolate to hit desired output points. (All calculations were done in Matlab on a Pentium III.)

The last line in Table 1 with $M = 2560$ corresponds to an outer integrator step size of 0.256 which is much too large for a first order method in a problem with active eigenvalues of the size of this problem. (At $X = 0.49, Y = 2.7, B = 3$ the eigenvalues of the Jacobian are $-0.7677 \pm 4.8395i$ and -1004.9)

We estimate the actual errors from these results as follows. Assume that the error in each step of the FE inner integrator is Ch^2 , and in each step of the outer integrator is $C(Mh)^2$. (The same C is used in both error estimates. C is also assumed to be constant

Table 2: Error coefficients for various M with $k = 4$

M	10	20	40	80	160	320	640	1280
$\frac{M^2+5}{M+5}$	7.0	16.2	35.7	75.4	155.2	315.1	635.0	1275.0

over the range of integration. Both assumptions are correct for problems whose solution is a quadratic and allow us to get a rough estimate of the error dependence on M .)

Over an integration interval of size L we make $L/[h(M+k+1)]$ outer steps and $(k+1)$ times as many inner steps. Thus the total error - if it is simply the sum of the errors in each step - is

$$Err = C \frac{L(k+1)}{h(M+k+1)} h^2 + C \frac{L}{h(M+k+1)} (Mh)^2 = ChL \frac{M^2+k+1}{M+k+1} \quad (29)$$

Thus, for large M we will see the linear error growth we expect from a first order method, but not for smaller M . In our example, $k = 4$. The errors should be proportional to $(M^2+5)/(M+5)$. The values of this function for the values of M used in the integration (10 to 1280) are shown in Table 2.

The errors in Table 1 will be approximately in these ratios. If we assume that the first two rows of Table 1 satisfy these ratios, the differences between its first two rows will be $(16.2 - 7.0)/7.0$ of the errors in row 1. From this assumption we can estimate the error in row 1 as 0.000208 in X and -0.001318 in Y . (We are showing more digits here because these values were computed from stored numbers held to greater precision than those printed in the tables.) Hence the “true” solution at $t = 10$ can be estimated as $X = 0.48745, Y = 2.7247$. (Integration using a version of LSODE with tolerance of 10^{-10} gives the answers $X = 0.487424, Y = 2.724937, B = 2.999854$.) From these we can estimate the errors in X and Y in Table 1 to be as shown in Table 3. In that table, the final two columns show the estimated errors divided by the error coefficients shown in Table 2. If the analysis were exact, all entries in a column would be the same. The nearly equal values indicated that the analysis is a good approximation until M becomes quite large.

For $k = 1$ we get the very similar results shown in Table 4.

A similar analysis verifies that these results have the expected error behavior from eq. (29).

We see that $k = 1$ is adequate for this problem. This occurs because (a) there is a single large eigenvalue so we can chose an h for the FE inner integrator to make $\rho = 0$, (b) the eigenvector corresponding to the large eigenvalue (the *dominant* eigenvector) does not vary much, and (c) the problem is only mildly non-linear in its fast components. Real-world problems are unlikely to be this cooperative. To illustrate the impact of a non-zero ρ we ran the problem with an inner step size of $h = \epsilon/2$. This means that ρ for the large eigenvalue is about $1/2$ (actually $1 - 1004.9/2000$ at the end of the interval). For large M we must have $\rho^k < 1/M$ according to the theory. To test this, we integrated the problem using $Pk-1-M$ with $M = 320, 640, 1280$, and 2560 with various k to determine the first k for which the results were stable (which was evaluated to be the first k for which the computed X and Y were not NaNs (Not-a-Number’s). The results are shown in Table 5.

Table 3: Errors in X and Y for Example 1

M	X error	Y error	scaled X error	scaled Y error
10	0.000208	-0.001318	0.000030	-0.000188
20	0.000481	-0.003050	0.000030	-0.000188
40	0.001055	-0.006650	0.000030	-0.000186
80	0.002248	-0.013967	0.000030	-0.000185
160	0.004744	-0.028749	0.000031	-0.000185
320	0.010322	-0.058820	0.000033	-0.000187
640	0.023528	-0.121017	0.000037	-0.000191
1280	0.070972	-0.271102	0.000056	-0.000213

Table 4: Results for Example 1 with $k = 1$

M	X	Y	B
10	0.48772	2.7231	2.9999
20	0.48800	2.7213	2.9999
40	0.48859	2.7176	2.9999
80	0.48979	2.7102	2.9999
160	0.49231	2.6954	2.9999
320	0.49789	2.6653	2.9999
640	0.51139	2.6030	2.9998
1280	0.55357	2.4604	2.9998

Table 5: Minimum k for stability with various M when $\rho_{max} \approx 0.5$

M	320	640	1280	2560
k_{min}	8	9	10	10

Table 6: Theoretical minimum $k = k_1$ for stability with various M when $\rho_{max} = 0.5$

M	320	640	1280	2560
k_{min}	8.32	9.32	10.03	11.32

Table 7: Results for Example 1 with $k = 4$ and $\epsilon = 10^{-6}$

M	X	Y
100	8.5777e-8	-5.4912e-6
200	4.1974e-8	-5.3997e-6
400	2.3872e-9	-5.2624e-6
800	2.1634e-7	-5.6259e-6
1600	1.4048e-7	-4.5435e-6
3200	2.7148e-6	-7.9766e-6
6400	1.4444e-5	-2.5224e-5
12800	5.8682e-5	-8.6436e-5
25600	2.5100e-4	-3.6538e-4
51200	9.8082e-4	-1.2701e-3

(The numerical result for $M = 2560$ and $k = 10$ was grossly in error - the final value of B was 7.1. With $k = 11$ better results were obtained, although $M = 2560$ corresponds to such a large outer step size that accuracy can hardly be expected.)

The values of k_1 given by eq. (5) for these values of M are shown in Table 6.

We used the second order projective method on the same problem with $\epsilon = 0.000001$ and $k = 4$. The smaller value of ϵ was used so that the inner step (which was equal to ϵ) was small enough that the error from the inner steps did not dominate the error from the outer steps. The second order behavior can be observed for a short range of M as shown in Table 7. For smaller M the errors from the inner integrator appear to dominate. (The final value of B was 3.0 to the precision printed in all cases, so its error is not shown. The other errors are based on the final value computed by LSODE as $X = 0.48739228, Y = 2.725322$.)

Example 2: An Index-reduced Differential Algebraic Equation

We will consider a pendulum described by ODEs derived via the Euler-Lagrange equations with constraints. This is one of the simplest set of Euler-Lagrange systems. For a general Hamiltonian $H(p, q)$ subject to holonomic constraints $C(q) = 0$ we have

$$q' = H_p$$

$$p' = -H_q - \lambda C_q$$

$$C(q) = 0$$

This is an index-3 differential-algebraic equation (DAE). Alternatively, we can replace the last equation with

$$\epsilon^2 \lambda C_q^T C_q = C(q) + 2\epsilon C_q^T H_p$$

to get an easily-solved index-1 DAE. Under some conditions with a small ϵ there is a “boundary layer” (whose width is order ϵ) after which the solution is an order ϵ perturbation of the solution of the original problem. The eigenvalues corresponding to the fast components are approximately $-1/\epsilon$ so h can be chosen to be ϵ in the inner integrator. (Note that this technique is proposed only to give an illustration of the algorithm. It is not a good way to solve such problems over long time periods because it is equivalent to replacing the constraint with a “stiff” near-constraint that forces the solution back onto the manifold of the constraint very rapidly (rate $1/\epsilon$) and with damping so that the corresponding eigenvalues are real. However, the damping is absorbing small amounts of energy so the system is no longer Hamiltonian, and energy is slowly dissipated.)

The pendulum problem - normalized with unit length rod and unit mass - is given by $H(p, q) = (u^2 + v^2)/2 + y$ and $C(q) = x^2 + y^2 - 1$ where $q = (x, y)$ and $p = (u, v)$. Thus the equations of the modified system are

$$x' = u; y' = v$$

$$u' = -2\lambda x; v' = -1 - 2\lambda y$$

with

$$\lambda = \frac{(x^2 + y^2) - 1 + 4\epsilon_1(xu + yv)}{4\epsilon^2(x^2 + y^2)} \quad (30)$$

with $\epsilon_1 = \epsilon$.

This was integrated using the Pk - q - M method for several values of the parameters k , q , and M , and also several values of ϵ . The initial conditions at $t = 0$ were $x = 0$, $y = -1$, $u = 2$, and $v = 0$. The analytic solution for $\epsilon = 0$ is $x = \sin(\theta)$, $y = \cos(\theta)$ where $\theta(t) = 4 \tan^{-1}(\exp(-t))$. Hence at $t = t_{end} = -\log(\tan(\pi/8))$ y should be 0.

The system has four eigenvalues. They are, asymptotically in ϵ , $\pm i \cos(\theta)$ and two values of $-1/\epsilon$. (These equal values can be separated by changing ϵ_1 in eq. (30) to $\gamma\epsilon$ with γ close to one. Values of γ larger than one separate the values on the real axis, values less than one separate them into the complex plane. The effect is similar to that of γ in $[\epsilon\lambda]^2 + 2\gamma[\epsilon\lambda] + 1 = 0$.)

The inner step size was taken to be ϵ in an FE method and a range of M was chosen dependent on ϵ .

Tables 8 to 10 show the values of y at $t = t_{end}$ and $k = 3, 4, 5$, and 6. Note that in Table 8 the outer steps range from 10^{-3} to 0.512. Since the active eigenvalues have magnitude 1.0 at the start of the interval, we should not expect much accuracy for the larger values of M in this table. In Tables 9 and 10 the outer step size ranges from 10^{-4}

Table 8: Error for Example 2 with Pk-1-M, $\epsilon = 10^{-3}$

M	k			
	3	4	5	6
1	0.004329	0.004121	0.003973	0.003860
2	0.006189	0.005742	0.005406	0.005146
4	0.010851	0.009976	0.009274	0.008701
8	0.021797	0.020340	0.019065	0.017970
16	0.044852	0.043936	0.041960	0.040192
32	0.056660	0.095381	0.092707	0.089971
64	-0.402159	0.206593	0.202625	0.198990
128	failure	0.453308	0.407696	0.421408
256	failure	failure	-0.373756	0.638664
512	failure	failure	-1.900817	0.245407

to 0.0512. The error for different k and ϵ for the same outer step sizes is comparable when small k does not lead to failures.

The “failure” entries in the lower left of the table indicate that the errors grew so large that the solution is, if not actually unstable, meaningless. Even though the linear theory indicates that $k = 1$ should be sufficient in this case, non linearities or time dependency of the differential equation can require a larger k . In this case, the eigenspace corresponding to the two large eigenvalues is changing with time. The projective outer step amplifies an error in this eigenspace as it projects forward in time, but at the new time, some of this error is now in the eigenspace of the small eigenvalues, and the numerical method does not damp these rapidly. Hence more inner steps are needed to damp them.

6 Conclusion

We have shown that it is possible to integrate stiff equations with “explicit” methods provided that there is a gap in the spectrum. By “explicit” we mean that no use is made of any approximation to the Jacobian or spaces representative of some of the eigenvectors. The size of the gap will determine how efficient these methods are, compared to the underlying explicit integrator. Whether or not they are competitive with existing methods will depend very strongly on how difficult it is to handle the stiff components by more conventional methods, and so in effect it will depend on the size and structure of the Jacobian. It will, of course, also depend on how effectively the proposed methods can be implemented in automatic codes, and we are not yet ready to report on that.

Table 9: Error for Example 2 with $Pk-1-M$, $\epsilon = 10^{-4}$

M	k			
	3	4	5	6
1	0.000434	0.000413	0.000399	0.000388
2	0.000620	0.000576	0.000543	0.000517
4	0.001086	0.000999	0.000930	0.000874
8	0.002172	0.002029	0.001906	0.001799
16	0.004536	0.004335	0.004150	0.003983
32	0.009415	0.009205	0.008967	0.008742
64	0.018221	0.019207	0.018932	0.018665
128	0.000363	0.039682	0.039371	0.039066
256	-0.417386	0.082059	0.081695	0.081346
512	failure	0.172150	0.170969	0.170939

Table 10: Error for Example 2 with $Pk-1-M$, $\epsilon = 10^{-5}$

M	k			
	3	4	5	6
10	0.000275	0.000258	0.000244	0.000232
20	0.000574	0.000552	0.000532	0.000514
40	0.001188	0.001162	0.001137	0.001114
80	0.002423	0.002399	0.002371	0.002344
160	0.004795	0.004887	0.004858	0.004829
320	0.006003	0.009892	0.009861	0.009830
640	-0.084920	0.019992	0.019960	0.019930
1280	-0.628337	0.040548	0.040514	0.040481
2560	failure	0.083090	0.082933	0.082959
5120	failure	0.216302	0.137220	0.172793

References

- [1] Burlisch, R. and Stoer, J., Numerical Treatment of Ordinary Differential Equations by Extrapolation Methods, *Num. Math.*, **8**, pp 1-13, 1966.
- [2] Edwards, W. S., Tuckerman, L. S., Friesner, R. A., and Sorensen, D. C., Krylov Methods for the Incompressible Navier-Stokes Equation, *JCP*, **110** pp 82-104, 1994.
- [3] Gear, C. W., Numerical Initial Value Problems in Ordinary Differential Equations, Prentice Hall, Englewood Cliffs, NJ, 1971.
- [4] Hochbruck, M. and Lubich, C.,H, On Krylov Subspace Approximations to the Matrix Exponential Operator, *SIAM J. Numer. Anal.*, **34** #5, pp 1911-1925, Oct 1997.
- [5] Hochbruck, M., Lubich, C., and Selhofer, H, Exponential Integrators for Large Systems of Differential Equations, *SIAM J. Sci. Computing*, **19**, #5, pp 1552-1574, Sept 1998.
- [6] Lebedev, V. I., Explicit Difference Schemes with Variable Time Steps for Solving Stiff Systems of Equations, in *Numerical Analysis and Its Applications*, Proc. WNAA '96, Bulgaria, ed. Vulkov, Waśniewski, and Yalamov, Springer, Berlin, 1997, pp274 - 283.
- [7] Loh, Y. L., Taraskin, S. N, and Elliott, S. R., Fast Time-Evolution Method for Dynamical Systems, *Physical Review Letters* **84** #11, pp 2290 - 2293, 13 Mar 2000.
- [8] Medovikov, A. L., High-Order Explicit Methods for Parabolic Equations, *BIT* **38**, #2, pp372 - 390, 1998.
- [9] Thompson J. M. T. and Stewart, H. B., Nonlinear Dynamics and Chaos, John Wiley, NY, 1986 pp56-60.
- [10] Verwer, J. G., Explicit Runge-Kutta Methods for Parabolic Partial Differential Equations, *Applied Numerical Mathematics* **22**, pp359-379, 1996.




Figure 1: Regions G_0 and G_1




Figure 2: Complex $h\lambda$ -plane stability for P2-5 Method with perfect inner integrator $[-i\pi, i\pi]$ strip only




Figure 3: Complex $h\lambda$ -plane stability for P2-7 Method with perfect inner integrator $[-i\pi, i\pi]$ strip only




Figure 4: Complex $h\lambda$ -plane stability for P2-9 Method with perfect inner integrator $[-i\pi, i\pi]$ strip only




Figure 5: Complex ρ -plane stability for P2-5 Method




Figure 6: Complex ρ -plane stability for P2-7 Method




Figure 7: Complex ρ -plane stability for P2-9 Method




Figure 8: Speed Improvement for $M = 5, 10, 15, 20, 25,$ and 30 (in increasing order from bottom).




Figure 9: ρ -plane stability for P4-4-6 Method




Figure 10: ρ -plane stability for P4-4-8 Method




Figure 11: ρ -plane stability for P4-4-10 Method




Figure 12: ρ -plane stability for P2-1-1-5 Method with $\alpha = 0$




Figure 13: ρ -plane stability for P2-1-1-100 Method with $\alpha = 0.49$




Figure 14: ρ -plane stability for P2-1-1-5 Method with $\alpha = 0.412$




Figure 15: ρ -plane stability for P2-1-1-5 Method with $\alpha = 0.413$




Figure 16: ρ -plane stability for P2-1-1-100 Method with α chosen for second order accuracy (0.5097)




Figure 17: $Mh\lambda \approx M(\rho - 1)$ -plane stability for P4-4-100 Method Section near $\rho = 1$ scaled by 100 (M).




Figure 18: ρ -plane stability for P4-4-100 Method Section near $\rho = 0$ scaled up by $M/(4!)^{1/4}$