Contemporary Mathematics
Volume 00, 0000

# Parallel Implicit Methods for Aerodynamics

X.-C. CAI, W. D. GROPP, D. E. KEYES, AND M. D. TIDRIRI

ABSTRACT. Domain decomposition (Krylov-Schwarz) iterative methods are natural for the parallel implicit solution of multidimensional systems of boundary value problems that arise, for instance, in aerodynamics. They provide good data locality so that even a high-latency workstation network can be employed as a parallel machine. Matrix-free (Newton-Krylov) methods are natural when it is unreasonable to compute or store a true Jacobian. We call their combination Newton-Krylov-Schwarz and report experimental progress on two algorithmic aspects: the use of a coarse grid in additive Schwarz preconditioning and the use of mixed discretization schemes in the (implicitly defined) Jacobian and its preconditioner. Two model problems in two-dimensional compressible flow are considered: the full potential equation, and the Euler equations.

## 1. Krylov-Schwarz Algorithms

Fully implicit linear solvers in aerodynamics allow more rapid asymptotic approach to steady states than time-explicit, approximate factorization, or relaxation solvers that hold the outer nonlinear iteration to small time steps. Nevertheless, the all-to-all data dependencies between the unknown fields in a fully implicit method have led to a resurgence of interest in less rapidly convergent methods in high-latency parallel environments. Resisting, we briefly overview two related efforts that lie along the route to parallel implicit computational aerodynamics. Though the governing equation formulations are mathematically very different – elliptic subsonic full potential and hyperbolic transonic Euler – a common implicit software core allows them to be treated together. Our ultimate

1991 *Mathematics Subject Classification.* Primary 65N55, 65N22, 76G25.

The work was supported in part by the National Science Foundation and the Kentucky EPSCoR Program under grant STI-9108764 (XCC); by the Office of Scientific Computing, U.S. Department of Energy, under Contract W-31-109-Eng-38 (WDG); by the National Science Foundation under contract number ECS-8957475, the State of Connecticut and the United Technologies Research Center (DEK); and by the National Aeronautics and Space Administration under NASA contract NAS1-19480 while three of the authors (XCC,DEK,MDT) were in residence at the Institute for Computer Applications in Science and Engineering.

This paper is in final form and no version of it will be submitted elsewhere.

©0000 American Mathematical Society
0000-0000/00 $1.00 + $.25 per page

interest is in applying Schwarzian domain decomposition techniques to industrial computations still being carried out in these (physically) primitive potential and Euler formulations, and in extending them to Navier-Stokes, for which implicit solvers are even more important. For a variety of reasons, industrial CFD groups are inclining towards the distributed network computing environment characterized by coarse to medium granularity, large memory per node, and very high latency, which creates a niche for domain decomposition methods.

Schwarz-preconditioned Krylov solvers for nonsingular linear systems, $Ax = b$, find the best approximation of the solution $x$ in a small-dimensional subspace that is built up from successive powers of the preconditioned matrix on the initial residual. Such systems, in which $A$ is a Jacobian matrix and $b$ is the nonlinear residual, arise from discretization and linearization of the governing PDEs. A variety of parallel preconditioners, whose inverse action we denote by $B^{-1}$, can be induced by decomposing the domain of the underlying PDE, finding an approximate representation of $A$ on each subdomain, inverting locally, and combining the results. Generically, we seek to approximate the inverse of $A$ by a sum of local inverses:

$$B^{-1} = \sum_k R_k^T A_k^{-1} R_k ,$$

where $R_k$ is a restriction operator that takes vectors spanning the entire space into the smaller dimensional subspace in which $A_k$ is defined.

The simplest domain decomposition preconditioner is block Jacobi, which can be regarded as a zero-overlap form of additive Schwarz [**2**]. The convergence rate of block Jacobi can be improved, at the price of a higher cost per iteration, with subdomain overlap and (for many problems) by solving an additional judiciously chosen coarse grid system.

## 2. Newton-Krylov Methods

Evaluation of the discrete residuals of $d$-dimensional compressible flow formulations requires a large number of arithmetic operations. (For instance, a $(d+2)$-dimensional eigendecomposition may be required at each grid point.) Their Jacobians, though block-sparse, have dense blocks and are usually an order of magnitude even more complex to evaluate, whether by analytical or numerical means. Hence, matrix-free Newton-Krylov methods, in which the action of the Jacobian is required only on a set of given vectors, are natural in this context. To solve the nonlinear system $f(u) = 0$, given $u^0$, let $u^{l+1} = u^l + \lambda^l \delta u^l$, for $l = 0, 1, \ldots$, until the residual is sufficiently small, where $\delta u^l$ approximately solves the Newton correction equation $J(u^l)\delta u^l = -f(u^l)$, and $\lambda^l$ is a damping parameter. The action of Jacobian $J$ on an arbitrary Krylov vector $w$ can be approximated by

$$J(u^l)w \approx \frac{1}{\epsilon}\left[f(u^l + \epsilon w) - f(u^l)\right] .$$

Finite-differencing with $\epsilon$ makes matrix-free methods potentially much more susceptible to finite word-length effects than ordinary Krylov methods [**5**]. Steady aerodynamics applications require the solution of linear systems that lack strong diagonal dominance, so a secondary goal of our investigation is to verify that properly-scaled matrix-free methods can be employed in this context. For brevity, details are deferred to a later paper. We simply note here that GMRES may have an advantage over other Krylov methods in the matrix-free context in that the vectors $w$ that arise in GMRES have unit two-norm, but may have widely varying scale in competing methods.

An approximation to the Jacobian can be used to precondition the Krylov process. Natural examples are: (1) the Jacobian of a related discretization that allows economical analytical evaluation of elements, (2) the Jacobian of a lower-order discretization, (3) a finite-differenced Jacobian computed with lagged values for expensive terms, and (4) domain-parallel preconditioners of the form

$$B^{-1} = R_0 J_{0,u^l}^{-1} R_0^T + \sum_{k=1}^{K} R_k J_{k,u^l}^{-1} R_k^T \ , \text{ where } J_{k,u^l} = \left\{ \frac{\partial f_i(u^l)}{\partial u_j} \right\}$$

is the Jacobian of $f(u)$ for $i$ and $j$ in subdomain $k$, and subscript "0" corresponds to a possible coarse grid. The Newton-Krylov-Schwarz method (case 4) can be combined with any other split-discretization technique (cases 1–3), in principle. Right preconditioning of the Jacobian with an operator $B^{-1}$ can be accommodated via

$$J(u^l)B^{-1}w \approx \frac{1}{\epsilon}\left[f((u^l + \epsilon B^{-1}w)) - f(u^l)\right].$$

## 3. Model Problems

For density $\rho$, velocity $\mathbf{v}$, specific internal energy $e$, and pressure $p$, the steady Euler equations of inviscid compressible flow are

$$\nabla \cdot (\rho \mathbf{v}) = 0, \quad \nabla \cdot (\rho \mathbf{v}\mathbf{v} + pI) = 0, \text{ and } \nabla \cdot ((\rho e + p)\mathbf{v}) = 0.$$

The full potential equation for velocity potential $\Phi$,

$$\nabla \cdot (\rho(||\nabla\Phi||)\nabla\Phi) = 0,$$

follows from the additional assumptions of irrotationality, $\mathbf{v} \equiv \nabla\Phi$, and isentropy, $\nabla(p/\rho^\gamma) \equiv 0$. The density is given in terms of the potential by

$$\rho = \rho_\infty \left(1 + \frac{\gamma - 1}{2} M_\infty^2 (1 - \frac{q^2}{q_\infty^2})\right)^{\frac{1}{(\gamma-1)}}$$

where $q = ||\nabla\Phi||$ and $M_\infty = q_\infty/a_\infty$. Here, $a$ is the sound speed, $q$ the flow speed, and $\infty$ refers to the freestream. When the flow is everywhere subsonic the full potential formulation fits within the monotone nonlinear elliptic framework of additive Schwarz methods [**1**]. For a simple non-lifting model problem of an

airfoil lying along the symmetry axis $y = 0$, we choose boundary conditions as follows:

- Upstream and Freestream: $\Phi = q_\infty x$ (zero angle of attack),
- Downstream: $\Phi_{,n} = q_\infty$,
- Symmetry: $\Phi_{,n} = 0$,
- On the parameterized airfoil with shape $y = f(x)$: $\Phi_{,n} = -q_\infty f'(x)$.

The farfield boundary conditions lead to inaccuracies if applied too near the airfoil, but our interest is in algebraic convergence rates.

Table 1 shows convergence performance for a fixed-size problem of $128 \times 128$ uniform cells with a fixed number of subdomains in an $8 \times 8$ array as the density of the unnested uniform coarse grid varies. Bilinear rectilinear elements are used for both coarse and fine grids, and bilinear interpolation for intergrid transfers. An overlap $2h$ is employed on each of the subdomains, which are solved exactly. $M_\infty$ is 0.1 and the airfoil is the scaled upper surface of a NACA0012. Nonlinear convergence is declared following a $10^{-3}$ relative reduction in the steady-state residual, which requires only three Newton steps independent of inner linear method. Inner iteration convergence is a relative residual reduction of $10^{-4}$. The Krylov solver used throughout this paper is GMRES [**7**], because of previous comparisons [**3**] with other modern Krylov solvers on the same problem class that showed CPU cost differences to be small and unsystematic when well-enough preconditioned that any of the methods were practical. Here, we restart GMRES every 20 iterations and precondition on the right, in order to keep the preconditioner out of the residual norm estimates used in the convergence test. Key observations from this example are: (1) even a modest coarse grid makes a significant improvement in an additive Schwarz preconditioner; (2) a law of diminishing returns sets in at roughly one point per subdomain; and (3) matrix-free "matvecs" degrade convergence as much as 15-20% in the less well-conditioned cases.

| Coarse Grid | $0 \times 0$ | $4 \times 5$ | $8 \times 9$ | $12 \times 13$ | $16 \times 17$ | $20 \times 21$ |
|---|---|---|---|---|---|---|
| Analytical | 177 | 35 | 28 | 27 | 24 | 21 |
| Matrix-free | 183 | 41 | 28 | 27 | 25 | 23 |

TABLE 1. Average number of GMRES steps per Newton step for full potential Newton-Krylov-Schwarz solver with varying coarse grid size.

Our Euler problem is a two-dimensional transonic airfoil flow modeled using an EAGLE-derivative code [**6**] that employs a finite volume discretization over a body-fitted coordinate grid. Only C-grids of $128 \times 16$ or $128 \times 32$ cells (from [**3**]) around a NACA0012 airfoil at an angle of attack of 1.25° and an $M_\infty$ of 0.8 are considered herein. To obtain a representative matrix/RHS pair on which to test the behavior of Euler Jacobians under Krylov-Schwarz, we first ran a demonstration case from [**6**] partway to convergence and linearized about the

resulting flow state. Following the defect correction practice of [**8**], a flux vector split scheme is employed for the implicit operators, and $f(u)$ itself is discretized by a flux difference split scheme. Characteristic variable boundary conditions are employed at farfield boundaries using an explicit, first-order accurate formulation. For a given granularity of decomposition, curvilinear "box" decompositions are generally better than curvilinear "strip" decompositions for this problem. Table 2 shows that the zero-overlap results are only slightly less convergent than the corresponding $h$-overlapped additive Schwarz results at high Courant-Friedrichs-Lewy (CFL) number, and that $h$-overlapped multiplicative Schwarz is significantly better, though the latter is a much less parallel algorithm. Though we have not yet experimented with a coarse grid in the Euler context, [**9**] shows that even a piecewise constant coarse grid operator substantially improves Krylov-Schwarz convergence rates in unstructured problems.

| Precond. | Block Jacobi | | Add. Schwarz | | Mult. Schwarz | |
|---|---|---|---|---|---|---|
| CFL No. | 1 | $10^2$ | 1 | $10^2$ | 1 | $10^2$ |
| $1 \times 1$ | 1 | 1 | 1 | 1 | 1 | 1 |
| $2 \times 2$ | 4 | 14 | 7 | 14 | 2 | 7 |
| $4 \times 4$ | 4 | 18 | 7 | 17 | 3 | 8 |
| $8 \times 8$ | 5 | 28 | 10 | 23 | 3 | 8 |

TABLE 2. Iteration counts for transonic flow Jacobians at local CFL numbers of 1 and $10^2$, for various preconditioners and decomposition into 4, 16, or 64 subdomains.

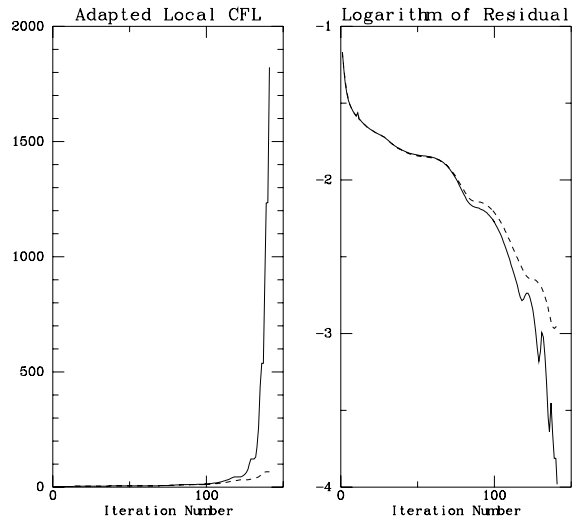


FIGURE 1. CFL and steady-state residual versus iteration count for defect correction and Newton-Krylov solvers.

To test the nonlinear matrix-free approach in a situation with four differently scaled components per gridpoint, we started over and approached the steady solution via a pseudo-transient continuation with an adaptively chosen local

CFL number, as described in [**3**]. Use of the baseline approximate factorization defect correction algorithm produces the dashed curves in Fig. 1. To obtain the solid curves, the explicitly available (Van Leer) flux vector split Jacobian ($J_{VL}$) is used to precondition the implicitly defined (Roe) flux difference split Jacobian ($J_R$) at each implicit time step. In matrix terms, the corrections $u$ are obtained as the approximate solutions of, respectively,

$$J_{VL}u = -f_R \quad \text{and} \quad (J_{VL})^{-1}J_R u = -(J_{VL})^{-1}f_R.$$

Unfortunately, in the retrofit of the existing code, transition to a full Newton method (CFL number approaching infinity) is precluded by explicit boundary conditions, but CFL number can be advanced, as shown in the figure, to $\mathcal{O}(10^3)$ with advantage.

Though space does not permit a meaningful discussion, we mention that both codes have been executed on an ethernet network of workstations using a package of distributed sparse linear system routines developed at Argonne National Laboratory by Gropp and Smith [**4**], with p4 as the data exchange layer. When exact solvers are used on each subdomain, speedups on a per iteration basis are seen on up to 16 processors, but exact solvers are an extreme case. As a serial preconditioner, global incomplete LU is superior to a Schwarz method using exact subdomain solvers.

## 4. Conclusions

By concentrating data dependencies locally, domain decomposition preconditioners exploit the two-level memory hierarchy of high-latency distributed-memory architectures. Low-communication zero or small overlaps between the preconditioner blocks are feasible with small convergence rate penalty, at least for intermediate granularities. The addition of a coarse grid has been shown to lead to major iteration count improvements, for a fixed problem size and algebraic residual reduction. Demonstrating the applicability of elliptic-based domain decomposition preconditioners to full potential and Euler problems is only a beginning. Further research will explore the limits of inconsistent preconditioners in matrix-free contexts, the cost versus benefits of the coarse grid solve in parallel contexts, and the relative tuning of inner and outer iteration convergence tolerances.

## References

1. X.-C. Cai and M. Dryja, *Domain Decomposition Methods for Monotone Nonlinear Elliptic Problems*, these proceedings, 1994.
2. M. Dryja and O. B. Widlund, *An Additive Variant of the Alternating Method for the Case of Many Subregions*, Courant Institute, NYU, TR 339, 1987.
3. W. D. Gropp, D. E. Keyes and J. S. Mounts, *Implicit Domain Decomposition Algorithms for Steady, Compressible Aerodynamics*, in Quarteroni et al., Sixth International Symposium on Domain Decomposition Methods for Partial Differential Equations (Como, Italy), AMS, Providence, 1994.

4. W. D. Gropp and B. F. Smith, *Simplified Linear Equations Solvers Users Manual*, ANL-93/8, Argonne National Laboratory, 1993.
5. P. R. McHugh and D. A. Knoll, *Inexact Newton's Method Solutions to the Incompressible Navier-Stokes and Energy Equations Using Standard and Matrix-Free Implementations*, AIAA Paper, 1993.
6. J. S. Mounts, D. M. Belk and D. L. Whitfield, *Program EAGLE User's Manual, Vol. IV – Multiblock Implicit, Steady-state Euler Code*, Air Force Armament Laboratory TR-88-117, Vol. IV, September 1988.
7. Y. Saad and M. H. Schultz, *GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems*, SIAM J. Sci. Stat. Comp. **7**(1986), 865–869.
8. J. L. Steger and R. F. Warming, *Flux Vector Splitting of the Inviscid Gasdynamics Equations with Applications to Finite-Difference Methods*, J. Comp. Phys. **40**(1981), 263–293.
9. V. Venkatakrishnan, *Parallel Implicit Unstructured Grid Euler Solvers*, AIAA Paper 94-0759, Reno, Nevada, January 1994.

Computer Science Department, University of Colorado, Boulder, CO, 80309.
*E-mail address*: cai@cs.colorado.edu

Argonne National Laboratory, 9700 South Cass Ave., Argonne, IL 60439-4844.
*E-mail address*: gropp@mcs.anl.gov

Computer Science Department, Old Dominion University, Norfolk, VA 23529-0162 and ICASE, MS 132C, NASA-LaRC, Hampton, VA 23681-0001
*E-mail address*: keyes@icase.edu

ICASE, MS 132C, NASA-LaRC, Hampton, VA 23681-0001
*E-mail address*: tidriri@icase.edu