Mixed-Precision Spectral Deferred Correction

Preprint

Ray W. S. Grout
National Renewable Energy Laboratory

Presented at the 27th International Conference on Parallel Computational Fluid Dynamics
Montreal, Canada
May 17–20, 2015

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Conference Paper
NREL/CP-2C00-64959
September 2015

Contract No. DE-AC36-08GO28308

NOTICE

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Available electronically at SciTech Connect http:/www.osti.gov/scitech

Available for a processing fee to U.S. Department of Energy and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
OSTI http://www.osti.gov
Phone: 865.576.8401
Fax: 865.576.5728
Email: reports@osti.gov

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312
NTIS http://www.ntis.gov
Phone: 800.553.6847 or 703.605.6000
Fax: 703.605.6900
Email: orders@ntis.gov

Cover Photos by Dennis Schroeder: (left to right) NREL 26173, NREL 18302, NREL 19758, NREL 29642, NREL 19795.

NREL prints on paper that contains recycled content.

MIXED-PRECISION SPECTRAL DEFERRED CORRECTION

RAY W.S. GROUT

National Renewable Energy Laboratory
Golden, CO, USA
e-mail: ray.grout@nrel.gov

Key words: Spectral deferred correction, mixed precision, combustion, CFD

Abstract. Convergence of spectral deferred correction (SDC), where low-order time integration methods are used to construct higher-order methods through iterative refinement, can be accelerated in terms of computational effort by using mixed-precision methods. Using ideas from multi-level SDC (in turn based on FAS multigrid ideas), some of the SDC correction sweeps can use function values computed in reduced precision without adversely impacting the accuracy of the final solution. This is particularly beneficial for the performance of combustion solvers such as S3D [6] which require double precision accuracy but are performance limited by the cost of data motion.

1 INTRODUCTION

We are currently seeing a transformative increase in computing power where first-principles and high-fidelity simulations of turbulent combustion can now realistically address practical combustion problems. In contrast to the voltage and frequency scaling that gave rise to the increases in computational power 'passively' over the last decade, the current avenues for increasing computational power also present new challenges. Performance gains are coming largely from increased parallelism—at the vector and core levels as well as the node level. Critically, the power density of new machines is rapidly becoming a design constraint. The cost to move data— either between memory and processing units or across the network — is now of critical importance to both the energy cost and time to complete a given operation [1].

At the same time, many challenges remain to be solved in combustion science. Current topical trends towards lightweight vehicles motivate moving to downsized, boosted engines that operate in non-traditional modes, possibly with novel ignition methods such as cool plasma ignition to extend the lean limit of operation. The increasing diversity of 'designer fuels' made available by sophisticated engineering of bioenergy feedstocks may enable new engine designs operating where traditional design rules are no longer applicable. On the gas turbine side, natural gas is an increasingly important fuel due to the US oil shale boom. Hydrogen enriched fuels such as syngas arise from many carbon sequestration and storage processes and are of significant interests to reduce the harmful effects of coal combustion. As the fuel streams evolve, so to do the operating modes; high-dilution,

lean combustion reduces flame temperatures and associated NOx formation but pose new stability challenges.

Combustion CFD is typically performance limited by memory bandwidth and can benefit significantly from reduced data movement and improved cache utilization, as well as improvements in floating-point performance, when calculations can be performed in single precision. However, especially for combustion problems, the accuracy of a double-precision solution is essential. Iterative deferred correction methods such as spectral deferred correction (SDC) have an advantage in being able to treat terms in the balance equations (e.g., chemical reaction vs. advection) with an appropriate implicit/explicit treatment while controlling splitting error, but the drawback to SDC is an increase in the number of function evaluations necessary to execute a time-step. Other iterative methods have been reformulated to use mixed precision with greatly improved performance (e.g., [5]); a similar formulation for SDC would alleviate the performance penalty.

In the next section, the SDC approach will be outlined in both standard and multi-level forms. Following the presentation of the current state of the art, a new formulation that uses single-precision evaluation of the function to be integrated will be proposed. This will be followed with a synthetic convergence test as well as a realistic combustion case.

2 Spectral Deferred Correction

SDC is an iterative method proposed by Dutt et al. [7] to iteratively improve a lower order approximation at spectral collocation points using an approximate update formulae. The original method has been advanced significantly, notably by Minion at colleagues (e.g., [11]). Several attractive features of SDC have been developed, notably: high order, mixed implicit/explicit treatment [4], operator splitting for reacting flows [13], incorporation of algorithmic resilience [9], parallel in time [8] and multi-level forms [14, 3].

To setup the SDC system, the timestep is broken into nodes at spectral quadrature points; using Gauss-Lobatto points conveniently includes the boundaries [12]. The solution $\mathbf{U}$ at each of the nodes can be written as:

$$\mathbf{U} = \Delta t \mathbf{Q}\mathbf{F}(\mathbf{U}) + \mathbf{U}_0, \tag{1}$$

where $\mathbf{Q}$ is the matrix of integration weights, $\mathbf{F}$ is the equation(s) begin integrated and $\mathbf{U}_0$ is the solution at the start of the timestep. The solution methodology is to determine an approximation to $\mathbf{U}$, typically by a low order (Euler) method and then to use a fixed point iteration to improve the solution. Given an approximation $\tilde{u}(t)$ to the initial value problem specified by:

$$u'(t) = F(u(t), t) \qquad u(t_a) = u_a, \tag{2}$$

a measure of the error in the approximation is:

$$E(\tilde{u}(t), t) = u_a + \int_{t_a}^{t} F(\tilde{u}(\tau), \tau) d\tau - \tilde{u}(t). \tag{3}$$

In comparison to the exact solution:

$$u(t) = u_a + \int_{t_a}^{t} f(u(\tau), \tau) d\tau, \tag{4}$$

the correction $\delta(t) = u(t) - \tilde{u}(t)$ is, algebraically:

$$\delta(t) = \int_{t_a}^{t} \left[F(\tilde{u}(\tau) + \delta(\tau), \tau) - F(\tilde{u}(\tau), \tau) \right] d\tau + E(\tilde{u}(t), t). \tag{5}$$

R. Speck et al. [14] reworked the SDC correction based on non-linear multigrid ideas so that some of the corrections can be replaced by coarse evaluations, coupling the levels together with an full approximation scheme (FAS) approach. The resulting SDC correction is unchanged from the conventional update given except for the addition of a FAS correction term τ that accounts for the non-commutative nature of the multigrid restriction (coarsening) operator. Approximating the integral in Eq. 5 with an explicit Euler method and using the error metric from Eq. 3, the resulting update formulae for the $m+1^{\text{st}}$ node in the k^{th} the multi-level SDC correction is:

$$u_{m+1}^{k+1} = u_m^{k+1} + \Delta t_m \left[F(u_m^{k+1}, t_m) - F(u_m^k, t_m) \right] + I_m^{m+1}(u^k) + \tau_m. \tag{6}$$

where:

$$I_m^{m+1}(u^k) \equiv \sum_{l=1}^{p} q_m^l F(u^k(t_l), t_l) \approx \int_{t_m}^{t_{m+1}} F(u^k(\tau), \tau) d\tau. \tag{7}$$

The multi-level approach adopted by Speck et al. is reproduced in Algorithm 1. Whereas Speck et al. considered coarsening in terms of reducing the spatial resolution, reducing the order of the spatial discritization and reducing the accuracy of implicit solves on the coarse levels, here we are interested in coarsening in terms of reduced precision.

Considering the coarsening in the MLSDC approach in terms of reduced-precision intermediate storage and evaluation leads to construction of a mixed precision SDC (MPSDC) algorithm. Although it is possible to construct an entire hierarchy of 'grids' using mixed precision (e.g., double, single, half...), this effort is limited to two levels: double (8 byte) and single (4 byte) precision. The update algorithm is shown in Algorithm 2 for a 4^{th} order formulation.

In Algorithm 2, five single precision function evaluations (1 of which can be reused from previous timesteps) and four double precision function evaluations per time step replace nine double precision function evaluations per time step for a conventional 4^{th} order SDC integration.

3 Synthetic convergence test of MPSDC algorithm

A synthetic test case used by Ascher and Petzold [2] is used to test the MPSDC convergence properties. The test equation is:

$$\frac{df}{dt} = -5tx^2 + \frac{5}{t} - \frac{1}{t^2}, \qquad y(1) = 1, \tag{8}$$

Algorithm 1: MLSDC Iteration for L levels from R. Speck et al.

Data: Initial $U^k_{1,0}$, function evaluations F^k_1 from the last iterate on the fine level
Result: Solution U^{k+1}_l and function evaluations F^{k+1}_l on all levels
$U^{k+1}_l, F^{k+1}_1 \leftarrow \text{SDCSweep}(U^k_l, F^k_l)$ *# Perform fine sweep*
if *fine level has converged* return *# Check convergence criteria*
 # Cycle from fine to coarse
for *l=1...L-1* **do**
  # Restrict, re-evaluate and save restriction
  for *m=0...M* **do**
    $U^k_{l+1,m} \leftarrow \text{Restrict}(U^{k+1}_{l,m})$
    $F^k_{l+1,m} \leftarrow \text{FEval}(U^{k+1}_{l+1,m})$
    $\tilde{U}^k_{l+1,m} \leftarrow U^k_{l+1,m}$
  end
  Compute FAS correction and sweep
  $\tau_{l+1} \leftarrow \text{FAS}(F^{k+1}_l, F^k_{l+1}, \tau_{l+1})$
  $U^{k+1}_{l+1}, F^{k+1}_{l+1} \leftarrow \text{SDCSweep}(U^k_{l+1}, F^k_{l+1}, \tau_{l+1})$
end
 # Cycle from coarse to fine
for *l=L-1...2* **do**
  # Interpolate coarse correction and re-evaluate **for** *m=0...M* **do**
    $U^{k+1}_{l,m} \leftarrow U^{k+1}_{l,m} + \text{Interpolate}(U^{k+1}_{l+1,m} - \tilde{U}^k_{l+1,m})$
    $F^{k+1}_{l,m} \leftarrow \text{FEval}(U^{k+1}_{l,m})$
  end
  $U^{k+1}_l, F^{k+1}_l \leftarrow \text{SDCSweep}(U^{k+1}_l, F^{k+1}_l, \tau_l)$
end
Return to finest level before next iteration
for *m=0...M* **do**
  $U^{k+1}_{l,m} \leftarrow U^{k+1}_{1,m} + \text{Interpolate}(U^{k+1}_{2,m} - \tilde{U}^k_{2,m})$
  $F^{k+1}_{1,m} \leftarrow \text{FEval}(U^{k+1}_{1,m})$
end

which has the analytic solution $y(t) = \frac{1}{t}$. The integration is performed over the time interval $[1, 25]$.

From Fig. 1, several features of the algorithms involving single-precision are of evident. Firstly, they all exhibit the expected 4^{th}-order convergence when the absolute error is sufficiently large to be represented in single precision. However, when step size is reduced eventually the error from the reduced precision representation dominates and the solution error plateaus even as the step size is reduced further. In the 'SDC-mp-simple' formulation, the first two correction passes are performed in single precision and the latter two are performed in double precision. Although this uses the same number of single-

Algorithm 2: MPSDC for 4^{th} order

$\vec{U}_c^1, \vec{F}_c^1 \leftarrow$ Predict_coarse($U_{f,0}$) *# Construct initial prediction*
$\vec{F}_f^1 \leftarrow$ FEval_fine($\vec{U}_c^1$) *# Perform fine sweep*
$\vec{U}_f^2, \vec{F}_f^2 \leftarrow$ SDCSweep($\vec{U}_c^1, \vec{F}_f^1$)
if *correction is less than SP significance* continue with DP
$U_c^2 \leftarrow$ Restrict(U_f^2) *# Restrict, re-evaluate and save restriction*
$F_c^2 \leftarrow$ FEval_coarse(U_c^2)
$\tau \leftarrow$ FAS($\vec{F}_f^2, \vec{F}_c^2$) *# Compute FAS correction and sweep*
$\vec{U}_c^3 \leftarrow$ SDCSweep_coarse(U_c^2, F_c^2, τ)
$U_f^3 \leftarrow U_f^2 + (U_c^3 - U_c^2)$ *# Correct fine*
$F_f^3 \leftarrow$ FEval_fine(U_f^3)
$U_f^4 \leftarrow$ SDCSweep(U_c^3, F_f^3) *# Final sweep on fine*

precision/double precision function evaluations (and hence same computational cost) as the MPSDC algorithm and is an improvement over using only single precision, the convergence rate is negatively impacted relative to the MPSDC and straight double-precision algorithms.

4 S3D and Combustion Test Case

A canonical combustion test case is set up in the combustion direct numerical simulation (DNS) code S3D as a practical demonstration. S3D [6] uses a method of lines approach to integrate the compressible reacting Navier-Stokes and species transport equations using a 4^{th} order, 6-stage explicit algorithm that requires 6 function evaluations per time step from the family developed by Kennedy et al. [10]. Previous exploration has shown that, as typical in combustion problems, double precision accuracy is necessary for sufficient accuracy, however, truncation of work arrays to single precision is often acceptable. For the chemistry evaluation a large number of scratch variables typically results in a working set size that spills out of nearby cache/registers so reductions in working set size can be expected to have an outsize performance benefit. For the comparison described below, a single precision version of the 'right hand side' evaluator was constructed that, as closely as possible, duplicated the operations in the original *RHS* function that includes the physical processes (advection, diffusion, reaction) that form the time derivative of the conserved variables (momentum, energy, density, species mass fractions, etc.). Some optimization transformations applied to the chemical rate evaluation routine that are effective with double precision were not possible with the single precision version due to overflow, notably applying some operations on log-transformed quantities. In these cases analytically equivalent but non-overflowing expressions were used.

The test case is a one-dimensional ignition problem where a Gaussian temperature 'hot spot' is placed within a homogenous mixture of stochiometric C_2H_4 and air. The quantity of interest is the ignition delay time: the time that passes before the temperature

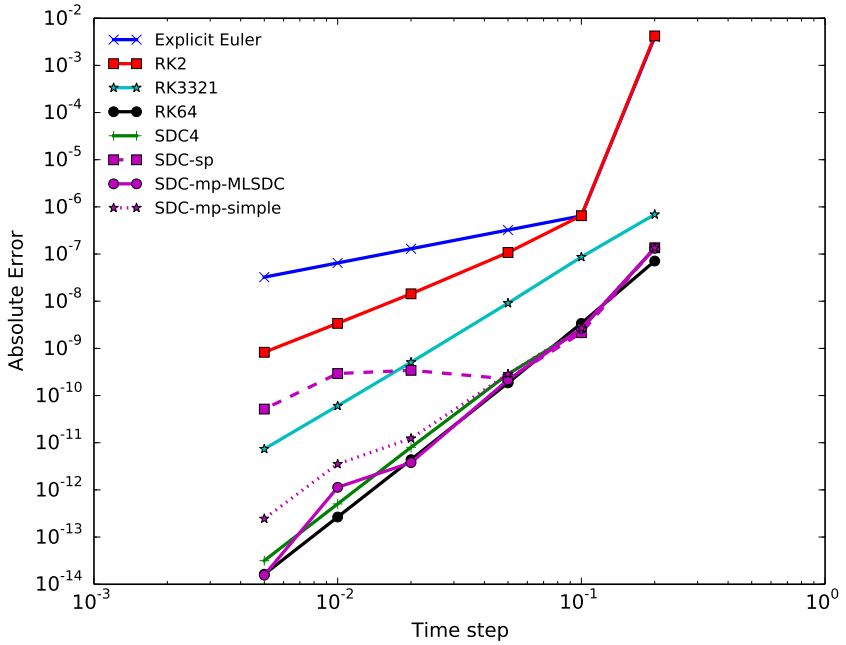

Figure 1: Convergence rates of MPSDC compared to several other algorithms: explicit Euler, second, third, and fourth order Runge-Kutta methods (RK2, RK33321, RK64, respectively) from Kennedy et al. [10], a conventional 4th order explicit SDC method (SDC4), a 4th order explicitly SDC method using single-precision function evaluation (SDC-sp), the MPSDC method (SDC-mp-MLSDC), and an alternate mixed precision form omitting the FAS correction (SDC-mp-simple).

exceeds the initial peak temperature. A feature of this problem is the slow accumulation of chemical intermediates before rapid temperature rise. Significant non-linearity in the coupling between chemical reaction and diffusion gives rise to a complicated relationship between the global error in the quantity of interest and the local error in the conserved quantities. Generally, as long as the local error is small enough, the global error is insensitive to the local error; however, if the local error grows too large the global error increases disproportionally. The result is the behavior observed in Figure 2, where the baseline double precision solution is shown alongside the (indistinguishable) MPSDC solution and a single-precision solution.

When the time step is reduced to $\Delta t = 5ns$, a similar behavior to that shown in Figure 2 is obeserved: the single precision solution has significant global error accumulation whereas the mixed precision solution has an error sufficiently controlled so that the global error is of the same order as the local error. However, for a timestep of $\Delta t = 5ns$, the change in the solution for each timestep is small relative to the error intrinsic in the single precision solution. For this situation, the 'coarse' correction does not make a useful contribution. This suggests that when the correction is larger than single-precision tolerance, it is effective to do iterates in single-precision to accelerate solution, however, it is pru-

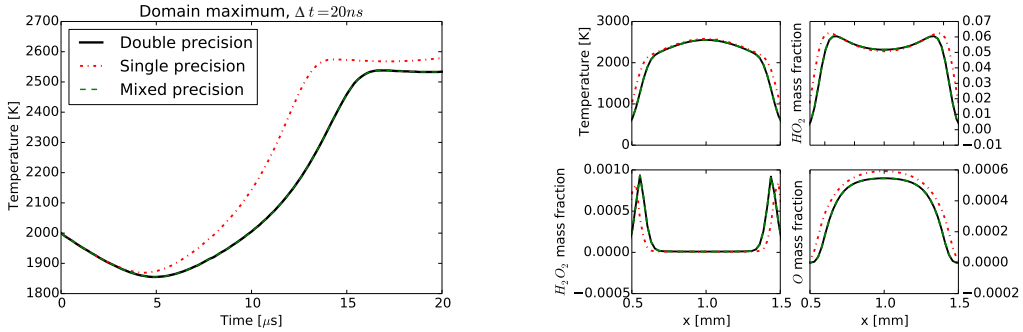

Figure 2: Combustion test problem - 1D ignition, $\Delta t = 20ns$

dent to do at least one double-precision correction pass to check the correction magnitude and abandon the MPSDC approach if the error is too small to be represented in single-precision. Further, while the single-precision iterations are computationally cheaper than double-precision corrections, including them does impact the convergence rate of subsequent double precision iterations. This is evident in Figure 3, where the local error for each correction pass is shown for the conventional and MPSDC algorithms. Whereas the local error decreases with each iterate in the conventional formulation, with the MPSDC formulation when the error approaches the level of single-precision representation applying a single precision correction appears to negatively impact the solution improvement resulting from the following double-precision iteration.

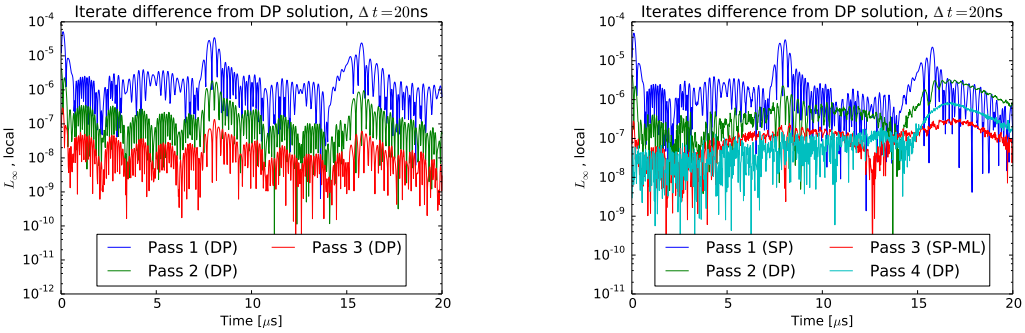

Figure 3: Local errors computed relative to DP SDC solution for each timestep

5 CONCLUSIONS

By reducing precision to coarsen the solution representation, the established multi-level spectral deferred correction can be used to formulate a mixed precision time advance algorithm. For a 'drop-in replacement' for the 4^{th} order integrator in the S3D combustion code, this allows half of the eight double precision evaluations of the time derivative to be replaced by single precision evaluations. This suggest a 25% reduction in memory traffic and 25% in increase FLOPS purely from data size considerations with larger benefits possible due to more efficient cache utilization. A complication is that the mixed

precision formulation introduces a potential inconsistency between the coarse problem and the fine solution. This limits the most significant benefit to situations where the size of the correction exceeds single precision accuracy yet the target local error is in the double-precision range. This would be expected where the stability limits exceed accuracy dictated Δt. Fortunately, the solution is easily monitored and the algorithm transparent when there is no benefit. For the combustion test case considered here, the mixed precision algorithm retains sufficient global error control whereas straight single precision does not with same Δt. This leads to effective use of mixed precision and significant computational savings.While the final solution accuracy does not depend directly on the intermediate reduced precision solution, in future work it would be beneficial to develop reduced precision formulations that improve consistency with the ultimate double precision solution to improve convergence rates.

REFERENCES

[1] The opportunities and challenges of exascale computing. Technical report, Summary Report of the ASCAC Subcommittee, 2010.

[2] Uri M. Ascher and Linda R. Petzold. *Computer methods for ordinary differential equations and differential-algebraic equations*. SIAM, 1998.

[3] W. Auzinger. Defect correction for nonlinear elliptic difference equations. *Numer. Math.*, 51:199–208, 1987.

[4] A. Bourlioux, A. T. Layton, and M. L. Minion. Higher-order multi-implicit spectral deferred correction methods for problems of reacting flow. *J. Comput. Phys.*, 189:351–376, 2003.

[5] Alfredo Buttari et al. Exploiting mixed precision floating point hardware in scientific computations. In *High Performance Computing Workshop*, pages 19–36, 2006.

[6] J. H. Chen, A. Choudhary, B. de Supinski, M. DeVries, E. R. Hawkes, S. Klasky, W. K. Liao, K. L. Ma, J. Mellor-Crummey, N. Podhorszki, R. Sankaran, S. Shende, and C. S. Yoo. Terascale direct numerical simulations of turbulent combustion using S3D. *Comput. Sci. Disc.*, 2:015001, 2009.

[7] A. Dutt, L. Greengard, and V. Rokhlin. Spectral deferred correction methods for ordinary differential equations. *BIT*, 40(2):241–266, 2000.

[8] M. Emmett and M. L. Minion. Towards an efficient parallel in time method for partial differential equations. *Commun. Appl. Math. Comp. Sci.*, 7(1):105–132, 2012.

[9] RW Grout, H Kolla, ML Minion, and JB Bell. Achieving algorithmic resilience for temporal integration through spectral deferred corrections. *arXiv preprint arXiv:1504.01329*, 2015.

[10] C. A. Kennedy, M. H. Carpenter, and R. M. Lewis. Low-storage, explicit Runge-Kutta schemes for the compressible Navier-Stokes equations. *Appl. Numer. Math.*, 35:177–219, 2000.

[11] A. T. Layton and M. L. Minion. Conservative multi-implicit spectral deferred correction methods for reacting gas dynamics. *J. Comput. Phys.*, 194(2):697–715, 2004.

[12] M. L. Minion. Semi-implicit spectral deferred correction methods for ordinary differential equations. *Commun. Math. Sci.*, 1(3):471–500, 09 2003.

[13] A. Nonaka et al. A deferred correction coupling strategy for low Mach number flow with complex chemistry. *Combust. Theory & Model.*, pages 1–36, 2012.

[14] Robert Speck, Daniel Ruprecht, Matthew Emmett, Michael Minion, Matthias Bolten, and Rolf Krause. A multi-level spectral deferred correction method. *BIT Numerical Mathematics*, pages 1–25, 2014.