

INTERWEAVING PFASST AND PARALLEL MULTIGRID

M. L. MINION*, R. SPECK^{†‡}, M. BOLTEN[§], M. EMMETT*, AND D. RUPRECHT[‡]

Abstract. The parallel full approximation scheme in space and time (PFASST) introduced by Emmett and Minion in 2012 is an iterative strategy for the temporal parallelization of ODEs and discretized PDEs. As the name suggests, PFASST is similar in spirit to a space-time Full Approximation Scheme (FAS) multigrid method performed over multiple time-steps in parallel. However, since the original focus of PFASST has been on the performance of the method in terms of time parallelism, the solution of any spatial system arising from the use of implicit or semi-implicit temporal methods within PFASST have simply been assumed to be solved to some desired accuracy completely at each sub-step and each iteration by some unspecified procedure. It hence is natural to investigate how iterative solvers in the spatial dimensions can be interwoven with the PFASST iterations and whether this strategy leads to a more efficient overall approach. This paper presents an initial investigation on the relative performance of different strategies for coupling PFASST iterations with multigrid methods for the implicit treatment of diffusion terms in PDEs. In particular, we compare full accuracy multigrid solves at each sub-step with a small fixed number of multigrid V-cycles. This reduces the cost of each PFASST iteration at the possible expense of a corresponding increase in the number of PFASST iterations needed for convergence. Parallel efficiency of the resulting methods is explored through numerical examples.

Key words. Parallel in time, PFASST, multigrid

1. Introduction. The past decade has seen a growing interest in the development of parallel methods for temporal integration of ordinary differential equations (ODEs), particularly in the context of temporal strategies for partial differential equations (PDEs). One factor fueling this interest is related to the evolution of supercomputers during this time. Since the end of the exponential increase in individual processors speeds, increases in supercomputer speeds have been mostly due to increases in the number of computational cores, and current projections suggest that the first exaflop computer will contain on the order of a billion cores [14]. The implication of this trend is that increasing concurrency in algorithms is essential, and in the case of time-dependent PDE simulations, the use of space-time parallelism is an attractive option.

Time-parallel methods have a long history dating back at least to the work of Nievergelt [26]. In the context of space-time multigrid, Hackbusch noted already in 1984 that relaxation operators in parabolic multigrid can be employed on multiple time steps simultaneously [15]. The 1997 review article by Burrage [7] provides a summary of early work on the subject. More recently, the parareal method proposed in 2001 [22] has renewed interest in temporal parallelization methods. In 2012 the parallel full approximation scheme in space and time (PFASST) was introduced by Emmett and Minion [9], and performance results for PFASST using space-time parallelization with hundreds of thousands of cores can be found in [31, 28].

The PFASST algorithm is based on a type of deferred corrections strategy for ODEs [8], with corrections being applied on multiple time steps in parallel. As such, there are similarities between parareal and PFASST (see [25, 24]). On the other hand, the parallel efficiency of PFASST depends on the construction of a hierarchy of space-time discretizations, hence there are also similarities between PFASST and space-

*Center for Computational Sciences and Engineering, Lawrence Berkeley National Laboratory, USA

[†]Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany

[‡]Institute of Computational Science, Università della Svizzera italiana, Lugano, Switzerland

[§]Department of Mathematics, Bergische Universität Wuppertal, Germany

time multigrid methods. However, in the original papers on PFASST, the solution of any spatial systems due to implicit time-stepping was assumed to be found to full precision since the interest was on the temporal accuracy and efficiency of the methods. From this point of view, PFASST is an iterative solver in the time direction but not in the spatial dimensions. This is, in a sense, orthogonal to the traditional use of multigrid solvers within PDE methods, where multigrid is used to iteratively solve spatial equations and the time direction is not iterative. One of the main goals of this paper is to investigate the use of both spatial and temporal iterative methods utilizing PFASST and multigrid.

To be more specific, the iterative strategy within the PFASST algorithm is derived from the method of Spectral Deferred Corrections [8], which is a variant of the defect and deferred correction methods developed in the 1960s [36, 27, 33, 2]. One advantage of SDC methods is that it is straightforward to construct methods of very high order of accuracy by iteratively applying low-order methods to a series of correction equations. This flexibility has been exploited to construct higher-order semi-implicit (IMEX) and multi-implicit SDC methods [23, 4, 20], as well as multirate SDC methods [5]. Such methods are very difficult to construct using traditional Runge-Kutta or linear-multistep approaches.

The main disadvantage of SDC methods is that they have a high cost per time step in the sense of the number of function evaluations (explicit or implicit) required per time step. When high accuracy is desired, this cost is generally offset by the use of relatively large time steps compared to lower-order methods for the same given accuracy. Nevertheless, approaches to reducing the cost of each SDC iteration have been proposed, including those generally referred to as *multi-level* SDC methods (MLSDC). The main idea in MLSDC is to perform some SDC iterations on coarser discretizations of the problem, and methods that coarsen the temporal discretization, the spatial discretization, and the order of the spatial approximation have recently been investigated [30]. SDC iterations are then performed on the hierarchy of levels in much the same way as V-cycles in traditional multigrid. The PFASST algorithm can be considered a time parallel version of MLSDC.

Using implicit SDC methods (as well as popular methods like backward Euler, trapezoid rule, Diagonally Implicit Runge-Kutta, or Backward Difference Formula) for grid-based PDEs simulations requires the solution of implicit equations in the spatial domain. If an iterative solver is used for these equations in the context of SDC, one advantageous result is that the initial guess for each solve becomes better as the SDC iterations converge. This raises the possibility of reducing the cost of SDC or MLSDC iterations further by limiting the number of spatial iterations used for each implicit solve rather than requiring each to be done to some specified tolerance. A recent paper explores this possibility in the context of SDC methods [32]. We refer to variants in which spatial solves are done only incompletely with a prepended capital “I” (ISDC, IMLSDC, IPFASST), resulting in the interweaving of spatial and temporal iterations. This paper adopts the ISDC strategy for MLSDC and PFASST and explores the performance of both IMLSDC and IPFASST.

As a starting point, we focus here on the the linear test problem resulting from a finite difference approximation of the heat equation. Variants of this example are considered in early papers on space-time multigrid [17, 18], waveform relaxation [34, 13], block-parallel methods [1], as well as more recent papers on parallel space-time multigrid [12, 11]. In all the cases above, a second-order finite-difference approximation is employed in space and first- or second-order methods used in time. Although this is

not the optimal setting for the PFASST algorithm, we present results using second-order space-time discretization to compare with other published results. We also show how, for a given accuracy, using fourth-order methods in space and/or time results in significant computational savings compared to second-order methods.

The remainder of this paper is organized as follows. We first present the SDC algorithm for linear problems in a compact notation that can be interpreted as matrix iterations and highlight the similarities between SDC iterations and classical relaxation schemes. We then discuss how IMLSDC and IPFASST are constructed. In Section 3, we provide the relevant analysis of SDC methods as a relaxation operator for parabolic problems. In Section 4 we examine the scaling of the IPFASST method in terms of problem size and number of parallel time steps, consider the effect of limiting the number of multigrid V-cycles per implicit solve, and provide numerical examples comparing PFASST with IPFASST. Finally, in Section 4.2 we provide timing comparisons for three-dimensional examples scaling to hundreds of thousands of cores.

2. Collocation and Spectral Deferred Corrections. This section briefly describes the methods used later. In Section 2.1, the collocation formulation for initial value problems is first reviewed. SDC as an iterative solver for a collocation solution is discussed briefly in Section 2.2. A compact notation of SDC for linear problems is given, and its interpretation as relaxation is discussed. The extension of SDC to multi-level SDC (MLSDC) is outlined in Section 2.3, including a number of strategies to coarsen the representation of the problem in order to reduce the overall cost. The possibility of solving linear systems approximately on all levels leads to so-called "inexact spectral deferred corrections" (ISDC) and their multi-level counterparts IMLSDC and IPFASST.

2.1. The Collocation Formulation. This paper concerns methods for the solution of initial value ODEs, particularly those arising from the spatial discretization of PDEs through the method of lines technique. To set notation, consider the scalar ODE of the form

$$(2.1) \quad y'(t) = f(t, y(t))$$

$$(2.2) \quad y(t_0) = y_0,$$

where $y(t), y_0 \in \mathbf{C}$ and $f : \mathbb{R} \times \mathbf{C} \rightarrow \mathbf{C}$. Similarly for systems of equations

$$(2.3) \quad \mathbf{y}'(t) = \mathbf{f}(t, \mathbf{y}(t))$$

$$(2.4) \quad \mathbf{y}(t_0) = \mathbf{y}_0,$$

where $\mathbf{y}(t), \mathbf{y}_0 \in \mathbf{C}^D$ and $\mathbf{f} : \mathbb{R} \times \mathbf{C}^D \rightarrow \mathbf{C}^D$. An equivalent form is given by the Picard equation, which for (2.3) is

$$(2.5) \quad \mathbf{y}(t) = \mathbf{y}_0 + \int_{t_0}^t \mathbf{f}(\tau, \mathbf{y}(\tau)) d\tau.$$

Since the goal here is to describe numerical methods for a given time step, consider the time interval $t \in [t_n, t_{n+1}]$ with $\Delta t = t_{n+1} - t_n$. Define the set of points t_m for $m = 0, \dots, M$ to be quadrature nodes scaled to $[t_n, t_{n+1}]$, so that $t_n = t_0 < t_1 < t_2 < \dots < t_M = t_{n+1}$. Given a scalar function $g(t)$, define the approximations to the integral of $g(t)$ over the intervals $[t_0, t_m]$ by choosing the coefficients $q_{m,i}$ such that

$$\int_{t_0}^{t_m} g(\tau) d\tau \approx \Delta t \sum_{i=0}^M q_{m,i} g(t_i).$$

The coefficients $q_{m,i}$ that give the highest order of accuracy given the points t_m can be derived using standard techniques. The quadrature rules used here have the property that $q_{m,0} \equiv 0$ (see [21] for a discussion of different choices of quadrature rules for SDC methods). Let the matrix Q of size $(M+1) \times (M+1)$ be composed of the coefficients $q_{m,i}$. The first row of Q contains only zeros. Q will be referred to here as the integration matrix.

Consider for the time being scalar equations, the integration matrix Q can be used to discretize (2.5) directly over the interval $[t_n, t_{n+1}]$. Let $y_0 = y_n \approx y(t_n)$ be the initial condition, and define $y_m \approx y(t_m)$ by

$$(2.6) \quad y_m = y_n + \Delta t \sum_{i=0}^M q_{m,i} f(t_i, y_i) \quad \text{for } m = 0, \dots, M.$$

We can write this equation in a compact form by denoting the vector of unknowns by $Y = [y_0 \dots y_M]^T$, and the vector of function values $F(Y) = [f(y_0, t_0) \dots f(y_M, t_M)]^T$. Furthermore, let Y_0 be the $(M+1) \times 1$ column vector with each entry equal to y_0 . Then (2.6) is equivalent to

$$(2.7) \quad Y = Y_0 + \Delta t Q F(Y).$$

This coupled (generally nonlinear) equation can be solved directly for the values Y , resulting in a collocation scheme for the ODE.

It should also be noted that the collocation form (2.7) corresponds to a fully implicit Runge-Kutta (IRK) method given by the Butcher tableau

$$(2.8) \quad \begin{array}{c|c} \mathbf{c} & \mathbf{A} \\ \hline & \mathbf{b} \end{array},$$

where $\mathbf{c}$ are the nodes t_m scaled to the unit interval, $\mathbf{A} = Q$, and the vector $\mathbf{b}$ corresponds to the last row of Q .

For the simplest linear scalar equation where $f(y, t) = \lambda y$ in (2.1), the collocation formulation given in (2.7) becomes

$$(2.9) \quad Y = Y_0 + \lambda \Delta t Q Y,$$

or in more standard matrix form

$$(2.10) \quad (I_{M+1} - \lambda \Delta t Q) Y = Y_0,$$

with I_{M+1} being the $(M+1) \times (M+1)$ identity matrix. In the case of a linear system of equations $\mathbf{f}(t, \mathbf{y}) = A \mathbf{y}$, where A is a $D \times D$ matrix, the Picard equation becomes

$$(2.11) \quad \mathbf{y}(t) = \mathbf{y}_n + \int_{t_n}^t A \mathbf{y}(\tau) d\tau.$$

where the integration is done term by term. To discretize (2.11), the solution is the vector

$$(2.12) \quad \mathbf{Y} = [y_0^1 \dots y_0^D \dots y_M^1 \dots y_M^D]^T$$

of length $D(M+1)$. Here the subscripts m correspond to the quadrature nodes in time and superscripts j correspond to the component of the solution vector.

To apply the numerical quadrature matrices to $\mathbf{Y}$, let $\mathbf{Q} = Q \otimes I_D$ where I_D is the $D \times D$ identity matrix. Likewise let $\mathbf{A} = I_{M+1} \otimes A$, and $\mathbf{Y}_0$ be the vector of length $(M+1)D$ consisting of M copies of the $\mathbf{y}_0$. Then the analogous form of (2.7) is

$$(2.13) \quad \mathbf{Y} = \mathbf{Y}_0 + \Delta t \mathbf{Q} \mathbf{A} \mathbf{Y}$$

or

$$(2.14) \quad (\mathbf{I} - \Delta t \mathbf{Q} \mathbf{A}) \mathbf{Y} = \mathbf{Y}_0,$$

where $\mathbf{I} = I_{M+1} \otimes I_D$.

2.2. Single level SDC. SDC can be understood as a preconditioned fixed point iteration to solve (2.14), which avoids treating the full system directly by computing a series of corrections node by node (see e.g. [19, 35]). Note that originally SDC was introduced in a different way, as a variant of earlier deferred and defect correction schemes [36, 27] designed to achieve a fixed order of accuracy for a fixed number of correction sweeps [8]. The multi-level SDC methods (and PFASST) described below move away from the idea of a fixed number of iterations in favor of the convergence toward the collocation (or IRK) solution.

Here we provide a short review of the SDC method. For more details see [8, 23]. Let the superscript k denote the numerical approximation at the k th SDC iteration. Using backward-Euler as the base method, one iteration of SDC can be written as a sweep through the quadrature nodes, successively updating the solutions. In the case of scalar equations, this takes the form

$$(2.15) \quad y_{m+1}^{k+1} = y_m^{k+1} + \Delta t_m (f(t_{m+1}, y_{m+1}^{k+1}) - f(t_{m+1}, y_m^k)) + Q_m^{m+1}(Y^k)$$

for $m = 0, \dots, M-1$ where

$$(2.16) \quad Q_m^{m+1}(Y^k) \approx \int_{t_m}^{t_{m+1}} f(\tau, y^k(\tau)) d\tau.$$

The values of $Q_m^{m+1}(Y^k)$ can easily be constructed using the integration matrix $\mathbf{Q}$

$$(2.17) \quad Q_m^{m+1}(Y^k) = \Delta t \sum_{i=0}^M (q_{m+1,i} - q_{m,i}) f(t_i, y_i^k).$$

Here we consider only implicit SDC methods, even though one particularly attractive feature of SDC is the flexibility to use different base methods in order to create e.g. high-order implicit-explicit or multi-rate methods [23, 4, 16, 5].

2.2.1. Compact notation. We will now present a compact notation of the SDC iterations. Given the points $t_m \in [t_n, t_{n+1}]$ as discussed above, let $\Delta t_m = t_m - t_{m-1}$, and let $\gamma_m = \Delta t_m / \Delta t$, where again $\Delta t = t_{n+1} - t_n$. We begin by defining the lower-order integration matrix that has the same dimensions as Q

$$(2.18) \quad Q_I = \begin{bmatrix} 0 & 0 & \cdots & 0 & 0 \\ 0 & \gamma_1 & \cdots & 0 & 0 \\ \cdot & \cdot & \cdots & 0 & 0 \\ 0 & \gamma_1 & \cdots & \gamma_{M-1} & 0 \\ 0 & \gamma_1 & \cdots & \gamma_{M-1} & \gamma_M \end{bmatrix}.$$

Using the same notation as in Section 2.1, M steps of backward Euler for (2.1) can be written as

$$(2.19) \quad Y = Y_0 + \Delta t Q_I F(Y).$$

This differs from the collocation formulation (2.7) only in that Q_I replaces Q .

Since Q_I is lower triangular with non-zero diagonal entries, (2.19) can be solved sequentially for the values y_m , where each value requires an implicit equation to be solved of the form

$$y_{m+1} = y_m + \Delta t_m f(t_{m+1}, y_{m+1}).$$

Using the matrix-vector notation above, one SDC sweep for (2.3) can be compactly written as

$$(2.20) \quad \mathbf{Y}^{k+1} = \mathbf{Y}_0 + \Delta t \mathbf{Q}_I (\mathbf{F}(\mathbf{Y}^{k+1}) - \mathbf{F}(\mathbf{Y}^k)) + \Delta t \mathbf{Q} \mathbf{F}(\mathbf{Y}^k)$$

or

$$(2.21) \quad \mathbf{Y}^{k+1} = \mathbf{Y}_0 + \Delta t \mathbf{Q}_I \mathbf{F}(\mathbf{Y}^{k+1}) + \Delta t (\mathbf{Q} - \mathbf{Q}_I) \mathbf{F}(\mathbf{Y}^k).$$

2.2.2. SDC for Linear Problems. The compact formulas derived above can be recast as matrix-vector operations when the governing equation is linear. In the case of the linear system, $\mathbf{y}' = \mathbf{A}\mathbf{y}$, (2.20) simplifies to

$$(2.22) \quad \mathbf{Y}^{k+1} = \mathbf{Y}_0 + \Delta t \mathbf{Q}_I (\mathbf{A} \mathbf{Y}^{k+1} - \mathbf{A} \mathbf{Y}^k) + \mathbf{Q} \mathbf{A} \mathbf{Y}^k$$

or

$$(2.23) \quad (\mathbf{I} - \mathbf{Q}_I \mathbf{A}) \mathbf{Y}^{k+1} = \mathbf{Y}_0 + \Delta t (\mathbf{Q} - \mathbf{Q}_I) \mathbf{A} \mathbf{Y}^k.$$

As in the case of backward Euler, this system of equations can be solved by substepping, requiring the solution of

$$(2.24) \quad (\mathbf{I} - \Delta t_m \mathbf{A}) \mathbf{y}_{m+1}^{k+1} = \mathbf{y}_0 - \Delta t_m \mathbf{A} \mathbf{y}_{m+1}^k + \mathbf{Q}_m^{m+1} (\mathbf{Y}^k).$$

This equation involves the inversion of the same operator that arises from a backward Euler method with a modified right-hand side. Note that as the SDC iterations converge in k , an increasing good approximation to this solution is provided by $\mathbf{y}_{m+1}^k$. This fact has two relevant implications when iterative methods are employed to solve the system. First, when considering the total cost of one time-step of SDC, the reduced cost of the implicit solves as k increases should be taken into consideration. Second, and more relevant for this paper, is that an iterative method need not solve (2.24) to full precision at each iteration. Instead, a fixed number of iterations could be performed during each SDC iteration, or each could be done so that the residual is reduced by some set tolerance. In the numerical studies presented in Section 4, multigrid methods are employed for solving (2.24) and we investigate limiting the number of multigrid V-cycles instead of solving to full precision.

Finally, note that in the SDC iterations, it is straightforward to compute the residual in terms of the solution to (2.13), specifically

$$(2.25) \quad \mathbf{r}^k := \mathbf{Y}_0 - (\mathbf{I} - \Delta t \mathbf{Q} \mathbf{A}) \mathbf{Y}^k.$$

In Section 4, the norm of $\mathbf{r}^k$ is used to monitor convergence of the different choices of methods.

2.2.3. SDC as a relaxation. For a linear problem, it is possible to write an SDC sweep described above as a relaxation operator applied to the linear collocation equation

$$(2.26) \quad (\mathbf{I} - \Delta t \mathbf{Q} \mathbf{A}) \mathbf{Y} = \mathbf{Y}_0.$$

As in classical iteration methods for linear systems, we look for a decomposition of the matrix $(\mathbf{I} - \Delta t \mathbf{Q} \mathbf{A}) = \mathbf{M} - \mathbf{K}$ such that $\mathbf{M}$ is relatively less expensive to invert than $(\mathbf{I} - \Delta t \mathbf{Q} \mathbf{A})$. Then a classical relaxation scheme based on this splitting would be

$$(2.27) \quad \mathbf{Y}^{k+1} = \mathbf{M}^{-1} \mathbf{K} \mathbf{Y}^k + \mathbf{M}^{-1} \mathbf{Y}_0.$$

Choosing

$$(2.28) \quad \mathbf{M} = \mathbf{I} - \Delta t \mathbf{Q}_I \mathbf{A} \quad \text{and} \quad \mathbf{K} = \Delta t (\mathbf{Q} - \mathbf{Q}_I) \mathbf{A}$$

produces the SDC sweep as given by (2.23), hence inverting $\mathbf{M}$ can be done by substepping on the SDC nodes solving the appropriate version of (2.24).

In the next section we carry this analogy further by introducing variants of SDC that utilize multiple levels of resolution as in classical multigrid methods.

2.3. Multi-level SDC Methods (MLSDC). The goal of MLSDC methods is that by introducing a hierarchy of levels from fine to coarse, some of the expensive fine level correction sweeps can be replaced with sweeps on coarser levels, so that the runtime required for convergence of SDC iterations is reduced (see recent results in [30]). An space-time FAS term is employed when forming coarser level equations, which is the difference between the coarsening of the latest fine function values and the coarse level function applied to a coarsening of the latest fine solution. The result of including the FAS term is that the accuracy on the coarse level approaches that of the fine level as the MLSDC iterations converge (see for example [6]). Different strategies for reducing the cost of SDC sweeps on the coarser levels are explored in [30], including using a lower-order spatial discretization, using a spatial mesh with fewer points, or performing only incomplete implicit solves. The structure of an MLSDC level hierarchy is sketched in Figure 1. For a detailed explanation of the method including performance results we refer to [30]. MLSDC is the basic building block for the time-parallel PFASST method summarized in Section 2.4.

2.3.1. Inexact MLSDC (IMLSDC). In [32], another approach to reducing the cost of SDC methods is discussed, so called "inexact spectral deferred corrections" or ISDC. Here, implicit solves in SDC sweeps (2.15) are computed only approximately, e.g. by a fixed small number of V-cycles of multigrid. While this strategy makes each SDC sweep less expensive, it can increase the number of SDC iterations required to reduce the residual to a given threshold. If the increase in iterations is not too large, ISDC can in total be faster than SDC. In [32] it is demonstrated that ISDC can save up to 50% of the V-cycles required by regular SDC while converging to the same tolerance. It is also shown that the reason this strategy works is the increasingly accurate initial guesses provided for the iterative solver by SDC (see the discussion toward the end of Section 2.2.2).

In regular MLSDC, doing incomplete solves on the coarse levels is already discussed and explored in [30] as a means to reduce the cost of coarse level sweeps. However, the ISDC strategy can also be incorporated into MLSDC, resulting in an




FIG. 1. A MLSDC level hierarchy over a time-step $[T^n, T^{n+1}]$. The upper fine level (red) uses five quadrature nodes, the medium level (blue) three and the bottom coarsest level (black) two. Also sketched is the possibility of using a spatial discretization with fewer degrees of freedom on the coarser levels: Depicted is an 8×8 point mesh on the fine level, a 4×4 point mesh in the middle, and a 2×2 point mesh on the coarse level. A detailed description of MLSDC can be found in [30]. Figure reprinted from [28].

IMLSDC method. Here, all implicit solves are done only approximately, including the finest level. Ideally, IMLSDC would save on runtime compared to ISDC, just as MLSDC does compared to SDC. However, the results presented in Section 4 suggest that this is not necessarily the case. However, MLSDC serves as the basis of the parallel-in-time method PFASST, and hence IMLSDC could be used in parallel as well. In Section 4 we demonstrate that IPFASST (see Section 2.4.1) can provide a significant reduction of runtime by exploiting temporal concurrency. A detailed description of MLSDC including pseudo code can be found in [30], and we refer the reader there for details.

2.4. The Parallel Full Approximation Scheme in Space and Time. The parallel full approximation scheme in space and time (PFASST), introduced in [24, 9] is an iterative time-parallel method for PDEs that has similarities to both the parareal method and space time multigrid methods. In PFASST each time step is assigned its own processor or, if combined with spatial parallelization, its own communicator (for a more detailed explanation of the latter case, see [31]). PFASST can be considered a time-parallel extension of MLSDC, where after an initialization procedure, MLSDC iterations are performed on multiple time steps in parallel with updates to initial conditions being passed between processors as each SDC sweep is completed.

PFASST typically starts by distributing the initial value over all time ranks as an initial guess which is then refined by a number of sweeps on the coarse level, where processors handling time steps later in time do more sweeps (this is usually referred to as predictor phase). After completing the predictor phase, each processor starts with its own MLSDC iterations while, after each sweep, sending forward an updated initial value for the current level to the processor handling the next time step, cf. Figure 2. Blocking communication is required on the coarsest level only, so that minimal synchronicity between the different MLSDC iterations is required, see [10]. Benchmarks for PFASST in large-scale parallel simulations illustrate how PFASST can extend strong scaling limits [31] or improve utilization of large parallel installations in comparison to codes utilizing only spatial parallelism [28].

A detailed description of PFASST including a sketch of the algorithm in pseudo code can be found in [9], and we refer the reader there for more details.

2.4.1. IPFASST. Just as incomplete implicit solves can be used in MLSDC, they can also be used in PFASST. Essentially, each processor now performs IMLSDC iterations instead of iterations with full solves on the fine level. Apart from that,




FIG. 2. PFASST performs *MLSDC* iterations concurrently on multiple time steps. For simplicity, only two time steps $[T^n, T^{n+1}]$ and $[T^{n+1}, T^{n+2}]$ are sketched here. After each sweep, an updated initial value on the corresponding level is sent forward to the processor handling the next time step. In the setup above, processor P_n handles time step $[T^n, T^{n+1}]$ and sends forward updates to processor P_{n+1} which handles $[T^{n+1}, T^{n+2}]$. Figure reprinted from [28].

IPFASST proceeds the same as PFASST. In Section 4, performance of IPFASST will be studied through numerical examples.

3. Analysis of SDC on the Linear Model Problem. In Section 2.2.3, the analogy between a single SDC sweep and a classical relaxation scheme derived from a splitting of the linear system is presented. In this section we examine the effect of a single SDC sweep on the scalar linear model problem

$$(3.1) \quad y' = \lambda y$$

$$(3.2) \quad y(0) = 1.$$

In the context of parabolic problems, the relative set of λ is the negative real axis, and we are interested in how one sweep of SDC reduces the error in the discrete approximation.

In this case, the compact form of the SDC sweep from (2.23) becomes

$$(3.3) \quad (I - \lambda \Delta t Q_I) Y^{k+1} = Y_0 + \lambda \Delta t (Q - Q_I) Y^k.$$

The corresponding solution to the collocation equation satisfies

$$(3.4) \quad (I - \lambda \Delta t Q) Y = Y_0,$$

which implies

$$(3.5) \quad (I - \lambda \Delta t Q_I) Y = Y_0 + \lambda \Delta t (Q - Q_I) Y.$$

In words, the SDC sweep acts as the identity operator for the exact solution Y .

Subtracting equation (3.5) from (3.3) gives

$$(3.6) \quad (I - \lambda \Delta t Q_I) (Y^{k+1} - Y) = \lambda \Delta t (Q - Q_I) (Y^k - Y)$$

or

$$(3.7) \quad Y^{k+1} - Y = (I - \lambda \Delta t Q_I)^{-1} \lambda \Delta t (Q - Q_I) (Y^k - Y).$$

Hence we can study the converge properties of SDC methods for the scalar model problem by examining the largest eigenvalue of the matrix

$$(3.8) \quad (I - \lambda \Delta t Q_I)^{-1} \lambda \Delta t (Q - Q_I).$$




FIG. 3. Damping factors for second- and fourth-order SDC methods used in Section 4.

This matrix depends in general on the form of the sub-stepping encapsulated in the approximate quadrature matrix (here backward Euler in Q_I), the number and type of quadrature nodes, and the product $\lambda\Delta t$. Here we examine the cases corresponding to the second- and fourth-order methods used in the numerical results in Section 4. These correspond to uniform quadrature nodes and a quadrature rule which does not use the left-hand endpoint in the quadrature rule (see [21] for a discussion of different quadrature rules). These methods are used together in the fourth-order PFASST example as well, where the second-order method is the time coarsened version of the fourth-order method. This choice requires 2 implicit sub-steps for second order and 4 for fourth order and hence is not “spectral” in the sense of using Gaussian quadrature rules. This choice does however provide good damping factors for low-order methods¹.

Figure 3 shows the maximum magnitude of the eigenvalues for both second- and fourth-order cases as a function of $\lambda\Delta t$. Clearly in both limits $\lambda\Delta t \rightarrow 0$ and $\lambda\Delta t \rightarrow -\infty$, the damping factor goes to zero. This behavior differs from the traditional analysis of point relaxation for elliptic equations like the Laplace equation where high-frequency eigen-components are damped more rapidly than low frequency components for the classical relaxation schemes. This feature is the reason why multigrid methods provide a tremendous speedup compared to relaxation alone. In the context of SDC sweeps, Figure 3 demonstrates that although MLSDC with temporal coarsening may provide some increase in efficiency compared to higher-order SDC, SDC alone will still converge rapidly in the stiff limit for this problem.

For any linear problem, one can represent one iteration of PFASST or IPFASST as a matrix vector multiplication on a vector composed of the solution at each SDC node within each parallel time step. Examining the eigenvalues of this matrix would then give an indication of how PFASST iterations would converge for a given choice of parameters. One could then examine, for a given linear PDE, different numbers and types of nodes; types of sweeps and quadrature rules; number of levels, refinement

¹A recent paper by Weiser [35] studies the optimization of more general sub-stepping rules for SDC methods. These ideas are not pursued used here despite their apparent promise.

factors in space and time, and type of interpolation and restriction between levels; and the size and number of parallel time-steps. In the case of IPFASST, the type of relaxation, the order of interpolation and restriction, and the number of V-cycles per implicit-solve could also be varied². A paper presenting a systematic study of the linear convergence of PFASST and IPFASST is in preparation.

4. Numerical Results. In this section, numerical results illustrating some features of the IPFASST method are presented on the model problem given by the linear heat equation. First, results analyzing the convergence behavior of IPFASST are reported. Strong scaling timing results for three-dimensional problems using space-time parallelism are then presented in Section 4.2.

Despite the popularity of the heat equation as a test case for space-time parallel methods (e.g. [17, 18, 34, 13, 1, 12]), there are several reasons why it is not a particularly good test case for time parallelization. First, when considering time-parallel methods, one would like to show how the method scales as the number of parallel time steps grows. Hence for many time steps, one must either choose a relatively long interval of integration or a relatively small time step. If a long time interval is chosen, the solution will decay to a value close to zero which complicates any discussion of accuracy and convergence. If on the other hand, very small time steps are chosen, then the temporal accuracy of the method will likely far exceed the spatial accuracy, which brings into question why such a small time step is being considered. Finally, it is important to note that the performance of parallel-in-time methods applied to the linear heat equation may not be indicative of performance on other problems that are not strictly parabolic, have dynamic spatial features, or have nontrivial boundary conditions. Despite these drawbacks, we will investigate the performance of IPFASST on this test case. The main motivation for doing so is to provide a comparison of the PFASST method with other published space-time parallel methods using this example.

4.1. One-Dimensional Convergence Studies. In this section we consider the accuracy and convergence of IPFASST, including the dependence on the number of V-cycles. For reasons addressed below, we use here the simple model problem consisting of the one-dimensional heat equation with homogeneous Dirichlet boundary conditions. The equation in a general form prescribed on the space-time domain $[0, L] \times [0, T]$ is given by

$$\begin{aligned} u_t &= \nu u_{xx} \\ u(x, 0) &= u_0(x) \\ u(0, t) &= u(L, t) = 0. \end{aligned}$$

Choosing the initial conditions

$$u_0(x) = \sin(k\pi x/L),$$

gives the exact solution

$$u(x, t) = e^{-\nu(k\pi/L)^2 t} u_0(x).$$

²It is also worth noting that the use of non-uniform sub-steps in SDC removes the equivalence between Fourier and eigenmode decomposition in the time direction.

Using the method of lines and a second-order finite-difference approximation of the spatial derivatives gives the linear system of ODEs

$$u_i'(t) = \nu \frac{u_{i-1} - 2u_i + u_{i+1}}{\Delta x^2},$$

where $u_i(t) \approx u(i\Delta x, t)$ for $i = 1 \dots N-1$, $\Delta x = L/N$, and $u_0 = u_N = 0$. The exact solution of the systems of ODEs given initial conditions

$$u_i(0) = \sin(k\pi x_i/L),$$

is

$$u_i(t) = e^{-\nu d(k)t} u_0(x_i),$$

where

$$d(k) = \frac{-2 + 2 \cos(k\pi \Delta x/L)}{\Delta x^2}.$$

Note that for k even moderately large, the solution decays very rapidly to zero. In the following one-dimensional examples, we choose $k = L = \nu = T = 1$, which means the solution decays to approximately 5×10^{-5} during the time interval.

One of the convenient features of SDC methods is that it is simple to construct higher-order methods simply by increasing the number of quadrature nodes being used. Even restricting the discussion to second-order spatial discretizations, using a first-order or second-order temporal integration method is very inefficient. As a simple demonstration of this, consider the following numerical experiment. We apply SDC methods of different orders to (4.1) with $k = \nu = L = T = 1$. Specifically we consider SDC methods of order 1, 2, 4, and 8 where the number of SDC nodes per time step corresponds to the formal order. We chose $\Delta x = 1/128$ and compute the L_∞ error of the solution at the final time $T = 1$ compared to the exact solution of the PDE and the exact solution of the discrete ODE for various values of $\Delta t = 1/N_t$ for $N_t = 2^p$ with p ranging from 2 to 12. In addition, we compute the residual at the final time step as in (2.25).

The results are displayed in Figure 4. Several comments can be made. Most obviously, as the order of the time integration increases, the number of time steps required to reduce the PDE error to the minimum possible given the spatial resolution goes down dramatically. For the eighth-order method, two time steps of size 0.5 yields a smaller error than 128 time steps of the second-order method. Of course, higher-order methods require more work per time step due to the increase in substeps on the SDC nodes, but as shown in Section 4.2 higher-order discretizations in space and time lead to much reduced computational times for problems in three dimensions. Note also that the residual and error in the discretized ODE continue to converge to zero well past the minimum error due to the spatial discretization. This implies that a judicious choice of residual tolerance is needed for iterative approaches like IPFASST to avoid wasted iterations.

In the PFASST algorithm, parallel efficiency depends on the availability of a hierarchy of space-time discretizations as in MLSDC, where coarsening in the temporal direction is achieved by reducing the number of nodes used in the underlying SDC method on each processor. Obviously for a second-order temporal discretization, only one coarsening step in time is possible leading to a coarse level based on the backward-Euler method. In the spatial directions, one way to reduce the work on PFASST levels




FIG. 4. Convergence in error and residual for serial SDC methods of different orders for the 1-D heat equation.

is to reduce the order of the spatial approximation on coarser levels, however, this is only possible when higher-order spatial approximation is being used on the finest level. Hence the options for coarsening in PFASST when restricted to second-order discretizations are more limited than for higher-order methods.

4.1.1. Comparison of IPFASST With Different Number of V-cycles. In this numerical experiment, we compare the convergence of IPFASST with different numbers of prescribed V-cycles for each approximate linear solve at each sub-step. A standard geometric multigrid method is used for these tests with linear interpolation, full-weighting restriction, and two pre- and two post-smoothing sweeps using a Jacobi smoother. IPFASST will only be more efficient than PFASST if the reduced cost of fewer V-cycles is not offset by an increase in the number of iterations required for convergence.

We use here $N_x = N_t = 128$ with 128 parallel time steps and the same parameters as in the previous examples, namely $k = \nu = L = T = 1$. Figure 5 shows the convergence of IPFASST in terms of ODE error, PDE error and residual versus the number of V-cycles used for the approximate implicit solve in every sub-step. The residual decays faster as the number of V-cycles is increased from 1 to 3, but using more than 3 V-cycles does not yield further improvement. In contrast, both the ODE and the PDE error are less effected by the number of V-cycles. Using two V-cycles per solve gives virtually the same behavior as using 10. We stress again that these




FIG. 5. Comparison of the convergence of IPFASST using different numbers of V-cycles per solve.

results for the linear heat equation may not hold for PDEs of different mathematical type (see [32] for some preliminary results on advection diffusion equations).

4.1.2. Convergence of IPFASST With Weak Scaling. An important motivation for the study of time-parallel methods is the “trap of weak scaling” given current supercomputer evolution where cores counts are increasing, but processor speed is not. For a given application, if the spatial resolution is increased as core counts increase, the cost of spatial operators will remain close to constant assuming good weak scaling of the spatial operations. However, the time step size will eventually need to also be reduced, either because of stability constraints and/or to match the increased spatial accuracy (depending on whether explicit or implicit methods are used). This means that more time steps are necessary for the same simulation time, and therefore, the total runtime will increase even with perfect spatial scaling unless some parallelism in the time direction is employed. For time-parallel methods to be effective, it is necessary that the performance does not deteriorate as resolution is refined. Therefore, in this test we consider how the convergence of the IPFASST iterations scales with the problem size and number of parallel time steps.

Three runs are performed with $N_x = N_t = 32, 64$, and 128 again choosing $k = \nu = L = T = 1$. The number of parallel time-steps in IPFASST is equal to the number of time steps N_t . We use three levels with coarsening by a factor two in space and second-order finite-difference approximations on all levels. On the coarsest level, the collocation rule corresponds to first-order backward Euler. On the middle and




FIG. 6. Convergence of IPFASST in terms of ODE error, PDE error, and residual in the final time step for increasing resolution in both space and time.

finest level, we use 3 uniform nodes corresponding to a second-order collocation rule. Based on the previous experiments, implicit solves are approximated by two V-cycles with two pre- and post-smoothing steps with a Gauss-Seidel smoother.

Figure 6 shows how IPFASST converges for the different resolutions in terms of the ODE error (left), the PDE error (middle) and the residual (right). Errors are reported for the last time step. As the resolution increases in space or time, the error level decreases up to some minimum level. Depending on which discretization error is dominant, this level is either the discretization error of the collocation rule or of the spatial discretization, cf. the discussion in [30, Sect. 3.2]. This level is reached at iteration 5 for $N = 32$ and at iteration 3 for $N = 64$ and $N = 128$. Therefore, increasing the resolution and with it the number of parallel time steps does not increase the number of iterations required by IPFASST. Note that this does not cover the case where the number of concurrently computed steps is increased while the time step size is kept constant in order to compute over a longer time interval.

4.2. Three-dimensional Strong Scaling Studies. To compare the computational cost of IPFASST with serial methods we consider the 3D heat equation

$$(4.1) \quad u_t(\mathbf{x}, t) = \nu \Delta u(\mathbf{x}, t), \quad \mathbf{x} \in \Omega, \quad 0 \leq t \leq T$$

on the unit cube $\Omega = [0, 1]^3$ with $T = 1.0$, initial conditions

$$(4.2) \quad u(\mathbf{x}, 0) = \sin(\pi x) \sin(\pi y) \sin(\pi z),$$

and homogeneous Dirichlet boundary conditions. We choose $\nu = \frac{1}{3}$ so that the solution decays at the same rate as the previous one-dimensional tests. The Laplacian

is discretized with either a second- or fourth-order finite difference stencil, and either ISDC or IPFASST is used to solve the resulting initial value problem. In all cases an implicit Euler sub-stepping is used for the SDC sweeps, and PMG [3] is used for parallel multigrid V-cycles. Simulations are run on the IBM BlueGene/Q installation JUQUEEN at Jülich Supercomputing Centre. Timing results are displayed in Figure 7, and the specifications of the different runs are provided below.

4.2.1. Problem and Method Specifications. ISDC, IMLSDC and IPFASST are run in three different configurations, which all give approximately identical errors measured in the infinity norm against the analytical solution of (4.1). The tolerance for the residual of ISDC and IPFASST is set to 10^{-9} , which in all three configurations results in an error of about 1.5×10^{-7} . PMG uses a tolerance of 10^{-12} whenever a full solve is performed, and a stalling criterion that stops the iteration if the new residual is not smaller than 75% of the previous one. Two levels are used in IMLSDC and IPFASST, with fourth-order spatial interpolation and point-wise restriction in both space and time.

In all runs, approximate implicit solves consisting of two PMG V-cycles are used in the SDC sweeps. PMG V-cycles use two pre- and two post-smoothing sweeps consisting of JOR red-black smoothers. Linear interpolation and full-weighting restriction are used in the V-cycles.

Second-order spatial and second-order temporal discretization. (Marked by circles in Figure 7). For these runs, ISDC and the fine level of IMLSDC/IPFASST use two quadrature nodes corresponding to the midpoint of the time step and the right-hand value t_{n+1} . Therefore the method converges to a second-order collocation scheme. The coarse levels in IMLSDC and IPFASST correspond to backward Euler. The spatial mesh has $N = 128^3$ points and 128 time steps are performed. The coarse level uses a 64^3 point spatial mesh. For ISDC, runs are performed with the number of cores used by PMG varying between 16 and 32,768. For IMLSDC and IPFASST, the number of cores for PMG is fixed to 4,096 and the number of parallel time-ranks varied between 1 (for IMLSDC) up to 32, for a total of $4,096 \times 32 = 131,072$ cores. On average, ISDC requires about 3.5 iterations and IMLSDC about 3.7 for convergence. The last time-rank in IPFASST requires between about 3.4 (for two parallel time-ranks) and 4.0 iterations (for 32 time-ranks).

Second-order spatial and fourth-order temporal discretization. (Marked by diamonds in Figure 7). Here, ISDC and the fine level of IPFASST use 4 quadrature nodes, corresponding to fourth-order collocation. The coarse level of IPFASST is the same as the second-order runs above. Only 24 time steps are required to achieve the same overall error because of the higher-order temporal discretization. The spatial mesh remains as described above. As before, scaling of ISDC is measured using 16 up to 32,768 cores for PMG while the number of cores for PMG in IMLSDC and IPFASST is fixed at 4,096. Because here only 24 time steps have to be performed, the number of time-ranks is varied only from 1 up to 24, for a total maximum number of cores of $24 \times 4,096 = 98,304$ cores. IMLSDC needs an average of 3.8 iterations, and the last time-rank in IPFASST between 4.0 (for two time-ranks) up to 6.0 (for 24 time-ranks).

Fourth-order spatial and fourth-order temporal discretization. (Marked by squares in Figure 7). As above, 24 time-steps are used with a fourth-order temporal discretization. $N = 32^3$ spatial points suffice to maintain the same error as above. On the coarse level, IPFASST uses only $N = 16^3$ points and a second-order finite-difference stencil. Here, because much fewer spatial degrees of freedom are needed, scaling of ISDC is




FIG. 7. Run times for the 3D heat equation depending on the total number of cores used for different configurations of ISDC and IPFASST. The first entry for each configuration of IPFASST corresponds to a run with only a single-time rank, that is IMLSDC. All setups result in an error of about 1.5×10^{-7} , making a comparison of their runtimes meaningful.

measured only using 16 to 512 cores in PMG. IMLSDC and IPFASST use 64 cores for PMG and, as before, up to 24 time-ranks. The maximum total number of cores here is therefore only $24 \times 64 = 1,536$.

4.2.2. Results. Figure 7 displays the run times for the different configurations of ISDC and IPFASST laid out above. In addition, Table 1 lists parallel speedup and efficiency (speedup divided by number of processors) provided by IPFASST compared to the ISDC counterpart with the same number of spatial processors for PMG. Note that the first marker for each IPFASST line corresponds to a run with only a single time-rank, which is IMLSDC. For each of the three setups, IMLSDC by itself is somewhat slower than ISDC using the same number of cores for PMG due to a slight increase in the number of iterations required by IMLSDC and the overhead of coarsening and interpolation between SDC levels. This is in contrast to the SDC/MLSDC method using full solves on the fine level as studied in [30], where MLSDC could significantly reduce run times compared to SDC. This means that the most efficient parallel variant for this test case (namely IPFASST) is not a direct parallelization of the most efficient serial variant (namely ISDC). Hence while the use of incomplete solves reduces the computational cost for both ISDC and IPFASST, it actually decreases the parallel efficiency of IPFASST because we have to compare to ISDC. We should also note that ISDC is not necessarily the most efficient serial method for this problem, but we use it as a comparison to the parallel performance of IPFASST based on ISDC sweeps.

As already observed in [29] for the PFASST method, it is more efficient to begin space-time parallelism with IPFASST using fewer spatial cores for PMG as for when the parallel speedup saturates. At the limit of spatial scaling, the spatial coarsening within PFASST does not make the coarse levels sweeps much cheaper than the fine

2 nd / 2 nd order			2 nd / 4 th order			4 th / 4 th order		
N_p	Speedup	Efficiency	N_p	Speedup	Efficiency	N_p	Speedup	Efficiency
1	0.8	–	1	0.9	–	1	0.7	–
2	1.2	60.0 %	2	1.4	70.0 %	2	1.0	–
4	2.1	53.5 %	4	2.6	65.0 %	4	1.9	47.5 %
8	3.5	43.5 %	8	4.6	57.5 %	8	3.2	40.0 %
16	5.6	35.0 %	12	5.9	49.2 %	12	4.6	38.3 %
32	8.6	26.9 %	24	7.2	30.0 %	24	6.6	27.5 %

TABLE 1

Additional speedup provided by IPFASST compared to the ISDC run with the same number of spatial processors for PMG. N_p indicates the number of parallel time steps.

level and therefore reduces the parallel efficiency of PFASST or IPFASST.

While the second/second-order and the second-/fourth-order versions of ISDC and IPFASST scale approximately equally well, the higher-order methods result in shorter run times. For fully fourth-order ISDC, spatial scaling is of course significantly worse than for the second-order spatial methods because the size of the problem is drastically smaller. However, no matter how many processors are used for the second-order version, the fourth-order ISDC is always significantly faster. The same is true for IPFASST. While all three configurations of IPFASST scale approximately equally well, the fourth-order version requires significantly fewer cores to achieve the same runtime as the second-order version. The smallest time-to-solution is about 2.5 seconds provided by fourth-order IPFASST using a total of 1,536 cores. While the mixed second/ fourth-order version is only slightly slower at 2.7 seconds, it requires 98,304 cores to achieve this runtime. We stress again that these results may not translate to other problems, particularly where higher-order methods are not appropriate due to a lack of smoothness in the solutions.

5. Discussion and Outlook. The two main goals for this paper are (1) to investigate combining spatial and temporal iterative strategies for space-time parallelization and (2) to provide some data on the performance of the resulting IPFASST method applied to the linear heat equation for comparison with other published results. The IPFASST algorithm introduced here interweaves the iterations of a spatial multigrid solver (PMG) with the temporal iterations in the SDC methods to achieve space-time parallelism, where the resulting hybrid space-time iterations are constructed by considering the spatial and temporal dimensions independently. Numerical results suggest that reducing the number of V-cycles for implicit spatial solves in the IPFASST can be done without a significant impact on the convergence of the time-parallel iterations (see e.g. Fig. 5). The scaling results shown in Fig. 7 demonstrate that incorporating temporal parallelization as in IPFASST can extend strong scaling and further reduce the time-to-solution when spatial parallelism is close to saturation, assuming more resources are available.

One interesting result suggested by the numerical experiments is that although MLSDC might be more efficient than SDC for the heat equation, it seems that ISDC is more efficient than IMLSDC. The difference between ISDC and IMLSDC is essentially only the order in which SDC relaxation sweeps and multigrid V-cycles are performed on various levels, and hence it is likely that even more efficient variations using a more general ordering of space-time relaxations could be found. A careful analysis of different variations of the methods presented here, and the extension of the linear analysis in Sect. 3 to PFASST and IPFASST is currently underway. As mentioned, given the number of different options already possible, finding optimal configurations

may be difficult, and more importantly these optimal configurations are probably problem-dependent.

6. Acknowledgments. Authors Minion and Emmett were supported by the Applied Mathematics Program of the DOE Office of Advanced Scientific Computing Research under the U.S. Department of Energy under contract DE-AC02-05CH11231. Speck and Ruprecht acknowledge support by the Swiss National Science Foundation (SNSF) under the lead agency agreement through the project “ExaSolvers” (SNF-145271) within the Priority Programme 1648 “Software for Exascale Computing” (SPPEXA) of the Deutsche Forschungsgemeinschaft. Computing time on JUQUEEN at the Jülich Supercomputing Centre (JSC) was provided by project HWU12.

REFERENCES

- [1] P. AMODIO AND L. BRUGNANO, *Parallel solution in time of odes: some achievements and perspectives*, Applied Numerical Mathematics, 59 (2009), pp. 424 – 435.
- [2] K. BÖHMER, P. HEMKER, AND H. J. STETTER, *The defect correction approach*, in Defect Correction Methods. Theory and Applications, K. Böhmer and H. J. Stetter, eds., Springer-Verlag, 1984, pp. 1–32.
- [3] M. BOLTEN, *Evaluation of a multigrid solver for 3-level Toeplitz and circulant matrices on Blue Gene/Q*, in NIC Symposium 2014, K. Binder, G. Münster, and M. Kremer, eds., John von Neumann Institute for Computing, 2014, pp. 345–352. (to appear).
- [4] ANNE BOURLIOUX, ANITA T. LAYTON, AND MICHAEL L. MINION, *High-order multi-implicit spectral deferred correction methods for problems of reactive flow*, Journal of Computational Physics, 189 (2003), pp. 651 – 675.
- [5] ELIZABETH L. BOUZARTH AND MICHAEL L. MINION, *A multirate time integrator for regularized Stokeslets*, Journal of Computational Physics, 229 (2010), pp. 4208 – 4224.
- [6] WILLIAM L. BRIGGS, VAN EMDEN HENSON, AND STEVE F. MCCORMICK, *A Multigrid Tutorial*, SIAM, 2000.
- [7] KEVIN BURRAGE, *Parallel methods for ODEs*, Advances in Computational Mathematics, 7 (1997), pp. 1–3.
- [8] ALOK DUTT, LESLIE GREENGARD, AND VLADIMIR ROKHLIN, *Spectral deferred correction methods for ordinary differential equations*, BIT Numerical Mathematics, 40 (2000), pp. 241–266.
- [9] M. EMMETT AND M. L. MINION, *Toward an efficient parallel in time method for partial differential equations*, Communications in Applied Mathematics and Computational Science, 7 (2012), pp. 105–132.
- [10] MATTHEW EMMETT AND MICHAEL L. MINION, *Efficient implementation of a multi-level parallel in time algorithm*, in Domain Decomposition Methods in Science and Engineering XXI, vol. 98 of Lecture Notes in Computational Science and Engineering, Springer International Publishing, 2014, pp. 359–366.
- [11] R. D. FALGOUT, S. FRIEDHOFF, TZ. V. KOLEV, S. P. MACLACHLAN, AND J. B. SCHRODER, *Parallel time integration with multigrid*, SIAM Journal on Scientific Computing, 36 (2014), pp. C635–C661.
- [12] S FRIEDHOFF, R D FALGOUT, T V KOLEV, S MACLACHLAN, AND J B SCHRODER, *A multigrid-in-time algorithm for solving evolution equations in parallel*, in Presented at: Sixteenth Copper Mountain Conference on Multigrid Methods, Copper Mountain, CO, United States, Mar 17 - Mar 22, 2013, 2013.
- [13] MARTIN J. GANDER AND ANDREW M. STUART, *Space-Time Continuous Analysis of Waveform Relaxation for the Heat Equation*, SIAM Journal on Scientific Computing, 19 (1998), pp. 2014–2031.
- [14] AL GEIST, *Paving the roadmap to exascale*, SciDAC Review, 16 (2010), pp. 53–59.
- [15] W. HACKBUSCH, *Parabolic multi-grid methods*, Computing Methods in Applied Sciences and Engineering, VI, (1984), pp. 189–197.
- [16] THOMAS HAGSTROM AND RUHAI ZHOU, *On the spectral deferred correction of splitting methods for initial value problems*, Communications in Applied Mathematics and Computational Science, 1 (2006), pp. 169–205.
- [17] G. HORTON AND R. KNIRSCH, *A time-parallel multigrid-extrapolation method for parabolic partial differential equations*, Parallel Computing, 18 (1992), pp. 21 – 29.

- [18] G. HORTON AND S. VANDEWALLE, *A space-time multigrid method for parabolic partial differential equations*, SIAM Journal on Scientific Computing, 16 (1995), pp. 848–864.
- [19] JINGFANG HUANG, JUN JIA, AND MICHAEL MINION, *Accelerating the convergence of spectral deferred correction methods*, Journal of Computational Physics, 214 (2006), pp. 633 – 656.
- [20] ANITA T. LAYTON AND MICHAEL L. MINION, *Conservative multi-implicit spectral deferred correction methods for reacting gas dynamics*, Journal of Computational Physics, 194 (2004), pp. 697 – 715.
- [21] ANITA T. LAYTON AND MICHAEL L. MINION, *Implications of the choice of quadrature nodes for picard integral deferred corrections methods for ordinary differential equations*, BIT Numerical Mathematics, 45 (2005), pp. 341–373.
- [22] J.-L. LIONS, Y. MADAY, AND G. TURINICI, *A "parareal" in time discretization of PDE's*, Comptes Rendus de l'Académie des Sciences - Series I - Mathematics, 332 (2001), pp. 661–668.
- [23] MICHAEL L. MINION, *Semi-implicit spectral deferred correction methods for ordinary differential equations*, Communications in Mathematical Sciences, 1 (2003), pp. 471–500.
- [24] M. L. MINION, *A hybrid parareal spectral deferred corrections method*, Communications in Applied Mathematics and Computational Science, 5 (2010), pp. 265–301.
- [25] MICHAEL L MINION AND SARAH A WILLIAMS, *Parareal and spectral deferred corrections*, in AIP Conference Proceedings, vol. 1048, 2008, p. 388.
- [26] J. NIEVERGELT, *Parallel methods for integrating ordinary differential equations*, Commun. ACM, 7 (1964), pp. 731–733.
- [27] VICTOR PEREYRA, *On improving an approximate solution of a functional equation by deferred corrections*, Numerische Mathematik, 8 (1966), pp. 376–391.
- [28] D. RUPRECHT, R. SPECK, M. EMMETT, M. BOLTEN, AND R. KRAUSE, *Poster: Extreme-scale space-time parallelism*, in Proceedings of the 2013 Conference on High Performance Computing Networking, Storage and Analysis Companion, SC '13 Companion, 2013.
- [29] R. SPECK, D. RUPRECHT, M. EMMETT, M. BOLTEN, AND R. KRAUSE, *A space-time parallel solver for the three-dimensional heat equation*, in Parallel Computing: Accelerating Computational Science and Engineering (CSE), vol. 25 of Advances in Parallel Computing, IOS Press, 2014, pp. 263 – 272.
- [30] R. SPECK, D. RUPRECHT, M. EMMETT, M. MINION, M. BOLTEN, AND R. KRAUSE, *A multi-level spectral deferred correction method*, BIT Numerical Mathematics, (2014).
- [31] R. SPECK, D. RUPRECHT, R. KRAUSE, M. EMMETT, M. MINION, M. WINKEL, AND P. GIBBON, *A massively space-time parallel N-body solver*, in Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC '12, Los Alamitos, CA, USA, 2012, IEEE Computer Society Press, pp. 92:1–92:11.
- [32] R. SPECK, D. RUPRECHT, M. MINION, M. EMMETT, AND R. KRAUSE, *Inexact spectral deferred corrections*, in Domain Decomposition Methods in Science and Engineering XXII, Lecture Notes in Computational Science and Engineering, 2014. Accepted.
- [33] H. J. STETTER, *Economical global error estimation*, in Stiff Differential Systems, R. A. Willoughby, ed., 1974, pp. 245–258.
- [34] S. VANDEWALLE AND G. HORTON, *Fourier mode analysis of the multigrid waveform relaxation and time-parallel multigrid methods*, Computing, 54 (1995), pp. 317–330.
- [35] MARTIN WEISER, *Faster SDC convergence on non-equidistant grids with DIRK sweeps*. ZIB Report 13–30, 2013.
- [36] P. E. ZADUNAISKY, *A method for the estimation of errors propagated in the numerical solution of a system of ordinary differential equations*, in Proceedings from Symposium no. 25 held in Thessaloniki, August 17–22, 1964, Georgios Ioannou Kontopoulos, ed., vol. 25 of The Theory of Orbits in the Solar System and in Stellar Systems, Academic Press, London, 1966, p. 281.