

MPI–OpenMP algorithms for the parallel space–time solution of Time Dependent PDEs

Ronald D. Haynes¹ and Benjamin W. Ong²

1 Introduction

Modern high performance computers offer hundreds of thousands of processors that can be leveraged, in parallel, to compute numerical solutions to time dependent partial differential equations (PDEs). For grid-based solutions to these PDEs, domain decomposition (DD) is often employed to add spatial parallelism [19].

Parallelism in the time variable is more difficult to exploit due to the inherent causality. Recently, researchers have explored this issue as a means to improve the scalability of existing parallel spatial solvers applied to time dependent problems. There are several general approaches to combine temporal parallelism with spatial parallelism. Waveform relaxation [15] is an example of a “parallel across the problem” method. The “parallel across the time domain” approaches include the parareal method [11, 17, 16]. The parareal method decomposes a time domain into smaller temporal subdomains and alternates between applying a coarse (relatively fast) sequential solver to compute an approximate (not very accurate) solution, and applying a fine (expensive) solver on each temporal subdomain in parallel. Alternatively, one can consider “parallel across the step” methods. Examples of such approaches include the computation of intermediate Runge–Kutta stage values in parallel [18], and Revisionist Integral Deferred Correction (RIDC) methods, which are the family of parallel time integrators considered in this paper. Parallel across the step methods allow for “small scale” parallelism in time. Specifically, we will show that if a DD implementation scales to N_x processors, a RIDC-DD parallelism will scale to $N_t \times N_x$ processors, where $N_t < 12$ in practice. This contrasts with parallel across the time domain approaches, which can potentially utilize $N_t \gg 12$.

This paper discusses the implementation details and profiling results of the parallel space–time RIDC-DD algorithm described in [5]. Two hybrid OpenMP – MPI frameworks are discussed: (i) a more traditional fork-join approach of combining threads before doing MPI communications, and (ii) a threaded MPI communications framework. The latter framework is highly desirable because existing (spatially parallel) legacy software can be easily integrated with the parallel time integrator. Numerical experiments measure the communication overhead of both frameworks, and demonstrate that the fork-join approach scales well in space and time. Our results indicate that one should strongly consider temporal parallelization for the solution of time dependent PDEs.

¹ Memorial University of Newfoundland, St. John’s, Newfoundland, Canada e-mail: rhaynes@mun.ca .² Michigan State University, Institute for Cyber-Enabled Research, East Lansing, MI, USA e-mail: ongbw@msu.edu

2 Review

This paper is interested in parallel space-time solutions to the linear heat equation. We describe the application of our method to the linear heat equation in one spatial dimension $x \in [0, 1]$ and $t \in [0, T]$,

$$u_t = u_{xx}, u(t, 0) = g_0(t), u(t, 1) = g_1(t), u(0, x) = u_0(x). \quad (1)$$

The actual numerical results in §4 are presented for the 2D heat equation.

2.1 RIDC

RIDC methods [6, 7] are a family of parallel time integrators that can be broadly classified as predictor corrector algorithms [10, 2]. The basic idea is to simultaneously compute solutions to the PDE of interest and associated error PDEs using a low-order time integrator. We first review the derivation of the error equation.

Suppose $v(t, x)$ is an approximate solution to (1), and $u(t, x)$ is the (unknown) exact solution. The error in the approximate solution is $e(t, x) = u(t, x) - v(t, x)$. We define the residual as $\varepsilon(t, x) = v_t(t, x) - v_{xx}(t, x)$. Then the time derivative of the error satisfies $e_t = u_t - v_t = u_{xx} - (v_{xx} + \varepsilon)$. The integral form of the error equation,

$$\left[e + \int_0^t \varepsilon(\tau, x) d\tau \right]_t = (v + e)_{xx} - v_{xx}, \quad (2)$$

can then be solved for $e(t, x)$ using the initial condition $e(0, x) = 0$. The correction $e(t, x)$ is combined with the approximate solution $v(t, x)$ to form an improved solution. This improved solution can be fed back in to the error equation (2) and the process repeated until a sufficiently accurate solution is obtained. It has been shown that each application of the error equation improves the order of the overall method, provided the integral is approximated with sufficient accuracy using quadrature [8].

We introduce some notation to identify the sequence of corrected approximations. Denote $v^{[p]}(t, x)$ as the approximate solution which has error $e^{[p]}(t, x)$, which is obtained by solving

$$\left[e^{[p]} + \int_0^t \varepsilon^{[p]}(\tau, x) d\tau \right]_t = (v^{[p]} + e^{[p]})_{xx} - v_{xx}^{[p]}, \quad (3)$$

where $v^{[0]}(t, x)$ denotes the initial approximate solution obtained by solving the physical PDE (1) using a low-order integrator. In general, the error from the p th correction equation is used to construct the $(p + 1)$ st approximation, $v^{[p+1]}(t, x) = v^{[p]}(t, x) + e^{[p]}(t, x)$. Hence, equation (3) can be expressed as

$$\left[v^{[p+1]} - \int_0^t v_{xx}^{[p]}(\tau, x) d\tau \right]_t = v_{xx}^{[p+1]} - v_{xx}^{[p]}. \quad (4)$$

We compute a low-order prediction, $v^{[0],n+1}$, for the solution of (1) at time t_{n+1} using a first-order backward Euler discretization (in time):

$$v^{[0],n+1} - \Delta t v_{xx}^{[0],n+1} = v^{[0],n}, \quad v^{[0],n+1}(a) = g_0(t^{n+1}), \quad v^{[0],n+1}(b) = g_1(t^{n+1}), \quad (5)$$

with $v^{[0],0}(x) = u_0(x)$. With some algebra, a first-order backward Euler discretization of equation (4) gives the update, $v^{[p+1],n+1}$, as

$$v^{[p+1],n+1} - \Delta t v_{xx}^{[p+1],n+1} = v^{[p+1],n} - \Delta t v_{xx}^{[p],n+1} + \int_{t^n}^{t^{n+1}} v_{xx}^{[p]}(\tau, x) d\tau, \quad (6)$$

with $v^{[p+1],n+1}(a) = g_0(t^{n+1})$ and $v^{[p+1],n+1}(b) = g_1(t^{n+1})$. The integral in equation (6) is approximated using a sufficiently high-order quadrature rule [8].

Parallelism in time is possible because the PDE of interest (1) and the error PDEs (4) can be solved simultaneously, after initial startup costs. The idea is to fill out the memory footprint, which is needed so that the integral in equation (6) can be approximated by high-order quadrature, before marching solutions to (5) and (6) in a pipe–line fashion. See Figure 1 for a graphical example, and [6] for more details.




Fig. 1 The black dots represent the memory footprint that must be stored before the white dots can be computed in a pipe. In this figure, $v^{[0],n+2}(x)$, $v^{[1],n+1}(x)$, $v^{[2],n}(x)$ and $v^{[3],n-1}(x)$ are computed simultaneously.

2.2 RIDC–DD

The RIDC–DD algorithm solves the predictor (5) and corrections (6) using DD algorithms in space. The key observation is that (5) and (6) are **both** elliptic PDEs of the form $(1 - \Delta t \partial_{xx})z = f(x)$. The function $f(x)$ is known from the solution at the previous time step and previously computed lower-order approximations. DD algorithms for solving elliptic PDEs are well known [3, 4]. The general idea is to replace the PDE by a coupled system of PDEs over some partitioning of the spatial domain using overlapping or non–overlapping subdomains. The coupling is provided by necessary transmission conditions at the subdomain boundaries. These transmission conditions are chosen to ensure the DD algorithm converges and to optimize the con-

vergence rate. In [5], as a proof of principle, (5-6) are solved using a classical parallel Schwarz algorithm, with overlapping subdomains and Dirichlet transmission conditions. Optimized RIDC-DD variants are possible using an optimized Schwarz DD method [13, 12, 9], to solve (5-6). The solution from the previous time step can be used as initial subdomain solutions at the interfaces. We will use RIDC p -DD to refer to a p th-order solution obtained using $p - 1$ RIDC corrections in time and DD in space.

3 Implementation Details

We view the parallel time integrator reviewed in §2.1 as a simple yet powerful tool to add further scalability to a legacy MPI or modern MPI-CUDA code, while improving the accuracy of numerical solution. The RIDC integrators benefit from access to shared memory because solving the correction PDE (6) requires both the solution from the previous time step and previously computed lower-order subdomain solution. Consequently, we propose two MPI-OpenMP hybrid implementations which map well to multi-core, multi-node compute resources. In the upcoming MPI 3.0 standard [1], shared memory access within the MPI library will provide alternative implementations.

Implementation #1: The RIDC-DD algorithm can be implemented using a traditional fork join approach, as illustrated in Program 1. After boundary information is exchanged, each MPI task spawns OpenMP threads to perform the linear solve. The threads are merged back together before MPI communication is used to check for convergence. The drawback to this fork-join implementation, is that the parallel space-time algorithm becomes tightly integrated, making it difficult to leverage an existing spatially parallel DD implementation.

```

1. MPI Initialization
2. ...
3.   for each time step
4.     for each Schwarz iteration
5.       MPI Comm (exchange boundary info)
6.       OMP Parallel for each prediction/correction
7.         linear solve
8.       end parallel
9.       MPI Comm (check for convergence)
10.    end
11.  end
12. ...
13. MPI Finalize

```

Program 1: RIDC-DD implementation using a fork-join approach. The time parallelism occurs *within* each Schwarz iteration, requiring a tight integration with an existing spatially parallel DD implementation.

Implementation #2: To leverage an existing spatially parallel DD implementation, a non-traditional hybrid approach must be considered. By changing the order of the loops, the Schwarz iterations for the prediction and the correction loops can be evaluated independently of each other. This is realized by spawning individual OpenMP threads to solve the prediction and correction loops on each sub-domain; the Schwarz iterations for the prediction/correction step run independently of each other until convergence. This implementation (Program 2) has several consequences: (i) a thread safe version of MPI supporting `MPI_THREAD_MULTIPLE` is required. (ii) In addition, we required a thread-safe, thread-independent version of `MPI_BARRIER`, `MPI_BROADCAST` and `MPI_GATHER`. To achieve this, we wrote our own wrapper library using the thread safe `MPI_SEND`, `MPI_RECV` and `MPI_SENDRECV` provided by (i).

```

1. MPI Initialization
2. ...
3.   for each time step
4.     OMP Parallel for each prediction/correction level
5.       for each Schwarz iteration
6.         MPI Comm (exchange boundary info)
7.         linear solve
8.         MPI Comm (check for convergence)
9.       end
10.    end parallel
11.  end
12. ...
13. MPI Finalize

```

Program 2: RIDC-DD implementation using a non-traditional hybrid approach. Notice that lines 5-9 are the Schwarz iterations that one would find in an existing spatially parallel DD implementation. Hence, provided the DD implementation is thread-safe, one could wrap the time parallelization around an existing parallel DD implementation.

4 Numerical Experiments

We show first that RIDC-DD methods converge with the designed orders in space and time. Then, we profile communication costs using TAU [14]. Finally, we show strong scaling studies for the RIDC-DD algorithm. We compute solutions to the heat equation in \mathbb{R}^2 , where centered finite differences are used to approximate the second derivative operator. Errors are computed using the known analytic solution. The computations are performed at the High Performance Computing Center at Michigan State University, where nodes (consisting of two quad core Intel Westmere processors) are interconnected using infiniband and a high speed Lustre file system.

4.1 Convergence Studies and Profile Analysis

In Figure 2, the convergence plots show that our classical Schwarz RIDC-DD algorithm converges as expected in space and time. In general, one would balance the orders of the errors in space and time appropriately for efficiency. Here we pick RIDC4 since it mapped well to our available four core sockets and to demonstrate the scalability of our algorithm in time. We could, of course, use a fourth order method in space. The Schwarz iterations are iterated until a tolerance of 10^{-12} is reached for the predictors and correctors (which explains why the error in the fourth-order approximation levels out as the time step becomes small).




Fig. 2 (a) Classical Schwarz RIDC p -DD algorithms, $p = 1, 2, 3, 4$, converge to the reference solution with the designed orders of accuracy. Here Δx is fixed while Δt is varied. (b) Second-order convergence in space is demonstrated for the fourth-order RIDC-DD algorithm. Here, Δt is fixed while Δx is varied.

The communication costs for our two implementations of RIDC4-DD are profiled using TAU [14]. We see in Figure 3, communication costs are minimal for implementation #1, and scales nicely as the number of nodes is increased, but the communication cost is significant for implementation #2. In Figure 3(a,c), the domain is discretized into 180×180 grid nodes, which are split into a 3×3 configuration of subdomains. In Figure 3(b,d), the domain is discretized into 360×360 grid nodes, which are split into a 6×6 configuration of subdomains. This keeps the number of grid points per subdomain constant so that the computation time for the matrix factorization and linear solve are the same.

4.2 Characterizing Parallel Performance

Due to the better communication profile, we use framework #1 for our experiments. We fix $\Delta x = \frac{1}{180}$, $\Delta y = \frac{1}{180}$, $\Delta t = \frac{1}{1000}$, and TOL= 10^{-12} (the Schwarz iteration tolerance). We consider three configurations of subdomains: 2×2 , 4×4 and 6×6 . For each configuration we illustrate the speedup and efficiency due to the time parallelism in Figure 4. We choose to fix the ratio between the overlap and subdomain

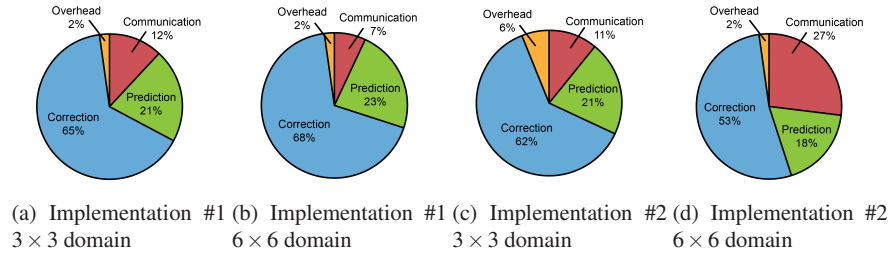


Fig. 3 Profile of the RIDC4-DD algorithm using both implementations. Overhead and communication costs are reasonable for implementation #1, but are high for implementation #2.

size to ensure the number of unknowns on each subdomain scales appropriately as the number of subdomains is increased.

In Figure 4 we show three curves corresponding to a 2×2 , 4×4 and a 6×6 configuration of subdomains. For each configuration we compute a fourth order solution in time using 1, 2 and 4 threads. The 6×6 configuration of subdomains with 4 threads uses a total of 144 cores. We plot the efficiency (with respect to the one thread run) as a function of the number of threads. Speedup is evident as temporal parallelization is improved, however, efficiency decreases as the number of subdomains increases.




Fig. 4 Scaling study (in time) for a RIDC4-DD algorithm.

5 Conclusions

This paper has presented the implementation details and first reported profiling results for a newly proposed space–time parallel algorithm for time dependent PDEs. The RIDC–DD method combines traditional domain decomposition in space with a new family of deferred correction methods designed to allow parallelism in time. Two possible implementations are described and profiled. The first, a traditional hybrid OpenMP–MPI implementation, requires potentially difficult modifications of an existing parallel spatial solver. Numerical experiments verify that the algorithm achieves its designed order of accuracy and scales well. The second strategy allows a relatively easy reuse of an existing parallel spatial solver by using OpenMP to spawn threads for the simultaneous prediction and correction steps. This non–traditional hybrid use of OpenMP and MPI currently requires writing of custom thread–safe and thread–independent MPI routines. Profile analysis shows that our non–traditional use of OpenMP–MPI suffers from higher communication costs than the standard use of OpenMP–MPI. An inspection of the prediction and correction

equations indicates that optimized Schwarz variants of the algorithm are possible and will enjoy nice load balancing. This work is ongoing.

Acknowledgements This work was supported by the Institute for Cyber-Enabled Research (iCER) at MSU, NSERC Discovery Grant 311796, and AFOSR Grant FA9550-12-1-0455.

References

1. Mpi 3.0 standardization effort. http://meetings.mpi-forum.org/MPI_3.0_main_page.php. Accessed 10/25/2012
2. Böhmer, K., Stetter, H.: Defect correction methods. Theory and applications. Computing Supplementum, 5 (1984)
3. Cai, X.C.: Additive Schwarz algorithms for parabolic convection-diffusion equations. Numer. Math. **60**(1), 41–61 (1991)
4. Cai, X.C.: Multiplicative Schwarz methods for parabolic problems. SIAM J. Sci. Comput. **15**(3), 587–603 (1994)
5. Christlieb, A., Haynes, R., Ong, B.: A parallel space-time algorithm. SIAM J. Sci. Comput. **34**(5), 233–248 (2012)
6. Christlieb, A., Macdonald, C., Ong, B.: Parallel high-order integrators. SIAM J. Sci. Comput. **32**(2), 818–835 (2010)
7. Christlieb, A., Ong, B.: Implicit parallel time integrators. J. Sci. Comput. **49**(2), 167–179 (2011)
8. Christlieb, A., Ong, B., Qiu, J.M.: Comments on high order integrators embedded within integral deferred correction methods. Comm. Appl. Math. Comput. Sci. **4**(1), 27–56 (2009)
9. Dubois, O., Gander, M., Loisel, S., St-Cyr, A., Szyld, D.: The optimized Schwarz method with a coarse grid correction. SIAM J. Sci. Comput. **34**(1), A421–A458 (2012)
10. Dutt, A., Greengard, L., Rokhlin, V.: Spectral deferred correction methods for ordinary differential equations. BIT **40**(2), 241–266 (2000)
11. Gander, M., Vandewalle, S.: On the superlinear and linear convergence of the parareal algorithm. Lecture Notes in Computational Science and Engineering **55**, 291 (2007)
12. Gander, M.J.: Optimized Schwarz methods. SIAM J. Numer. Anal. **44**(2), 699–731 (2006)
13. Gander, M.J., Halpern, L.: Optimized Schwarz waveform relaxation methods for advection reaction diffusion problems. SIAM J. Numer. Anal. **45**(2), 666–697 (2007)
14. Koehler, S., Curreri, J., George, A.: Performance analysis challenges and framework for high-performance reconfigurable computing. Parallel Computing **34**(4), 217–230 (2008)
15. Lelarasme, E., Ruehli, A.E., Sangiovanni-Vincentelli, A.L.: The waveform relaxation method for time-domain analysis of large scale integrated circuits. IEEE Trans. on CAD of IC and Syst. **1**, 131–145 (1982)
16. Lions, J., Maday, Y., Turinici, G.: A “parareal” in time discretization of PDEs. Comptes Rendus de l’Academie des Sciences Series I Mathematics **332**(7), 661–668 (2001)
17. Minion, M., Williams, S.: Parareal and spectral deferred corrections. In: NUMERICAL ANALYSIS AND APPLIED MATHEMATICS: International Conference on Numerical Analysis and Applied Mathematics 2008. AIP Conference Proceedings, vol. 1048, pp. 388–391 (2008)
18. Nievergelt, J.: Parallel methods for integrating ordinary differential equations. Communications of the ACM **7**(12), 731–733 (1964)
19. Toselli, A., Widlund, O.: Domain decomposition methods—algorithms and theory, *Springer Series in Computational Mathematics*, vol. 34. Springer-Verlag, Berlin (2005)