

# PSEUDO-TRANSIENT CONTINUATION AND DIFFERENTIAL-ALGEBRAIC EQUATIONS \*

TODD S. COFFEY <sup>†</sup>, C. T. KELLEY <sup>†</sup>, AND DAVID E. KEYES <sup>‡</sup>

**Abstract.** Pseudo-transient continuation is a practical technique for globalizing the computation of steady-state solutions of nonlinear differential equations. The technique employs adaptive time-stepping to integrate an initial value problem derived from an underlying ODE or PDE boundary value problem until sufficient accuracy in the desired steady-state root is achieved to switch over to Newton’s method and gain a rapid asymptotic convergence. The existing theory for pseudo-transient continuation includes a global convergence result for differential equations written in semidiscretized method-of-lines form. However, many problems are better formulated or can only sensibly be formulated as differential-algebraic equations (DAEs). These include systems in which some of the equations represent algebraic constraints, perhaps arising from the spatial discretization of a PDE constraint.

Multirate systems, in particular, are often formulated as differential-algebraic systems to suppress fast time scales (acoustics, gravity waves, Alfvén waves, near equilibrium chemical oscillations, etc.) that are irrelevant on the dynamical time scales of interest. In this paper we present a global convergence result for pseudo-transient continuation applied to DAEs of index 1, and we illustrate it with numerical experiments on model incompressible flow and reacting flow problems, in which a constraint is employed to step over acoustic waves.

**Key words.** pseudo-transient continuation, nonlinear equations, steady-state solutions, global convergence, differential-algebraic equations, multi-rate systems

**AMS subject classifications.** 65H10, 65J15, 65L80, 65N12, 65N22

**1. Introduction.** Pseudo-transient continuation ( $\Psi tc$ ) is a globalization technique for the nonlinearly implicit computation of steady-state solutions of partial differential equations. Mathematically, it is a homotopy that embeds the steady-state problem in a space-time setting. The technique is popular in the aerodynamics [38], magnetohydrodynamics [24], radiation transport [35], reacting flow [37], structural analysis [17], and circuit simulation [14] modeling communities, among others, because conventional methods for solving nonlinear equations when the initial iterate is far from a solution, such as line-search and trust region methods [9, 18, 32], can converge to nonphysical solutions or local minima of the norm of the steady-state residual [7, 22]. This is particularly the case when the solution has complex features, such as shocks or discontinuities, that are not present in the initial iterate [31].  $\Psi tc$  is natural to apply in systems that originally derive from transient dynamics. While following a physical transient accurately may be inefficient when only the steady state is required, the existence of a physical space-time trajectory provides a basis for robust nearby pseudo-transient approaches. As a result,  $\Psi tc$  is available as a solver option in numerous general use simulator software packages, e.g., for fluid flows [1, 30] and electrical circuits [28, 29].

In general terms, pseudo-transient continuation is a predictor-corrector method for temporal integration in which the time step is increased as the iteration progresses with the objective of fast convergence near a solution. Temporal accuracy is not the objective, but in most cases the method begins by integrating accurately in time, closely following the transient behavior of the solution in the early stages of the iteration until an approximate steady state is reached. The time step is thereafter increased, sacrificing temporal accuracy in the terminal phase in favor of rapid convergence to steady state.

There are many other forms of continuation, some of which can be employed directly on the steady-state boundary value problem of interest. Parameter continuation [34], mesh sequencing [36], and model sequencing [22] can be effective in globalizing Newton’s method applied to large-scale nonlinear problems. Another complementary approach towards globalizing Newton’s method is nonlinear preconditioning of additive Schwarz type, recently introduced in [6]. In practice, we advocate mesh sequencing as a primary strategy in any highly-resolved nonlinear problem — solving first on a relatively coarse mesh, interpolating to the next finer mesh to produce an initial iterate for Newton’s method that lies within or close to the ball of convergence on the finer mesh (where iteration is more expensive), and so forth recursively [36]. Any of these globalization techniques can, in principle, be combined with each other and with  $\Psi tc$ .

---

\*Submitted for publication in a special issue of SIAM Journal of Scientific Computing (SISC) dedicated to papers from the 2002 Copper Mountain Conference on Iterative Methods.

<sup>†</sup> Center for Research in Scientific Computation, Mathematics Department, Box 8205, North Carolina State University, Raleigh, NC 27695-8205, (tscoffe2@math.ncsu.edu, tim\_kelley@ncsu.edu). The research of these authors was supported in part by National Science Foundation grant DMS-0070641. Current address of first author: Sandia National Laboratories, New Mexico, PO Box 5800, Albuquerque, NM 87185-1110.

<sup>‡</sup> Department of Mathematics and Statistics, Old Dominion University, Norfolk, VA 23529-0077, (dkeyes@odu.edu). The research of this author was supported in part by Lawrence Livermore National Laboratory under subcontract B347882, by the U.S. Department of Energy under cooperative agreement DE-FC02-01ER25476, and by the National Aeronautics and Space Administration under grant NAS1-19480, while the author was in residence at ICASE.

A particularly effective algorithmic combination is  $\Psi$ tc on each level of a mesh-sequenced approach. Depending upon the physical understanding that can be brought to bear on a specific application, any of the techniques that work directly on the steady-state formulation might be more effective than  $\Psi$ tc alone, but  $\Psi$ tc may have a more complete general mathematical theory than some of these other methods, especially with the recent additions of this paper and [10, 16]. Theoretical exploration of the interaction of  $\Psi$ tc with these other techniques is likely to be fruitful, but is beyond the scope of the present contribution.

**1.1. ODE Dynamics.** For problems with ODE dynamics, sufficient conditions for convergence of  $\Psi$ tc are given in [20], using rules for growth of the time step that are common in practice. We give a brief review of those results here, both to introduce notation and to put the new results in perspective.

If we express the initial value problem as  $u' = F(u)$ ,  $u(0) = u_0$ , we seek a root, if it exists, of  $F(u)$  that is the limit of  $u(t)$  as  $t \rightarrow \infty$ . Conventional nonlinear equation algorithms can (and do) find other solutions; hence the need for a continuation method that respects the physical solution.

Assuming that a stable steady-state solution exists, global convergence and local superlinear convergence to that solution are proved in [20] for a class of methods that integrate

$$u' = -V^{-1}F(u), \quad u(0) = u_0 \quad (1.1)$$

by a variable time step method that attempts to increase the time step as the integration progresses and steady state is approached. In (1.1),  $V$  is a scaling or preconditioning matrix. One method considered in [20] is

$$u_{n+1} = u_n - (\delta_n^{-1}V + F'(u_n))^{-1}F(u_n), \quad (1.2)$$

where  $F'$  is the Jacobian and  $\delta_n$  is the time step. The Newton step may be computed inexactly in the sense of [8, 18] and the convergence results change in an expected way.

A common choice for  $\delta_n$ , covered by the results in [20], is the “switched evolution relaxation” (SER) method [27],

$$\delta_n = \delta_{n-1} \|F(u_{n-1})\| / \|F(u_n)\| = \delta_0 \|F(u_0)\| / \|F(u_n)\|. \quad (1.3)$$

SER is commonly used in computational aerodynamics [21, 31, 38]. In practice, the growth of  $\{\delta_n\}$  may be limited from above, relative to (1.3), which can otherwise be rather aggressive. The results in the present paper, as with those in [20], assume an update of the form

$$\delta_n = \phi \left( \delta_{n-1} \frac{\|F(u_{n-1})\|}{\|F(u_n)\|} \right). \quad (1.4)$$

In (1.4)

$$\phi(\xi) = \begin{cases} \xi & \xi \leq \xi_t \\ \delta_{\max} & \xi > \xi_t \end{cases}. \quad (1.5)$$

where either  $\xi_t = \delta_{\max}$  or  $\xi_t < \infty$  and  $\delta_{\max} = \infty$ .

As a method for integrating in time,  $\Psi$ tc is a Rosenbrock method [12] when  $\delta$  is fixed. One may also think of this as a predictor-corrector method, where the simple predictor (result from the previous time) and a Newton corrector are used.

The motivating example in [20] is Euler flow over an airfoil.  $\Psi$ tc is a natural for this problem since the time-dependent Euler equations have physical meaning. The main result of [20] is that if the desired steady-state solution is a stable steady state for the time-dependent problem and the initial step is sufficiently small, then either the  $\Psi$ tc iteration stagnates with  $\inf_n \delta_n = 0$  or the iteration converges to the desired steady-state solution and the convergence in the terminal phase is as fast as an inexact-Newton method with the same algorithmic parameters. Of course, stagnation is an undesirable outcome, and ultimately requires the numerical analyst to exercise ingenuity in finding a better initial iterate. However, this “graceful” termination is preferable to allowing an erratic Newton process to enter an infeasible regime, where perhaps certain operations would generate floating point exceptions, perhaps leading to the unexplained hanging of a large-scale parallel execution.

There have been other recent theoretical studies of  $\Psi$ tc. Trust-region methods are used in an analysis of gradient flows in [16] to control the growth of the time step and prove convergence. Affine invariance ideas are applied to

problems with linear dynamical invariants in [10]. The time step is controlled with a method that seeks to increase the step but, at the same time, demands a decrease in the norm of the nonlinear residual. We did not ask for such a decrease in [20], nor do we here. In the transient phase, early in the iteration, the norm of the steady-state residual may well increase, as indeed may happen in a time-accurate integration.

We close this section with a short discussion of the differences between  $\Psi$ tc and a more conventional Newton-Armijo [19] nonlinear solver. The Newton-Armijo solver insists on a reduction in the norm of the nonlinear residual at each step and will either converge to a solution, diverge to infinity, or stagnate at a point where  $F'$  is singular.  $\Psi$ tc will accept an increase in the residual, responding to that increase (in the case of (1.5)) by decreasing  $\delta$ . Most implementations, including the one in this paper, impose a lower bound on  $\delta$  and terminate the iteration unsuccessfully if  $\delta$  reaches that lower bound. The particular value of that lower bound depends on the problem. If the dynamic problem to be solved does not have a steady-state solution, a Newton-Armijo iteration is unaffected, while a  $\Psi$ tc iteration may fail. The failure could well be signaled by reaching the bound on the total number of iterations. We view this as the likely signal of such a failure since, in a time-accurate integration, the norm of the residual would neither diverge to infinity nor decrease to zero.

**1.2. DAE Dynamics.** One of the limitations of the theory developed in [10, 16, 20] is the restriction to ODE dynamics. Many problems are best formulated, or can only be formulated, as differential-algebraic equations (DAEs). In § 3 we give two examples from fluid mechanics, an incompressible flow and a variable-density reacting flow. We seek a steady-state balance of the discretized conservation laws of mass, momentum, energy, and (as applicable) chemical species. For all equations but the mass (“continuity”) equation, we evolve the solution through pseudo-time stepping. In effect, we parabolize these nonlinear elliptic PDEs by prepending a time derivative term in the appropriate dependent variable. This is not only a physically natural form of nonlinear relaxation; it also leads to superior linear conditioning of the discrete systems resulting at each implicit time step (see, e.g., [11]).

However, we argue on the basis of physical intuition that it could be harmful to allow the dependent variables of the mass conservation equation to relax by prepending a similar tendency term. Instead, we force the progression from initial iterate to steady-state solution to occur on a manifold of continuity-satisfying states. We require this because linear acoustic waves in a time-dependent hydrodynamical system arise from a combination of the tendency and convective terms in the continuity equation and the tendency and convective terms in the momentum equations [26]. We can suppress acoustic-like oscillations in the solution trajectory by not equipping our pseudo-transient system with this mechanism. Indeed, the formulations we choose, velocity-vorticity and streamfunction-vorticity, are historically motivated, in part, by their suppression of the acoustic mechanism. We do not want to readmit this carefully filtered mechanism by relaxing the continuity constraint in pursuit of nonlinear robustness. In computing the steady-states of other multi-rate physical phenomena possessing fast waves or stiff restoring forces, we may similarly choose to omit a tendency-term piece of the fast mechanism. A combination analogous to the continuity equation and acoustic waves is the hydrostatic potential equation and gravity waves in shallow water wave theory.

We therefore need to be allowed to formulate our governing system as a partial differential algebraic equation (PDAE) of index 1 [5]. Let the vector of unknowns be  $x = [u, v]^T$  and the conformally partitioned vector of steady conservation laws be  $F = [f(u, v), g(u, v)]^T$ , where the first equation set takes the pseudo-transient tendency terms. The steady-state equation is then  $F(x) = 0$  and the time-dependent equation is (allowing for different scalings of the transient terms)

$$Dx' = -F(x), \tag{1.6}$$

where

$$D = \begin{pmatrix} V & \\ & 0 \end{pmatrix}$$

and  $V$  is a nonsingular diagonal matrix. In fact, none of the results to follow depend on  $V$  being either constant or even diagonal.  $V$  may depend on the evolving solution vector, as it does in our second example, and  $V$  could be an unlumped mass matrix in the context of finite elements. We require only that  $V$  be easily invertible everywhere along the solution trajectory.

The purpose of this paper is to extend the results from [20] to the semi-explicit index-one DAE case and thereby understand the historical success of  $\Psi$ tc as a solver for such systems. In the remainder of this paper we state and prove new results for the DAE case in § 2, and illustrate with a pair of examples in § 3.

**2.  $\Psi$ tc for Differential Algebraic Equations.** In this section we extend the current theory for  $\Psi$ tc to include a global convergence result for semi-explicit index-1 DAEs. We consider DAEs of the form

$$\begin{aligned} Vu' &= -f(u, v) \\ 0 &= g(u, v) \end{aligned} \quad (2.1)$$

with initial value  $(u(0), v(0)) = (u_0, v_0)$ ,  $u \in R^{N_1}$ ,  $v \in R^{N_2}$ , where  $V$  is an  $N_1 \times N_1$ , nonsingular scaling matrix and  $\partial g/\partial v$  is nonsingular. Our result follows from the work in [20], and we prove a theorem corresponding to the final stage of the  $\Psi$ tc convergence.

We write the DAE (2.1) in the form of (1.6), where  $x = [u, v]^T \in R^N$ ,  $N = N_1 + N_2$ ,  $F = [f, g]^T$ , and

$$Dx' = -F(x), \quad x(0) = x_0. \quad (2.2)$$

The pseudo-transient continuation procedure will be defined by the iteration

$$x_{n+1} = x_n - (\delta_n^{-1}D + F'(x_n))^{-1}F(x_n), \quad (2.3)$$

where  $\{\delta_n\}$  is given by (1.4).

We assume that the DAE has an index of one, has a global solution in time, and that the solution converges to a steady state.

ASSUMPTION 2.1.


1. The initial values  $(u_0, v_0)$  are consistent, i.e.  $g(u_0, v_0) = 0$ .
2.  $F$  is twice Lipschitz continuously differentiable.
3. The solution  $x(t) = [u(t), v(t)]^T$  of (2.2) exists for all  $t > 0$  and  $\lim_{t \rightarrow \infty} x(t) = x^* \in R^N$ .
4.  $g_v(u(t), v(t)) = g_v(x(t))$  is a nonsingular  $N_2 \times N_2$  matrix for all  $t \geq 0$  and  $\|g_v(x(t))^{-1}\|$  is uniformly bounded on  $(0, \infty)$ .

The analysis hinges on the fact that the  $\Psi$ tc iteration remains in a neighborhood of the solution of (2.2),

$$S(\epsilon) = \{z \mid \inf_{t \geq 0} \|z - x(t)\| \leq \epsilon\}. \quad (2.4)$$

Given  $\epsilon > 0$ , we will show that if  $\delta_0$  is sufficiently small, then the  $\Psi$ tc iteration remains within  $S(\epsilon)$  for all  $n$  and accurately approximates  $\{x(t_n)\}$  until the iteration is within  $B(\epsilon_T)$  the ball of convergence for Newton's method. Figure 2.1 illustrates the idea.

FIG. 2.1.  $S(\epsilon)$  and  $B(\epsilon_T)$ .



This corresponds to the first phase of the iteration as identified in [20]. The complete convergence analysis requires more. Once the iteration is near  $x^*$ , one must show that the  $\delta_n$  increases so as to obtain fast convergence and that the  $\Psi$ tc iteration does not diverge while this increase is in progress. Once this intermediate phase is complete, the final phase is a rapidly convergent Newton or inexact Newton iteration [20]. We describe the latter two phases first. These can be studied with a local convergence analysis that depends on the DAE dynamics only in a trivial way.

In order for the  $\Psi$ tc iteration to be well defined, we must assume that the linear systems to be solved are nonsingular and that the steady-state solution is stable.

ASSUMPTION 2.2. *There are  $\epsilon_S, \epsilon_G > 0$  such that*

1. *For all  $z_0 \in S(\epsilon_S)$ , the solution of  $Dz' = -F(z)$ ,  $z(0) = z_0$  exists,  $z(t) \in S(\epsilon_G)$  for all  $t$ , and  $\lim_{t \rightarrow \infty} z(t) = x^*$ ,  $t \rightarrow \infty$ .*
2. *Moreover, there are  $M_D, M_I > 0$  such that for all  $\delta > 0$ ,*
  - (a)  *$(\delta^{-1}D + F'(x))$  is nonsingular for all  $x \in S(\epsilon_G)$ ,*
  - (b)  *$\|(\delta^{-1}D + F'(x))\| \leq M_D$  for all  $x \in S(\epsilon_G)$ , and*
  - (c)  *$\|(\delta^{-1}D + F'(x))^{-1}\| \leq M_I$  for all  $x \in S(\epsilon_G)$ .*

The foregoing assumptions are mathematically reasonable and conducive to a compact theory. However, they may be too severe in practice, even for relatively simple nonlinear problems containing material interfaces, discretized with nondifferentiable flux limiters, etc. Pseudo-transient continuation has been effectively used well beyond these restrictions on the model for  $F(u)$ , including in mild violations in the examples of this paper.

**2.1. Local Convergence.** The assumptions for local convergence differ from the standard ones for nonlinear equations [18] in that we must allow for the possibility that  $\delta$  can be small even if  $x$  is near  $x^*$ .

Let

$$B(\epsilon) = \{x \mid \|x - x^*\| < \epsilon\}.$$

Our assumptions on continuity and differentiability of  $F$  and nonsingularity of  $F'$  near  $x^*$  are

ASSUMPTION 2.3. *There are  $\beta, \epsilon_L > 0$  such that for all  $x \in B(\epsilon_L)$  and all  $\delta > 0$*

$$\|(D + \delta F'(x))^{-1}D\| \leq 1/(1 + \beta\delta).$$

Assumption 2.3 is the DAE analog of case 3 of Assumption 2.1.3 in [20], and is critical for stability in the middle phase of the iteration.

Theorem 2.1 describes the middle and terminal phases of the iteration in the way that Theorems 2.1 and 2.3 of [20] do in the ODE case. The proof differs only in the local convergence estimate.

Once the approximate solution is near  $x^*$ , then the iteration will converge and  $\delta$  will increase.

THEOREM 2.1. *Let Assumptions 2.1, 2.2 and 2.3 hold. Let  $\{\delta_n\}$  be given by (1.4). Then there are  $C_T, \epsilon_T > 0$  so that if  $x_0 \in B(\epsilon_T)$  then either  $\inf_n \delta_n = 0$  or  $\delta_n \rightarrow \delta_{max}$ , the  $\Psi$ tc iteration converges, and for  $n$  sufficiently large*

$$\|x_{n+1} - x^*\| \leq C_T \|x_n - x^*\| (\delta_{max}^{-1} + \|x_n - x^*\|). \quad (2.5)$$

*Proof.* As is standard [18], we describe the progress of the iteration in terms of the transition from a current iterate  $x_c$  to a new one  $x_+$ . We let  $e = x - x^*$ . Let  $\epsilon_T < \epsilon_L$  be small enough so that the local convergence theory for Newton's method holds for  $F(x) = 0$ . Then, if  $x_c \in B(\epsilon_T)$

$$\begin{aligned} e_+ &= e_c - (\delta_c^{-1}D + F'(x_c))^{-1}F(x_c) \\ &= e_c - (\delta_c^{-1}D + F'(x_c))^{-1}(\delta_c^{-1}D + F'(x_c))e_c \\ &\quad + (\delta_c^{-1}D + F'(x_c))^{-1}(F'(x_c)e_c - F(x_c)) \\ &\quad + (\delta_c^{-1}D + F'(x_c))^{-1}\delta_c^{-1}De_c. \end{aligned} \quad (2.6)$$

The standard local convergence theory for Newton's method and Assumption 2.3 imply that there is  $C_N$  such that

$$\|(\delta_c^{-1}D + F'(x_c))^{-1}(F'(x_c)e_c - F(x_c))\| \leq C_N \|e_c\|^2,$$

Assumption 2.3 implies that

$$\|(\delta_c^{-1}D + F'(x_c))^{-1}\delta_c^{-1}De_c\| \leq \|e_c\|/(1 + \beta\delta_c).$$

Now assume that  $\{\delta_n\}$  is bounded from below by  $\delta^*$ . Since  $\delta_c \geq \delta^*$ ,

$$\|e_+\| \leq (2C_N\|e_c\| + (1 + \beta\delta^*)^{-1})\|e_c\|.$$

Reducing  $\epsilon_T$  if needed so that

$$(2C_N\epsilon_T + (1 + \beta\delta^*)^{-1}) < (1 + \beta\delta^*/2)^{-1}$$

implies that the local convergence is q-linear provided  $\{\delta_n\}$  is bounded from below.

The q-linear convergence will eventually drive  $\delta_n$  to  $\delta_{max}$  if  $\delta_{max} < \infty$ . Hence, the standard local convergence theory for Newton's method [18] will imply (2.5).  $\square$

**2.2. Global Convergence.** The objective of this section is to show that if  $\delta_0$  is sufficiently small, then  $x_n \in S(\epsilon)$  for some  $\epsilon < \min(\epsilon_S, \epsilon_T)$  until  $x_n \in B(\epsilon_T)$ , where the local convergence theory holds.

**THEOREM 2.2.** *Let Assumptions 2.1, 2.2, and 2.3 hold. If  $\delta_0$  is sufficiently small and  $\{\delta_n\}$  is bounded from below, then  $x_n \rightarrow x^*$  and (2.5) holds for  $n$  sufficiently large.*

*Proof.* Assume that  $\delta_n \geq \delta^*$ . Let

$$t_n = \sum_{j=0}^{n-1} \delta_j.$$

The assumptions imply that there is  $T$  such that  $x(t) \in B(\epsilon_T/2)$  for all  $t \geq T$ , where  $B(\epsilon_T)$  is the ball of local convergence from Theorem 2.1. Hence if  $N \geq T/(2\delta^*)$ ,  $x(t_N) \in B(\epsilon_T/2)$ . The proof will show that for  $\delta_0$  sufficiently small

$$\|x_n - x(t_n)\| < \epsilon_T/2$$

for all  $n \leq N$ . This will imply that  $x_N \in B(\epsilon_T)$  and complete the proof.

We will show that the  $\Psi$ tc iteration is a consistent first-order explicit method for the ODE

$$\begin{bmatrix} u \\ v \end{bmatrix}' = x' = G(x) = \begin{bmatrix} -V^{-1}f \\ g_v^{-1}g_u V^{-1}f \end{bmatrix}, \quad (2.7)$$

which is equivalent to (2.1). Our assumptions imply that  $G$  is Lipschitz differentiable and hence the explicit Euler method is first-order accurate.

We will show that the local error of Euler's method for (2.7) is within  $O(\delta^2)$  of that for the  $\Psi$ tc method. This will complete the proof.

Assume that  $x_n = x(t_n)$ , then the Euler step for (2.7) can be written as  $x_{n+1} = x_n + \delta_n s$ , where

$$s = \begin{bmatrix} s_u \\ s_v \end{bmatrix} = \begin{bmatrix} -V^{-1}f \\ g_v^{-1}g_u V^{-1}f \end{bmatrix}.$$

Hence, for the Euler discretization

$$g_u s_u + g_v s_v = 0. \quad (2.8)$$

The  $\Psi$ tc discretization is  $x_{n+1} = x_n + \delta_n \sigma$ , where

$$\sigma = \begin{bmatrix} \sigma_u \\ \sigma_v \end{bmatrix} = -(D + \delta_n F'(x_n))^{-1} F(x_n).$$

The equations for  $\sigma_u$  and  $\sigma_v$  are

$$V\sigma_u + \delta_n(f_u\sigma_u + f_v\sigma_v) = -f \quad (2.9)$$

and

$$\delta_n(g_u\sigma_u + g_v\sigma_v) = -g. \quad (2.10)$$

Assuming  $x_n = x(t_n)$ , we have  $g = 0$  and thus

$$g_u\sigma_u + g_v\sigma_v = 0,$$

which is similar to (2.8) and  $\sigma_v$  can be computed from  $\sigma_u$  in exactly the same way  $s_v$  can be computed from  $s_u$ . Therefore the proof will be complete if we can show that

$$\sigma_u = s_u + O(\delta_n). \quad (2.11)$$

Since  $Vs_u = -f$ , (2.11) follows from (2.9) and the nonsingularity of  $V$ .  $\square$

**3. Computational Examples.** In this section we illustrate the DAE form of  $\Psi\text{tc}$  with two simple examples of elliptically-constrained systems: incompressible Boussinesq flow in a lid- and buoyancy-driven cavity and reacting flow in a laminar diffusion flame with single-step infinitely fast kinetics — a classic flamesheet model. The first example is freely downloadable as part of the release of the PETSc toolkit [2]. For the first example, we compare a fully parabolized  $\Psi\text{tc}$  formulation and a PDAE version and show that the latter has an advantage in computational efficiency.

**3.1. Incompressible Flow.** Our incompressible flow example is a combination of two classic problems, lid-driven flow and buoyancy-driven flow in a two-dimensional rectangular cavity. The lid, moving with a steady and spatially uniform velocity, sets up a principal vortex and subsidiary corner vortices by viscous forces and the differentially heated lateral walls of the cavity set up a buoyant vortex flow, which opposes the principal lid-driven vortex. Our parameterization and discretization are inspired by [3]; however, in this example we do not exploit the local adaptive mesh refinement that is the main focus of [3]. The PETSc implementation contains uniform-refinement mesh-sequencing capabilities, applied both in a nonlinear continuation sense in the outer iteration, as mentioned in § 1, and also possibly in an inner iteration as part of a multigrid linear solver, both of which are incidental to our focus here.

The governing system consists of the following system of four elliptic partial differential equations in two dimensions:

$$-\Delta u - \frac{\partial \omega}{\partial y} = 0, \quad (3.1)$$

$$-\Delta v + \frac{\partial \omega}{\partial x} = 0, \quad (3.2)$$

$$-\Delta \omega + u \frac{\partial \omega}{\partial x} + v \frac{\partial \omega}{\partial y} - \text{Gr} \frac{\partial T}{\partial x} = 0, \quad (3.3)$$

$$-\Delta T + \text{Pr} \left( u \frac{\partial T}{\partial x} + v \frac{\partial T}{\partial y} \right) = 0, \quad (3.4)$$

where  $(u(x, y), v(x, y))$  are the velocity fields in the  $(x, y)$  directions,  $\omega(x, y) \equiv -\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}$  is the component of the vorticity normal to the  $xy$ -plane (representing the in-plane rotation of an infinitesimal fluid element), and  $T(x, y)$  is the temperature.  $\text{Pr}$  is a Prandtl number and  $\text{Gr}$  is a Grashof number. (For those who would rather see a Peclet number than a Prandtl number in the energy equation, Reynolds number is unity, so  $\text{Pe}$  and  $\text{Pr}$  are identical.)

These equations are subject to boundary conditions:

1. Along the bottom ( $0 < x < 1, y = 0$ ):  $u = v = 0, \frac{\partial T}{\partial y} = 0$ .
2. Along the top ( $0 < x < 1, y = 1$ ):  $u = V_{lid}, v = 0, \frac{\partial T}{\partial y} = 0$ .
3. Along the left ( $x = 0, 0 < y < 1$ ):  $u = v = 0, T = 0$ .

4. Along the right ( $x = 1, 0 < y < 1$ ):  $u = v = 0, T = 1$  if  $Gr > 0$  or 0 otherwise.

On each boundary,  $\omega$  is given by its definition,  $\omega(x, y) = -\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}$ , based on the normal gradient of either  $u$  and  $v$  (the tangential gradients being zero, according to the velocity boundary conditions). We make the standard assumption of Boussinesq perturbation theory, namely that the effect of density variation is significant only in the body force term, when multiplied by the acceleration of gravity (in  $Gr$ ), and not in the mass or momentum flux terms. (The effects of buoyancy are omitted if  $Gr = 0$ .)

Nonlinear iterative methods for (3.1)-(3.4) require an initial iterate, while  $\Psi tc$  methods require an initial condition. For either, we start with a motionless, vorticity-free initial flow field. This initial condition is incompatible with the forcing boundary conditions on a set of measure zero for the continuous problem. Parameter continuation could be used to eliminate the resulting ‘‘impulsive start’’ but we do not pursue this possibility here.

In the interests of space, we leave the derivation of this velocity-vorticity form of the Navier-Stokes and energy governing equations (3.1)-(3.4) to the literature (e.g., [3] and citations therein). However, to motivate our main point, we note that the two equations with velocity components under the Laplacian operators, (3.1) and (3.2), come from differentiating the continuity equation and substituting from the definition of vorticity. Thus, for instance, we differentiate the continuity equation,  $\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0$ , by  $x$  and the definition of vorticity,  $\omega(x, y) = -\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}$ , by  $y$  and eliminate the term in  $v$  to get (3.1).

We discretize the system (3.1)-(3.4) with a standard five-point stencil for each component, on a uniform Cartesian grid, using standard first-order upwinding for the convective terms and central differencing for the other first-order gradients. All four solution components are vertex centered. This discretization, the source code for which can be examined in the PETSc release, is famously diffusive at large Reynolds numbers and would not be of acceptable order for deriving accurate engineering results on efficient meshes. However, it serves our purpose (and that of the PETSc example) by leading straightforwardly to a model problem with interesting nonlinearities. These nonlinearities may be tuned to be strong or weak, depending upon the magnitude of  $V_{lid}$  and/or the Grashof number.

PETSc [2] allows an enormous variety of solver combinations to be assembled via command-line options to solve (3.1)-(3.4). The reader is encouraged to download and build PETSc and experiment<sup>1</sup>. We prefer various forms of preconditioned Jacobian-free Newton-Krylov solvers [25]. Such a solver has as its heart an inexact Newton method [8]. The Newton-correction equation is solved approximately with a Krylov method, such as GMRES [33], which requires the action of the Jacobian of  $F(u)$  only in the form of Jacobian-vector products. Such Jacobian-vector products are available either in approximate finite-difference (Fréchet derivative) form, via construction of an explicit Jacobian matrix from analytical or numerical means, or via automatic differentiation (AD). The PETSc distribution example is of the AD variety, exploiting the automatic differentiation package ADIC [4]. In practice, the Jacobian-based linear systems for the Newton corrections are ill-conditioned, and require preconditioning. PETSc is equipped with many forms of preconditioning for this Cartesian grid discretization. We use the default, which is a block ILU relaxation sweep within a multilevel context. The block ILU kernel treats all four components common to each vertex implicitly. This is a well known practice in multicomponent problems with cross-coupling in transport and source terms, as in (3.1)-(3.4). (It also has excellent cache locality properties, since it leads to small BLAS3-like kernels that can be unrolled into registers. Such solver issues have been extensively studied elsewhere, e.g., [15]. In the context of this paper, the only aspect of performance about which we are directly concerned is the outer nonlinear convergence.)

We readily observe from experimentation that Newton’s method struggles when this simulation is started at high values of  $V_{lid}$  and  $Gr$ , on any but the coarsest grids. For illustrative purposes, we use dimensionless values of  $V_{lid} = 100$  and  $Gr = 10^5$  below. Without  $\Psi tc$ , the standard way to solve this system in PETSc is to use mesh sequencing, working up from a grid containing just  $4 \times 4$  cells at its coarsest. Starting from a ‘‘cold’’ initial iterate even for a grid only as fine as  $32 \times 32$  with these high parameter values for  $V_{lid}$  and  $Gr$ , Newton stagnates.

We can prepend a tendency term to each equation in (3.1)-(3.4) and attempt to overcome this problem by  $\Psi tc$ .

$$\frac{\partial u}{\partial t} - \Delta u - \frac{\partial \omega}{\partial y} = 0, \quad (3.5)$$

$$\frac{\partial v}{\partial t} - \Delta v + \frac{\partial \omega}{\partial x} = 0, \quad (3.6)$$

<sup>1</sup>Upon unpacking the release `petsc-2.1.5`, which is available at <http://www.mcs.anl.gov/petsc>, see `PETSC/src/snes/examples/tutorials/ex27.c`, where PETSC represents the local root directory of the package. By default, in this example, pseudo-transient continuation is applied to the vorticity and energy equations. To force it to be applied to all four equations, one uses the command-line switch ‘‘-use\_parab’’.






FIG. 3.1. Convergence of residual norm (dashed curve) and growth of pseudo time step (solid curve) for two versions of  $\Psi tc$  on the driven cavity example. The DAE form of  $\Psi tc$  converges faster than the fully parabolized ODE form.

$$\frac{\partial \omega}{\partial t} - \Delta \omega + u \frac{\partial \omega}{\partial x} + v \frac{\partial \omega}{\partial y} - \text{Gr} \frac{\partial T}{\partial x} = 0, \quad (3.7)$$

$$\frac{\partial T}{\partial t} - \Delta T + \text{Pr} \left( u \frac{\partial T}{\partial x} + v \frac{\partial T}{\partial y} \right) = 0, \quad (3.8)$$

The result is plotted in Figure 3.1. This figure contains two convergence plots of the Euclidean norm of the steady state residual,  $\|F(x_n)\|$  (“res”), as a function of iteration  $n$ , and corresponding plots of  $\delta_n$  (“dt”). The longer-running history pertains to the case in which (3.1)-(3.4) are fully parabolized in the manner of (3.5-3.8). We note terminal Newton-like convergence of the Euclidean norm of the steady-state residual “res<sub>parab</sub>” and growth of the pseudo-timestep “dt<sub>parab</sub>” in accordance with the  $\Psi tc$  theory of [20]. If, instead, we constrain the two velocity equations to be satisfied “exactly” (to within the limits of the overall inexact Newton method linear convergence), by not adding any time-stepping to those equations, i.e., solving (3.1)-(3.2) together with (3.7-3.8), we obtain the shorter-running history in Figure 3.1. Both techniques (ODE-based and DAE-based  $\Psi tc$ ) lead to convergence in a case where Newton alone fails, apart from mesh-sequencing. However, the DAE form of the technique is evidently superior. It converges after half as many pseudo-time steps and, unlike the fully parabolized case, does not suffer any nonmonotonicity in  $\|F(x_n)\|$  or  $\delta_n$ .

**3.2. Variable-density Reacting Flow.** As a second example, we consider the axisymmetric flamesheet model from [23], expressed as a system of three steady-state PDEs. Instead of the velocity-vorticity treatment of the driven cavity above, we use a streamfunction-vorticity formulation, following the the development of [13]. With  $r$  and  $z$  for the radial and axial directions in cylindrical coordinates and primitive variables  $v_r$  and  $v_z$  as the radial and axial velocities, respectively, we define the vorticity component normal to the axisymmetric plane, representing in-plane rotation,  $\omega = \frac{\partial v_r}{\partial z} - \frac{\partial v_z}{\partial r}$ , and the variable-density Stokes streamfunction,  $\psi$ , in terms of which

$$\rho r v_r = -\frac{\partial \psi}{\partial z} \quad \text{and} \quad \rho r v_z = \frac{\partial \psi}{\partial r}.$$

The three equations to be solved are mass conservation:

$$-\frac{\partial}{\partial z} \left( \frac{1}{r\rho} \frac{\partial \psi}{\partial z} \right) - \frac{\partial}{\partial r} \left( \frac{1}{r\rho} \frac{\partial \psi}{\partial r} \right) - \omega = 0, \quad (3.9)$$

momentum conservation:

$$\begin{aligned}
& -\frac{\partial}{\partial r} \left( r^3 \frac{\partial}{\partial r} \left( \frac{\mu\omega}{r} \right) \right) - \frac{\partial}{\partial z} \left( r^3 \frac{\partial}{\partial z} \left( \frac{\mu\omega}{r} \right) \right) + r^2 \left[ \frac{\partial}{\partial z} \left( \frac{\omega}{r} \frac{\partial\psi}{\partial r} \right) - \frac{\partial}{\partial r} \left( \frac{\omega}{r} \frac{\partial\psi}{\partial z} \right) \right] \\
& + r^2 g \frac{\partial\rho}{\partial r} + r^2 \nabla \cdot \left( \frac{v_r^2 + v_z^2}{2} \right) \cdot \text{iso } \rho = 0,
\end{aligned} \tag{3.10}$$

and species conservation:

$$-\frac{\partial}{\partial r} \left( r\rho D \frac{\partial S}{\partial r} \right) - \frac{\partial}{\partial z} \left( r\rho D \frac{\partial S}{\partial z} \right) + \frac{\partial}{\partial z} \left( S \frac{\partial\psi}{\partial r} \right) - \frac{\partial}{\partial r} \left( S \frac{\partial\psi}{\partial z} \right) = 0. \tag{3.11}$$

Here,  $\mu$  is mixture viscosity,  $D$  is molecular diffusivity, and the notation “iso” is defined by  $\text{iso}\rho = \left( -\frac{\partial\rho}{\partial z}, -\frac{\partial\rho}{\partial r} \right)^T$ . The system is closed by constitutive laws and an equation of state that express  $\mu$ ,  $D$ , and  $\rho$  as nonlinear functions of  $S$ .

A single conserved scalar,  $S$ , is capable of describing the complete composition and thermodynamic state of the fluid under the assumption of infinitely fast single-step kinetics and the additional assumptions that the binary diffusion coefficients of all pairs of species are equal and that the Lewis number (ratio of thermal to species diffusivity) is unity. Therefore, this governing system contains no separate species or energy conservation equations.  $S$  is a so-called Shvab-Zeldovich or mixture-fraction variable, as defined, for instance, in [39]. The flame front,  $\Gamma$ , is defined as a level set of  $S$ :  $\Gamma \equiv \{(r, z) | S(r, z) = S_f\}$ , on one side of which fuel cannot exist and on the other side of which oxidizer cannot exist, because each is consumed at the interface. The value of  $S_f$  comes from the stoichiometry of the reaction. For a detailed derivation of this formulation in the combustion literature in the case of a methane-air flame, see [22].

The geometry is that of a co-flowing gaseous burner, with an inner tube out of which flows a methane-air mixture and an outer annulus out of which flows a standard atmosphere. The boundary conditions are those of symmetry for  $r = 0$ , a stress-free free-stream for  $r = r_{max}$ , inhomogeneous Dirichlet values for velocities and chemical species at the inflow where  $z = 0$ , and extrapolative exit conditions at  $z = z_{max}$ . The inflow velocities are piecewise constant in  $r$ , different on either side of the radius that divides the inner tube from the outer co-flowing annulus. The inflow condition on  $S$  is constructed to represent a Gaussian peak of combustion products (carbon dioxide and water vapor) centered at the dividing radius between the co-flowing fuel and oxidizer streams and a corresponding temperature profile achieving the stoichiometrically appropriate peak value at the dividing radius. These conditions are swept downstream to form an initial condition throughout the domain. Thus, the problem of ignition is avoided. The flame is “underventilated” in the language of combustion theory, which implies that the initial flame tube must find a final bulbous pinched steady-state shape characteristic of the flame attached to a bunsen burner. There are numerous technical issues in flameholding whose discussion we suppress here, with the justification that this model is well published, its limitations well understood, and various remedies for those limitations available. (See, e.g., [37] and [11].) Though highly simplified, the problem is practically motivated in that it serves as a useful asymptotic precursor for finite-rate detailed kinetics flame structure models. It has been used to generate an initial condition for the detailed kinetics problem, as an example of “model continuation”.

Because of the polynomial nonlinearities that are evident in (3.9)-(3.11) and the more insidious nonlinearity of the various state functions ( $\mu$ ,  $D$ , and  $\rho$ ) of  $S$  as a function of the *a priori* unknown interface where  $S(r, z) = S_f$ , this system is virtually impossible to solve by Newton’s method from any readily specified initial iterate, including the relatively sophisticated, already ignited swept tubular flame state described above. A “natural” pseudo-transient approach prepends a term  $r\rho \frac{\partial S}{\partial t}$  to (3.11) and a term  $r^3\rho \frac{\partial\omega}{\partial t}$  to (3.10). The resulting transient equations are dimensionally and physically correct and could be used to time-accurately evolve the unsteady flame. Observe that since  $\rho$  is a function of  $S$ , this is an instance of a varying (but always nonsingular diagonal)  $V$ .

We argue on the basis of physical intuition and the experience of the prior example that it might be harmful to allow  $\psi$  to relax by prepending a similar tendency term in the streamfunction to (3.9). Like (3.1)-(3.2), Eq. (3.9) derives from the underlying primitive variable expression of the conservation of mass for the reacting flow. Due to chemical reaction and heat release, the density  $\rho$  varies (by up to an order of magnitude in common terrestrial hydrocarbon-air or hydrogen-air diffusion flame applications). This variation cannot be neglected outside of the buoyancy term, as it was in the earlier example. However, we do not want to admit density variations arising from transient momentum

imbalances, which could lead to acoustic-like waves. While there are many applications, such as turbulent combustion, in which thermal fluctuations and acoustic waves interact in dynamically relevant ways, these applications require far more complex models than (3.9)-(3.11) to describe. On the other hand, there are many applications for which the set (3.9)-(3.11) is a useful model. For these situations, we exploit our theory for  $\Psi$ tc that retains (3.9) as it is, an elliptic constraint enforced at every (pseudo-)time step to implicitly equilibrate pressure and suppress acoustic stiffness in a variable density context.

A convergence history for DAE-based  $\Psi$ tc for the flamesheet problem is shown in Figure 3.2. There is an initial rapid plunge of the steady-state residual on the first pseudo-time step. This corresponds to creating a consistent initial condition for the DAE (Assumption 2.1, part 1). There is no assurance in the limited theory of this paper that this first Newton step should be successful, which is why we start with a relatively small initial time step,  $\delta_0 = 10^{-5}$ . Following this point,  $\Psi$ tc convergence is fairly uneventful. However, it should be noted that the success of the method is relatively sensitive to the initial time step. If we attempt  $\delta_0 \geq 10^{-4}$ , for instance, the Newton iteration dies on the first step with an infeasible iterate (components of  $S$  are driven out of bounds). In particular, unassisted Newton iteration on the steady-state system is hopeless. If we try  $\delta_0 < 10^{-5}$ , Newton's method assisted by SER-style  $\Psi$ tc will converge, but will take more pseudo-time steps to build  $\delta_n$  up to a value at which a switchover to a full Newton iteration can safely occur. It would be desirable to come up with theory, heuristics, or a database of historical convergence results for related problems, by which to guide the choice of  $\delta_0$ .

In this implementation, the maximum increase of  $\delta_n$  from one time step to the next is limited to a factor of 2, except at the final switchover to a full Newton step, which is triggered by the pseudo-time step reaching unity. This succession of doublings is observed in the semi-log plot of Fig. 3.2 between iterations 2 and 13, during which time the residual norm is struggling nonmonotonically, but at a level that is substantially below  $\|F(x_0)\|$ , so the criterion (1.3) permits a growth in  $\delta_n$  that is limited only by the ratio bound. Obviously, many refinements of the strategy to pick  $\delta_n$  are possible but the simple one we have adopted is sufficient, if possibly inefficient.

We remark that the individual rootfinding problems (at each pseudo-timestep) in this run enjoy reasonable linear conditioning. This is evidenced by the number of zero-fill block ILU-preconditioned Krylov subiterations required in each inexact Newton step, which varies between 8 and just 13, until the last five pseudo-transient Newton rootfinding problems, when  $\delta_n$  becomes close to unity.




FIG. 3.2. Convergence of residual norm (dashed curve) and growth of pseudo time step (solid curve) for the DAE form of  $\Psi$ tc on the flamesheet example.

We remark that  $\Psi$ tc is more crucial for this flamesheet example than it was for the driven cavity example. For the driven cavity, mesh sequencing leads to a steady-state solution with or without  $\Psi$ tc. For the flamesheet, the problem

has too rich a geometric structure to be represented on a grid much coarser than  $16 \times 16$  in the axisymmetric plane, and the choice of initial mesh must be constructed rather carefully (with refinement in  $r$  around the mixing radius and gradual expansion in the  $z$ -direction downstream) for the flame to be sufficiently accurately discretized with second-order finite differences on such a coarse mesh. Even on a  $16 \times 16$  grid,  $\Psi_{tc}$  is crucial. After interpolation of the  $16 \times 16$  solution to a refined  $32 \times 32$  grid,  $\Psi_{tc}$  is again crucial to the success of Newton's method, as it is following interpolation of the  $32 \times 32$  method to a  $64 \times 64$  grid. It takes many doublings of grid density before the bilinearly interpolated solution on the refined grid lies initially in the ball of convergence of Newton's method applied directly to the steady-state system.

**4. Conclusion.** We have shown, under natural assumptions, that  $\Psi_{tc}$  is a globally convergent method for semi-explicit index-1 differential-algebraic equations. We have also argued that the index-1 DAE formulation is an important complement to the standard method-of-lines ODE formulation of  $\Psi_{tc}$ . Both propositions are illustrated by numerical experiments on model nonlinear fluid mechanical applications, which show, first, that  $\Psi_{tc}$  is effective and, second, that the DAE formulation can be more effective than the ODE formulation. Our model problems are of relatively low discrete dimension but are already sufficient to illustrate the challenges of applying Newton's method without  $\Psi_{tc}$ .

The objective of this paper is to analyze and illustrate  $\Psi_{tc}$  by itself in the context of standard inexact Newton-Krylov solvers, without much attention to the preconditioning of the resulting inner linear problems and unadorned by other complications such as mesh sequencing. In large-scale practice, linear preconditioning is important and is made easier by the addition of the implicit time step. Finally, to recapitulate the introduction, we believe that  $\Psi_{tc}$  is best exploited in combination with other techniques, especially of the multilevel variety, both for nonlinear robustness and for solution efficiency.

**Acknowledgments.** The third author wishes to acknowledge the assistance of Paul Hovland, Dinesh Kaushik, Lois McInnes, and Barry Smith of Argonne National Laboratory for their maintenance of the demonstration cavity flow problem within PETSc, and for upgrading it over the years to exploit evolving PETSc functionality.

#### REFERENCES

- [1] ANSYS, *Flotran*. <http://www.ansys.com/ansys/flotran.htm>.
- [2] S. BALAY, W. D. GROPP, L. C. MCINNES, AND B. F. SMITH, *The Portable Extensible Toolkit for Scientific Computing (PETSc), version 2.1.5*. <http://www.mcs.anl.gov/petsc/>, 2003.
- [3] B. A. V. BENNETT AND M. D. SMOOKE, *Local rectangular refinement with application to nonreacting and reacting fluid flow problems*, J. Comput. Phys., 151 (1999), pp. 648–727.
- [4] C. BISCHOF, L. ROH, AND A. MAUER, *ADIC — An extensible automatic differentiation tool for ANSI-C*, Tech. Report ANL/MCS-P626-1196, Argonne National Lab, 1996.
- [5] K. E. BRENNAN, S. L. CAMPBELL, AND L. R. PETZOLD, *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations*, SIAM, Philadelphia, PA, 1996.
- [6] X.-C. CAI AND D. E. KEYES, *Nonlinear preconditioned inexact Newton algorithms*, SIAM J. Sci. Comput., 24 (2002), pp. 183–200.
- [7] T. S. COFFEY, R. J. MCMULLAN, C. T. KELLEY, AND D. S. MCRABE, *Globally convergent algorithms for nonsmooth nonlinear equations in computational fluid dynamics*, Tech. Report CRSC-TR01-28, North Carolina State University, Center for Research in Scientific Computation, November 2001. to appear in J. Comp. Appl. Math.
- [8] R. DEMBO, S. EISENSTAT, AND T. STEihaug, *Inexact Newton methods*, SIAM J. Numer. Anal., 19 (1982), pp. 400–408.
- [9] J. E. DENNIS AND R. B. SCHNABEL, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, SIAM, Philadelphia, PA, 1996.
- [10] P. DEUFLHARD, *Adaptive pseudo-transient continuation for nonlinear steady state problems*, Tech. Report 02-14, Konrad-Zuse-Zentrum für Informationstechnik, Berlin, March 2002.
- [11] A. ERN, V. GIOVANGIGLI, D. E. KEYES, AND M. D. SMOOKE, *Towards polyalgorithmic linear system solvers for nonlinear elliptic problems*, SIAM J. Sci. Comput., 15 (1994), pp. 681–703.
- [12] C. W. GEAR, *Numerical Initial Value Problems in Ordinary Differential Equations*, Prentice-Hall, Englewood Cliffs, New Jersey, 1971.
- [13] A. D. GOSMAN, W. N. PUN, A. K. RUNCHAL, D. B. SPALDING, AND M. WOLFSHTEIN, *Heat and Mass Transfer in Recirculating Flows*, Academic Press, New York, 1969.
- [14] K.-T. GRASSER, *Mixed-mode device simulation*, tech. report, Technical University of Vienna (doctoral dissertation), 1999. [http://www.iue.tuwien.ac.at/publications/PhD Theses/grasser/diss.html](http://www.iue.tuwien.ac.at/publications/PhD%20Theses/grasser/diss.html).
- [15] W. D. GROPP, D. K. KAUSHIK, D. E. KEYES, AND B. F. SMITH, *High performance parallel implicit CFD*, Parallel Comput., 27 (2001), pp. 337–362.
- [16] D. J. HIGHAM, *Trust region algorithms and time step selection*, SIAM J. Numer. Anal., 37 (1999), pp. 194–210.
- [17] T. KANT AND S. PATEL, *Transient/pseudo-transient finite element small/large deformation analysis of two-dimensional problems*, Computers and Structures, 36 (1990), pp. 421–427.
- [18] C. T. KELLEY, *Iterative Methods for Linear and Nonlinear Equations*, SIAM, Philadelphia, PA, 1995.
- [19] ———, *Iterative Methods for Linear and Nonlinear Equations*, no. 16, SIAM, Philadelphia, 1995.
- [20] C. T. KELLEY AND D. E. KEYES, *Convergence analysis of pseudo-transient continuation*, SIAM J. Numer. Anal., 35 (1998), pp. 508–523.

- [21] D. E. KEYES, *Aerodynamic applications of Newton-Krylov-Schwarz solvers*, in Proceedings of the 14th International Conference on Numerical Methods in Fluid Dynamics, R. Narasimha, ed., New York, 1995, Springer, pp. 1–20.
- [22] D. E. KEYES AND M. D. SMOOKE, *Flame sheet starting estimates for counterflow diffusion flame problems*, J. Comput. Phys., 72 (1987), pp. 267–288.
- [23] ———, *A parallelized elliptic solver for reacting flows*, in Parallel Computations and their Impact on Mechanics, A. K. Noor, ed., 1987, pp. 375–402.
- [24] D. KNOLL AND P. MCHUGH, *Enhanced nonlinear iterative techniques applied to a nonequilibrium plasma flow*, SIAM J. Sci. Comput., 19 (1998), pp. 291–301.
- [25] D. A. KNOLL AND D. E. KEYES, *Jacobian-free Newton-Krylov methods: a survey of approaches and applications*. submitted to J. Comput. Phys., 2002.
- [26] L. D. LANDAU AND E. M. LIFSCHITZ, *Fluid Mechanics*, Pergamon Press, London, 1959.
- [27] W. MULDER AND B. V. LEER, *Experiments with implicit upwind methods for the Euler equations*, J. Comput. Phys., 59 (1985), pp. 232–246.
- [28] NEWBURY TECHNOLOGY LTD., *Simetrix*. <http://www.newburytech.co.uk/>.
- [29] NUMBER ONE SYSTEMS, *Easy-spice*. <http://www.numberone.com/easy-spice.htm>.
- [30] NUMERICAL ALGORITHMS GROUP, *Fastflo*. <http://www.nag.com/simulation/Fastflo/ToolBox/toolbox.asp>.
- [31] P. D. ORKWIS AND D. S. MCRAE, *Newton's method solver for the axisymmetric Navier-Stokes equations*, AIAA Journal, 30 (1992), pp. 1507–1514.
- [32] J. M. ORTEGA AND W. C. RHEINBOLDT, *Iterative Solution of Nonlinear Equations in Several Variables*, SIAM, Philadelphia, PA, 2000.
- [33] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.
- [34] A. G. SALINGER, N. M. BOU-RABEE, R. P. PAWLOWSKI, E. D. WILKES, E. A. BURROUGHS, R. B. LEHOUCQ, AND L. A. ROMERO, *Loca: Library of continuation algorithms: Theory and implementation manual*, Tech. Report SAND2002-0396, Sandia National Laboratory, March 2002.
- [35] A. SHESTAKOV AND J. MILOVICH, *Applications of pseudo-transient continuation and Newton-Krylov methods for the Poisson-Boltzmann and radiation diffusion equations*, Tech. Report UCRL-JC-139339, Lawrence Livermore National Laboratory, 2000.
- [36] M. D. SMOOKE AND R. M. MATTHEIJ, *On the solution of nonlinear two-point boundary value problems on successively refined grids*, Appl. Numer. Math., 1 (1985), pp. 463–487.
- [37] M. D. SMOOKE, R. MITCHELL, AND D. KEYES, *Numerical solution of two-dimensional axisymmetric laminar diffusion flames*, Combust. Sci. and Tech., 67 (1989), pp. 85–122.
- [38] V. VENKATKRISHNAN, *Newton solution of inviscid and viscous problems*, AIAA Journal, 27 (1989), pp. 885–891.
- [39] F. A. WILLIAMS, *Combustion Theory*, Benjamin/Cummings Publishing, Menlo Park, 1985.