Padé approximant algorithm for solving nonlinear ordinary differential equation boundary value problems on an unbounded domain

John P. Boyd[a)]
Department of Atmospheric, Oceanic and Space Science, University of Michigan, 2455 Hayward Avenue, Ann Arbor, Michigan 48109

(Received 11 October 1996; accepted 7 March 1997)

We describe a four-step algorithm for solving ordinary differential equation nonlinear boundary-value problems on infinite or semi-infinite intervals. The first step is to compute high-order Taylor series expansions using an algebraic manipulation language such as Maple or Mathematica. These expansions will contain one or more unknown parameters z which will be determined by the boundary condition at infinity. The second step is to convert the Taylor expansions into diagonal Padé approximants. The boundary condition that $u(x)$ decays to zero at infinity becomes the condition that the coefficient of the highest power of x in the numerator polynomial must be zero. The third step is to solve this equation for the free parameter z. The final step is to evaluate each of the multiple solutions of this equation for physical plausibility and convergence (as N increases). This algorithm can be implemented in as few as seven lines of Maple (sample program provided!). We illustrate the method with three examples: the Flierl–Petviashvili vortex of geophysical fluid dynamics, the quartic oscillator of quantum mechanics, and the Blasius function for the boundary layer above a semi-infinite plate in fluid mechanics. Methods for nonlinear problems are almost always iterative and need a first guess to initialize the iteration. The Padé algorithm is unusual in that it is a direct method that requires no *a priori* information about the solution. [S0894-1866(97)01703-3]

INTRODUCTION: THE HIDDEN VIRTUES OF TAYLOR SERIES

Power series are a ubiquitous topic of low-level college mathematics courses. This popularity obscures the fact that for real applications, Taylor series *per se* are almost useless. The real power of power series is in combination with other ideas. Newton's iteration (power series plus iteration of the linear approximation) and finite differences (multiple Taylor expansions, each on its own subinterval and linked) are familiar examples.

Power series in isolation are almost never useful to solve boundary-value problems because it is only occasionally that the radius of convergence is sufficiently large to embrace both boundaries. Nevertheless, when combined with another idea, Padé approximants, power series become an effective tool for ordinary differential equation boundary-value problems.

The $[p/q]$ Padé approximant to a function $f(x)$ is a polynomial of degree p divided by a polynomial of degree q which is chosen so that the leading terms of the power series of the approximant match the first $(p+q+1)$ terms of the power series of $f(x)$. One might suppose that the approximant would be restricted to the same domain of convergence as the Taylor expansion from whence it came. In reality, the Padé approximant will usually converge on the entire real axis if $f(x)$ is free of singularities on the real axis.[1–3]

By matching powers in the equation $Q(x)f_T(x) = P(x)+O(x^{p+q+1})$ where f_T is the power series of $f(x)$ and $P(x)$ and $Q(x)$ are the numerator and denominator polynomials of the Padé approximant, one can derive a set of linear equations to determine the coefficients of the approximant. However, recurrence relations such as the "q-epsilon" algorithm do the same job even more efficiently than solving a matrix problem. Because Padé-finding algorithms are built-in utilities in most algebraic manipulation languages such as Maple and Mathematica and are also available in most Fortran and C software libraries, we shall assume that we have access to such utilities. The theory and practice of Padé approximants are described in many texts.[1–3]

One useful piece of folklore is that the most accurate approximant is usually the *diagonal* approximant, that is, from all combinations of p and q such that $p+q=N+1$, it is usually best to choose $p=q$. We shall form only diagonal approximants in what follows.

For two of our three examples, the Padé approximant in fact appears to converge over the entire physical domain. The Blasius function is tougher because the diagonal approximant does not have the correct asymptotic behavior as $x\rightarrow\infty$. It does, however, accurately describe the function $f(x)$ for larger and larger but finite x as the order of the

a)E-mail: jpboyd@engin.umich.edu

approximant increases. We can therefore apply our algorithm to this case, too, by applying the boundary condition for large but finite x, and then comparing results for different x and different $p=q$, until we obtain convergence.

I. IMPLEMENTATION AND FIRST EXAMPLE

The best way to describe the algorithm is to work an example. The Flierl–Petviashvili (FP) monopole $u(r)$ is the radially symmetric solution to the equation

$$u_{rr}+\frac{1}{r}u_r-u-u^2=0, \tag{1}$$

where the subscript denotes differentiation with respect to the subscripted coordinate. Flierl[4] and Petviashvili[5] independently derived this equation as a model for vortex solitons in the ocean and in plasmas, respectively. Because the solution is a function of radius only, it must be a function of r^2 rather than r and its first derivative at the origin must be zero (see Chap. 15 of Ref. 6, which shows that these statements must be true of the radially symmetric component of any well-behaved function when represented in polar coordinates). This implies that the only free parameter in the power series is

$$z\equiv u(0). \tag{2}$$

Maple's built-in utility to solve a differential equation in a power series gives

$$\begin{aligned}u&=z+\frac{z+z^2}{4}r^2+\frac{z+3z^2+2z^3}{64}r^4\\&+\frac{z+9z^2+16z^3+8z^4}{2304}r^6\\&+\frac{z+29z^2+106z^3+130z^4+52z^5}{147456}r^8+O(r^{10}).\end{aligned} \tag{3}$$

The second step is to convert this to the equivalent Padé approximant—a one-line command:

$$u_{[4/4]}(r)=\frac{p_0+p_2r^2+p_4r^4}{q_0+q_2r^2+q_4r^4}, \tag{4}$$

$$\begin{aligned}p_0&=20z+16z^2+16z^3,\\p_2&=\tfrac{17}{4}z+6z^2+\tfrac{7}{2}z^3+z^4,\\p_4&=\tfrac{79}{576}z+\tfrac{65}{288}z^2-\tfrac{35}{288}z^3-\tfrac{23}{72}z^4-\tfrac{7}{72}z^5,\\q_0&=20+16z+16z^2,\\q_2&=-\tfrac{3}{4}-3z-\tfrac{9}{2}z^2-3z^3,\\q_4&=\tfrac{7}{576}-\tfrac{7}{288}z+\tfrac{37}{288}z^2+\tfrac{11}{36}z^3+\tfrac{11}{72}z^4.\end{aligned} \tag{5}$$

The third step is to impose the boundary condition,

$$u(\infty)=0. \tag{6}$$

The limit of the Padé approximant as $r\rightarrow\infty$ is p_4/q_4. This will vanish, as demanded by the boundary condition, if and only if

$$p_4(z)=0. \tag{7}$$

Table I. Roots of the Padé approximants to the FP monopole in *z*.

Degree	Roots
[2/2]	-1.5
[4/4]	-2.50746, 0.84
[6/6]	-2390278, -2.02, -1.02, -0.92, -0.78, 1.44, 2.70
[8/8]	-2.392214, -2.396796, -3.46, -0.27, -0.017, 0.017
[10/10]	-2.3919746, -2.3969315, -2.42, -1.70, -1.12, -1.02, -1.0006, -0.99913, -0.9909, 1.36, 1.39
Exact	-2.3919564

Maple has a built-in command to solve polynomial equations.

The fourth step is to evaluate each of the roots. Two are complex valued, and therefore discarded. Since the Laplace operator, which in radial form is $u_{rr}+(1/r)u_r$, describes the curvature of u, it follows that when $u>0$, the solution has positive, upward curvature when it is already positive. Therefore, a solution with $u>0$ anywhere must blow up to ever-increasing positive values as $r\rightarrow\infty$. Thus, the positive root for $z=u(0)$ must be extraneous too.

The negative root approximates the true value of $u(0)$ to within about 4%. The Padé rational function approximates $u(r)$ with small absolute error for all $r\in[0,\infty]$. Table I illustrates the convergence of approximants with increasing order. The roots converge monotonically to the exact $z=u(0)$ with an error that decreases *exponentially* with N as illustrated by Fig. 1.

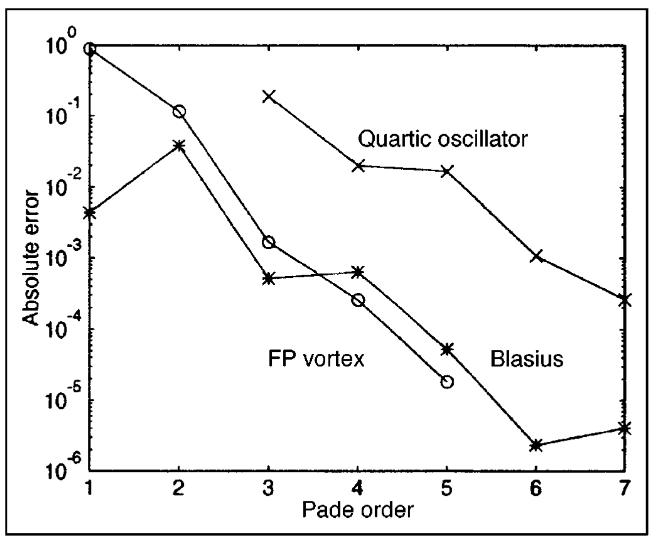

Figure 1. Absolute errors vs the order of the Padé approximant for each of the examples: FP vortex: circles; quartic oscillator: ×'s; Blasius function: ∗ The order is simply the row number in the tables for each case. (The Blasius errors are the result of imposing the boundary condition at x=8.) Note that on this log–linear graph, decay of the error as exp(−qN) for some constant q would appear as a straight line. Clearly, the error for all three examples can be bounded by a straight line of this form for some positive q, showing that the error is decaying exponentially with N.

```
# Step One: Compute power series
ODE:=diff(u(r),r$2) + (1/r)*diff(u(r),r)-u(r)-u(r)**2=0;
initvals:= u(0)=z, D(u)(0)=0;
Order:=9; # Must be of form 4*j+1;
solution:=dsolve(ODE,initvals, u(r),'series');
# Step Two: Convert to Padé Approximant
uPade:=convert(rhs(solution),'ratpoly');
# Step Three: Solve for z≡u(0).
highest_of_P:=coeff(collect(numer(uPade),r),r,(Order-1)/2);
zroots:=fsolve(highest_of_P,z,'complex');
```

Figure 2. Maple program to compute Padé approximants to the FP monopole.

There are two surprises. First, each Padé approximant has many roots in z even though the solution to the differential equation is unique. Second, the true zero is approximated for high order by a cluster of closely spaced roots. These are *generic* features of the Padé method that are complications for the next two examples as well. We shall return to the issues raised by multiple roots in Secs. IV and V.

Figure 2 is the complete Maple script for this calculation. Excluding comments, the program contains *just seven lines*.

II. SECOND EXAMPLE: QUANTUM QUARTIC OSCILLATOR

This is the linear differential equation eigenvalue problem

$$u_{xx}+(E-x^4)u=0, \tag{8}$$

where E is the eigenvalue. It is the stationary form of Schrödinger's equation with x^4 as the potential energy. In contrast to our other examples, which are posed on a semi-infinite interval, the quantum problem is defined for $x \in [-\infty,\infty]$. Only a slight variation from the algorithm of Sec. I is needed. Here, the eigenvalue E is the free parameter. Since the amplitude of the eigenfunction is arbitrary, we can set $u(0)=1$ without loss of generality. All the eigefunctions of the equation are either symmetric $[u(-x)=u(x)]$ or antisymmetric $[u(-x)=-u(x)]$ with respect to the origin.[6] For the even eigenfunctions, we can set $u_x(0)=0$ because the power series for an even function involves only even powers of x.

Applying the same algorithm as for the Flierl–Petviashvili monopole gives the results shown in Table II. The $[14/14]$ approximant gives the ground state eigenvalue to within 1 part in 5000.

In contrast to the Flierl–Petviashvili equation, which has only a single nontrivial solution that is bounded at infinity, the quartic oscillator has a countable infinity of even parity modes. The Padé method is also generating approximations, albeit much cruder than for the ground state, to the second symmetric eigenfunction. The $[14/14]$ approximation to the first symmetric excited state has a relative error of only 1.1%.

Table II. Roots of the Padé approximants to the anharmonic oscillator, lowest two even modes.

Degree	Roots	Absolute error	Roots	Error
[2/2]	0	⋯	⋯	⋯
[4/4]	0	⋯	⋯	⋯
[6/6]	1.248	0.18	⋯	⋯
[8/8]	1.04036	0.020	7.4017	0.054
[10/10]	1.0769	0.0165	7.0203	0.435
[12/12]	1.06144	0.00174	7.55627	0.101
[14/14]	1.060100	0.000262	7.3738	0.082
Exact	1.06036209		7.45569794	

III. THIRD EXAMPLE: BLASIUS FUNCTION

The boundary layer flow over a semi-infinite flat plate can be reduced to the ODE (Blasius),[8]

$$2\,f_{xxx}+ff_{xx}=0, \tag{9}$$

subject to the boundary conditions,

$$f(0)=f_x(0)=0, \quad f_x(\infty)=1. \tag{10}$$

One might suppose that one could impose the boundary condition at $x=\infty$ by differentiating the Padé approximant and then demanding that its difference from unity must vanish. Unfortunately, the power series for the Blasius function shows that it is of the form

$$f(x)=zx^2v(zx^3), \tag{11}$$

where the free parameter in the power series is

$$z\equiv f_{xx}(0). \tag{12}$$

This implies that the Padé approximants are x^2 times the ratio of two polynomials in x^3. This causes major difficulties because the diagonal Padé approximant for v will give an approximation that blows up as x^2 instead of x. The $[N/N+1]$ approximant will decay as $O(1/x)$, which also fails to match the correct asymptotic behavior of $f(x)$.

Nevertheless, for *intermediate* x, the Padé rational function will give a good approximation. If we compute z so that $f_x(x_j,z)=1$ for several large but finite x_j for each diagonal Padé approximant, then we can compare results. If some $z(x_j)$ agree, then we shall have presumably obtained an accurate approximation to the true z.

This apparent consistency is not a rigorous proof of small error, but we face the same dilemma in accepting the results of any numerical calculation once the numbers have apparently converged to a common value for large N where the grid spacing, Padé order or whatever is varied. At some point, one must make the subjective judgment that it is unlikely that a large set of numerical values will agree unless the calculation is accurate.

Table III shows that this strategy succeeds. The approximants converge rapidly for $x=6$, but to a value for $f_{xx}(0)$, which is too large by about 0.0005, a relative error of about 0.15%. The reason for this error is that the Blasius function has not quite asymptoted to its exact value at

Table III. Roots of the Padé approximants for $f_{xx}(0)$: Blasius equation. (Note: The highest power of x in the Taylor series together with the order of the diagonal approximant in x^3 are given in the leftmost column. The other columns are labeled by the values of x where the coefficient of the highest term in the numerator was required to vanish; the resulting roots, which approximate $f_{xx}(0)$, are given below.)

Order	$x=6$	$x=7$	$x=8$	$x=9$	$x=10$
14 [2/2]	0.3299	0.3292	0.3364	⋯	⋯
20 [3/3]	0.3347	0.3424	0.3699	0.5122	⋯
26 [4/4]	0.33257	0.33216	0.33257	0.3345	⋯
32 [5/5]	0.33257	0.33218	0.33269	0.3349	⋯
38 [6/6]	0.33257	0.33210	0.33211	0.33239	0.33344
44 [7/7]	0.33257	0.332096	0.332055	0.332033	0.33199
50 [8/8]	0.33257	0.3320964	0.3320614	0.3320846	0.3322356
Error [8,8]	5.1E-4	3.9E-5	4.1E-6	2.7E-5	1.8E-4
Exact	0.33205734				

$x=6$. However, one can show from the differential equation[7] that f asymptotes to its large-x form with an error that is proportional to $\exp(-0.25x^2)$. At $x=8$, the [8,8] approximant has an error only slightly larger than 1 part in a million. For $x>8$, the error in the last row of Table III increases with x because the [8,8] approximant is not fully converged to the exact $f(x)$; we need larger N for larger x.

A reviewer suggested an alternative procedure. First, write the solution as

$$f(x;z)=zx^2\{u(zx^3)\}^{1/3}. \qquad (13)$$

The power series for $u(zx^3)$ may be obtained by dividing that for $f(x;z)$ by zx^2 and cubing the result. The $[N/N+1]$ approximant to u in the variable $Z\equiv zx^3$ will give an approximation to f that grows *linearly* with x as it should. One can then impose the condition that, denoting the highest coefficients in the numerator and denominator polynomials by p_N and q_{N+1},

$$p_N(z)/q_{N+1}(x)-1=0, \qquad (14)$$

which is equivalent to demanding that

$$\lim_{x\to\infty} zx^2\{[N/N+1](zx^3)\}^{1/3}\sim x.$$

It is not necessary to evaluate the approximant for intermediate x.

Unfortunately, this procedure converges very slowly. It is not unusual for Padé approximants to converge even at infinity; the bad news is that it is common for the rate of convergence to deteriorate from exponential to algebraic, that is, the error decreases as N^{-r} for some small positive constant r.[1,2] This is apparently what happens with the cube-root-of-the-Padé approach. However, evaluation of the approximant at finite x as illustrated in Table III will usually succeed even when the limit of the approximant as $x\to\infty$ is unbounded or very slowly convergent.

IV. NONPHYSICAL ROOTS

The FP monopole is unique, but the $[N/N]$ Padé approximant yields a polynomial equation for z that has in general N roots. How does one separate the good root from a swarm of bad ones?

The first part of the answer is that, if $u(r)$ is real-valued, all complex roots can be immediately excluded. The second part is that sometimes real roots can be rejected on physical grounds. For example, as noted above, one can show that it is impossible for $u(0)$ to be positive for a solution of the Flierl–Petviashvili equation.

When these fail, the only remaining criterion is convergence. The good roots will not change with degree N when N is sufficiently large; the bad roots will hop around and never settle down as shown in Table I.

We have not shown the extraneous roots for the other tables for simplicity, but the problem of excluding extra roots is very general. Whatever the discretization, the finite dimensional matrix that approximates an eigenoperator will always have many eigenvalues, typically at least half, even for a Chebyshev polynomial discretization, which are terrible approximations to those of the differential operator. The system of N algebraic equations that results from discretizing a quadratically nonlinear differential equation such as the Flierl–Petviashvili equation will have in general 2^N solutions. Most are complex valued, thank goodness, but some are real valued and have to be excluded by brainwork and multiple computation, that is, physical reasoning and nonconvergence with increasing N.

V. SURPRISES

One surprise was that the Padé method, at high order, approximates $u(0)$ for the Flierl–Petviashvili vortex by a cluster of closely spaced roots as shown in Table I. We cannot rigorously explain this phenomenon, which appears for the other two examples also.

An even more unfortunate surprise is that when extraneous roots were substituted into the approximants of various orders, the result was a smooth curve that often could not be dismissed on any obvious physical grounds. The failure to converge with increasing degree is sometimes the only reason for rejecting roots.

VI. SUMMARY

Padé methods have been used with great success to sum divergent perturbation series for eigenvalues, to evaluate special functions, and so on. It is rather odd that we have not seen them used to solve boundary-value problems *ab initio*.

To show as dramatically as possible that the Padé approximant overcomes the difficulty of the finite radius of convergence of a power series, all three of our examples are posed on an unbounded spatial domain. However, the same algorithm should work just as well for problems on a finite interval.

The algorithm can be generalized to systems of ordinary differential equations. Indeed, we applied this method in an unsuccessful search for nonaxisymmetric solutions of the quadratic Poisson equation, whose radially symmetric solution is the Flierl–Petviashvili monopole. For a system of M equations in M unknowns, the result is normally a coupled system of polynomial equations in M unknowns. This can be solved by blackbox rootfinders such as the well-known Hompack code (in Fortran), or within Maple by computing either Grobner bases or repeatedly taking the resultant of pairs of polynomial equations to reduce the system to a single equation in a single unknown.

Both the theory of Padé approximants and empirical experience in applying them to a wide variety of functions in physics suggest that the error decreases *exponentially* fast with the degree N of the approximant, just as for a spectral method, and much better than a standard finite difference algorithm[1–3] (Fig. 1). However, a rigorous convergence proof is lacking.

We have not attempted to compare the Padé algorithm with other methods for solving boundary value problems. Historically, the Flierl–Petviashvili monopole was first found by shooting;[4] the ''exact'' values for the anharmonic oscillator quoted above were computed by a rational Chebyshev pseudospectral method; the Blasius problem was first solved by matching the power series to an asymptotic approximation,[8] then later by a steepest descent algorithm,[7,9–11] by conversion to a nonlinear integral equation and iteration,[12] and a variety of other schemes.[7]

Still, seven lines of Maple code is not much of a price for an exponential rate of convergence. At the very least, our algorithm shows that power series are a very good way to solve differential equations even when their radius of convergence fails to span the interval, but only when converted to Padé approximants.

It is also remarkable that the Padé algorithm is a *direct* method. Methods for solving nonlinear problems are almost always iterative and require a first guess for the solution to initialize the iteration. Polynomial equations are the great exception; robust root solvers are available for this special class of nonlinear equations that do not require any user input except the coefficients of the polynomial. The Padé algorithm is also a direct method because it reduces the differential equation to the much simpler problem of finding the roots of a polynomial equation.

ACKNOWLEDGMENT

This work was supported by the National Science Foundation through Grant Nos. OCE9119459 and OCE9521133.

REFERENCES

1. G. A. Baker, Jr., Padé approximants, *Advances in Theoretical Physics* (Academic, New York, 1965), No. 1, pp. 1–50.
2. G. A. Baker, Jr., *Essentials of Padé Approximants* (Academic, New York, 1975).
3. C. M. Bender and S. A. Orszag, *Advanced Mathematical Methods for Scientists and Engineers* (McGraw-Hill, New York, 1978).
4. G. R. Flierl, Dynam. Atmos. Oceans **3**, 15 (1979).
5. V. I. Petviashvili, Sov. Phys. JETP Lett. **32**, 619 (1981).
6. J. P. Boyd, *Chebyshev and Fourier Spectral Methods* (Springer, New York, 1989).
7. G. A. Baker, Jr., *New Methods in Laminar Boundary-Layer Theory* (Pergamon, New York, 1961).
8. H. Blasius, Z. Math. Phys. **56**, 1 (1908).
9. L. Bairstow, J. R. Aero. Soc. **19**, 3 (1925).
10. D. Meksyn, Proc. R. Soc. A London, Ser. **237**, 543 (1956).
11. D. Meksyn, C. R. Acad. Sci. **T248**, 2286 (1959).
12. H. Weyl, Proc. Natl. Acad. Sci. USA **27**, 578 (1941).