

PROJECTIVE METHODS FOR STIFF DIFFERENTIAL EQUATIONS: PROBLEMS WITH GAPS IN THEIR EIGENVALUE SPECTRUM*

C. W. GEAR[†] AND IOANNIS G. KEVREKIDIS[‡]

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 nonlinear 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.

Key words. integration, stiff, explicit, stability

AMS subject classifications. 65L05, 65L12

PII. S1064827501388157

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 nonautonomous 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. The existence of such spectral gaps is germane to the theory of inertial manifolds for dissipative PDEs [12, 2] and of approximate inertial forms for their discretizations.

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. 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 other 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” postprocessing through the “outer” integrator, as discussed in [6].

Once beyond the region where the fast components are active, we would like to

*Received by the editors April 17, 2001; accepted for publication (in revised form) May 20, 2002; published electronically January 31, 2003.

<http://www.siam.org/journals/sisc/24-4/38815.html>

[†]NEC Research Institute, Princeton, NJ 08540 (cwg@research.nj.nec.com).

[‡]Department of Chemical Engineering, Princeton University, Princeton, NJ 08544 (yannis@arnold.princeton.edu). This author’s research was supported in part by AFOSR (Dynamics and Control program, Dr. Mark Jacobs).

use step sizes in a numerical integration commensurate with the slow components. Conventional wisdom (see, e.g., Gear [4, Chapter 11]) tells us that we have to use implicit methods in this case, and then 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

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

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 legacy code that implements an explicit integration and it is not practical to reengineer 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 ODEs, the primary application is undoubtedly to problems arising from the semidiscretization of PDEs.)

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). The interesting feature of the method is that the linear stability region is (roughly) composed of the stability regions of the two (inner and outer) methods so that it can accommodate problems with two clusters of eigenvalues. (Of course, a linear stability analysis does not guarantee stability for nonlinear/strongly time-dependent problems, but it is seldom a good idea to use a method that is not stable for linear problems! Further analysis is needed to understand the method's behavior on nonlinear problems, but standard analysis will show that mild nonlinearities or time-dependencies will not destroy the stability if the linearized eigenvalues lie sufficiently inside the linear stability region.)

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 processed to get the derivatives needed for the outer integrator. This processing avoids some of the error amplification inherent in an evaluation of $f(y)$ due to the large eigenvalues of the Jacobian $\partial f/\partial y$.

The reader might think that these should be called "extrapolation methods," but that name has already been used [1] for methods which estimate the error terms by extrapolation 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 [7], [8], and [3]), or they may use Chebyshev polynomials to get a lower-dimensional approximation to the solution operator in the fast subspace [10].

FIG. 1. Regions G_0 and G_1 .

A related set of papers [9, 11, 14] 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 determines the polynomial approximation to the exponential operator and is chosen to make this approximation less than one in magnitude in the desired regions. Reference [9] indicates how the regions can take a number of shapes. Such methods could be used to generate the shape of the stability region that we address. However, the methods we will use are quite different and rely on only 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 is 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

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

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

$$y'_{n+k+1} = \frac{y_{n+k+1} - y_{n+k}}{h}.$$

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

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

for each of the eigenvalues λ of the problem.

One step of the inner integrator applied to (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, $\rho(h\lambda) = \exp(h\lambda)$, but this is not achievable unless the problem can be integrated explicitly.

We now consider the error propagation of PFE for a linear equation. Suppose that the error at t_n in an eigenvector 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 integration 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

$$(4) \quad \sigma(h\lambda) = [(M+1)\rho - M]\rho^k.$$

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 (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 Figure 2 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, for $M = 3$ and 5 , 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. When $M = 9$, the stability region has split into multiple parts—the interior of the closed contour on the right-hand 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, and thus 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, Figure 3 shows the stability regions for the same three methods (P2-5, P2-7, and P2-9) if the inner integrator is the forward Euler (FE) method. (For reasons which will be apparent shortly, the regions are displaced by distance of 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




FIG. 2. Complex $h\lambda$ -plane stability for the P2-M method with perfect inner integrator for $M = 5$ (dashes), 7 (dots), and 9 (solid lines); $[-i\pi, i\pi]$ strip only.

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 (4). The stability region 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 plots of the method using FE as the inner integrator, which we have shown in Figure 3. 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 3.

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 3 when $M = 9$. 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.

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 can be 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




FIG. 3. Complex ρ -plane stability for the P2-M method with $M = 5$ (dashes), 7 (dots), and 9 (solid lines).

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 3 with $M = 9$ illustrates this, although M is still relatively small there 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 having 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 its being 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

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

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

3. Higher-order 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.

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 series to compute y_{n+k+M} . Each additional derivative will require one more inner integration step.

If we do q additional inner integration 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, 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

$$(6) \quad r = \left[\frac{q!}{M^q} \right]^{1/k}$$

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

$$\left| 1 + \mu + \frac{\mu^2}{2!} + \frac{\mu^3}{3!} + \dots + \frac{\mu^q}{q!} \right| \leq 1$$

in the μ -plane. Figure 4 shows the stability regions of the $P4$ -4- M methods with $M = 6, 8$, and 10 . In the case of $M = 6$ and 8 there is a single stability region. In the case of $M = 10$ the region has split into two. For comparison purposes, the stability region of the four-stage, fourth-order Runge-Kutta method is also plotted, scaled by 0.1 (corresponding to $1/M$). Note that it is similar to the rightmost of the regions of the $P4$ -4-10 method, as predicted (although, since M is not large, it still differs significantly).

Note from (6) that we now need to choose a k such that

$$(7) \quad \rho_{max}^k < \frac{q!}{M^q}$$

or

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

Comparing this with (5), we see that

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




FIG. 4. ρ -plane stability for the $P4\text{-}4\text{-}M$ method with $M = 6$ (dashes), 8 (dots), and 10 (solid lines), as well as the scaled $RK4$ stability region (dot-dashed line).

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

3.1. Implicit methods. The $Pk\text{-}q\text{-}M$ method is closely related to q th-order Taylor series methods. The implicit extension of a Taylor 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 series method is

$$(8) \quad y_{n+1} = y_n + h\beta_{01}y'_n + \cdots + h^{q_0}\beta_{0q_0}y_n^{(q_0)} + h\beta_{11}y'_{n+1} + \cdots + h^{q_1}\beta_{1q_1}y_{n+1}^{(q_1)}.$$

Viewing the inner integrator as a device that can approximate derivatives, we can implement a projective analogue of the implicit method. It is described in more detail in the report on which this paper is based [5]. It might seem that an implicit method would require the implicit solution of nonlinear equations at each (outer) step—precisely what we are trying to avoid. It turns out that a “predictor-corrector” iteration can be used that will converge, provided that the eigenvalues are within the two clusters. We have not yet conducted experiments to determine whether the additional accuracy of a predictor-corrector method justifies the additional computation (which it often does in standard integrators).

4. Analysis of general methods. This section will sketch the 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 (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 (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 (4) as

$$(9) \quad \frac{1}{\rho^k} = M - (M+1)\rho.$$

The graph of the left-hand side is independent of M and composed of two pieces, one piece in the positive half-plane that is convex downwards 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, and 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 Figure 3 for $M = 5$ and 7 . For $M > M_{crit}$ the stability region boundary intersects the real axis in four places, and we have the situation shown in Figure 3 with $M = 9$. The critical value of M is at the point where (9) has repeated real roots.

We can compute this value by asking that the right-hand side of (9) be tangential to the left-hand side, or that (9) and

$$(10) \quad \frac{k}{\rho^{k+1}} = M+1$$

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

$$\gamma = \exp \left(1 + \frac{1}{\gamma} \right).$$

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

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

$$\rho_0 = 1 - \frac{1}{M} + \frac{\exp(i\theta)}{M} + O\left(\frac{1}{M^2}\right)$$

and

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

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

$$(11) \quad \sigma = \left[1 + (M+q)(\rho-1) + \dots + (M+q)(M+q-1) \dots \frac{(M+1)(\rho-1)^q}{q!} \right] \rho^k.$$

The stability region is the set of ρ for which σ is 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 in Figure 4 for $q = 4$, $k = 4$, and $M = 6, 8$, and 10 . 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 (11) are

$$(12) \quad \rho_i = \left[\frac{\sigma q!}{M^q} \right]^{1/k} + O(M^{-(q+1)/k}), \quad i = 1, \dots, k,$$

(the k th root yields k different values) and

$$(13) \quad \rho_j = 1 + \frac{\gamma_j}{M} + O(M^{-2}), \quad j = k+1, \dots, k+q,$$

where the γ_j are the q roots of

$$(14) \quad \sigma = 1 + \gamma + \frac{\gamma^2}{2!} + \dots + \frac{\gamma^q}{q!}.$$

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 (13) maps the usual stability plot into the scaled version in the ρ -plane, as stated earlier.

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

$$X' = A - (B+1)X + X^2Y,$$

$$Y' = BX - X^2Y,$$

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

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

a replenishment at an exponential rate as might be expected with diffusion from a reservoir. The additional equation is

$$B' = \frac{B_0 - B}{\epsilon} - BX.$$

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, or stiffness, into the reaction.

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 [13, pp. 56–60], for example.)

The system was integrated over the interval $[0, 10]$ starting from $X = A + 0.1$, $Y = B_0/A + 0.1$, $B = B_0$. In the first two tests the inner integrator was the FE method with $h = \epsilon = 10^{-4}$ and $k = 1$ and 4. (All calculations were done in Matlab on a Pentium III.)

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. 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 found in 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 .)

The behavior of the columns in Table 1 may not appear to represent a first-order method. This is because there is a significant influence of the errors from the inner method whose step size is not changing. An analysis in [5] shows that we are seeing a combination of first-order error change in the outer integrator plus constant error from the inner integrator.

The results for $k = 1$ for the above are similar. This is to be expected because (a) there is a single large eigenvalue, and thus we can choose 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 nonlinear in its fast components. Real-world problems are unlikely to be this cooperative. To illustrate the impact of a nonzero 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

TABLE 2
Minimum k for stability with various M when $\rho_{\max} \approx 0.5$.

M	320	640	1280	2560
k_{comp}	8	9	10	10
k_1	8.32	9.32	10.03	11.32

TABLE 3
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

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 $Pk\text{-}1\text{-}M$ 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 2 as k_{comp} . The k_1 shown on the third line in this table is the minimum predicted by (5).

(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.)

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 from $M = 6400$ to 51200, as shown in Table 3. 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 Euler–Lagrange equations with constraints. This is one of the simplest sets of Euler–Lagrange systems. We present this example because the eigenspace of the stiff components is slowly rotating (rotating with the pendulum, in fact), and this is an example of a time-dependent behavior that requires a larger value of k . 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

$$x' = u, \quad y' = v,$$

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

with

$$(15) \quad \lambda = \frac{(x^2 + y^2) - 1 + 4\epsilon_1(xu + yv)}{4\epsilon^2(x^2 + y^2)}$$

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 for 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 (15) 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 4 to 6 show the values of y at $t = t_{end}$ and $k = 3, 4, 5$, and 6. Note that in Table 4 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 5 and 6 the outer step size ranges from 10^{-4} 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, nonlinearities

TABLE 4
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

TABLE 5
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 6
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

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 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 with 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.

REFERENCES

- [1] R. BURLISCH AND J. STOER, *Numerical treatment of ordinary differential equations by extrapolation methods*, Numer. Math., 8 (1966), pp. 1–13.
- [2] P. CONSTANTIN, C. FOIAS, B. NICOLAENKO, AND R. TEMAM, *Integral Manifolds and Inertial Manifolds for Dissipative Partial Differential Equations*, Springer-Verlag, New York, 1988.
- [3] W. S. EDWARDS, L. S. TUCKERMAN, R. A. FRIESNER, AND D. C. SORESENSEN, *Krylov methods for the incompressible Navier–Stokes equation*, J. Comput. Phys., 110 (1994), pp. 82–104.
- [4] C. W. GEAR, *Numerical Initial Value Problems in Ordinary Differential Equations*, Prentice–Hall, Englewood Cliffs, NJ, 1971.
- [5] C. W. GEAR AND I. G. KEVREKIDIS, *Projective Methods for Stiff Differential Equations: Problems with Gaps in Their Eigenvalue Spectrum*, NEC Research Institute Report 2001-029, Princeton, NJ; also available online from <http://www.neci.nj.nec.com/homepages/cwg/projective.pdf>.
- [6] C. W. GEAR, I. G. KEVREKIDIS, AND C. THEODOROPOULOS, “Coarse” integration/bifurcation analysis via microscopic simulators: Micro-Galerkin methods, Computers and Chemical Engineering, 26 (2002), pp. 941–963.
- [7] M. HOCHBRUCK AND C. LUBICH, *On Krylov subspace approximations to the matrix exponential operator*, SIAM J. Numer. Anal., 34 (1997), pp. 1911–1925.
- [8] M. HOCHBRUCK, C. LUBICH, AND H. SELHOFER, *Exponential integrators for large systems of differential equations*, SIAM J. Sci. Comput., 19 (1998), pp. 1552–1574.
- [9] V. I. LEBEDEV, *Explicit difference schemes with variable time steps for solving stiff systems of equations*, in Numerical Analysis and Its Applications, Proceedings of the Workshop on Numerical Analysis and Its Applications (WNAA ’96), Rousse, Bulgaria, L. Vulkov, J. Waśniewski, and P. Yalamov, eds., Springer, Berlin, 1997, pp. 274–283.
- [10] Y. L. LOH, S. N. TARASKIN, AND S. R. ELLIOTT, *Fast time-evolution method for dynamical systems*, Phys. Rev. Lett., 84 (2000), pp. 2290–2293.
- [11] A. L. MEDOVIKOV, *High-order explicit methods for parabolic equations*, BIT, 38 (1998), pp. 372–390.
- [12] R. TEMAM, *Infinite Dimensional Dynamical Systems in Mechanics and Physics*, Springer-Verlag, New York, 1988.
- [13] J. M. T. THOMPSON AND H. B. STEWART, *Nonlinear Dynamics and Chaos*, John Wiley, NY, 1986.
- [14] J. G. VERWER, *Explicit Runge–Kutta methods for parabolic partial differential equations*, Appl. Numer. Math., 22 (1996), pp. 359–379.