



UNIVERSITY OF LEEDS

This is a repository copy of *Shared Memory Pipelined Parareal*.

White Rose Research Online URL for this paper:

<http://eprints.whiterose.ac.uk/115836/>

Version: Accepted Version

Proceedings Paper:

Ruprecht, D orcid.org/0000-0003-1904-2473 (2017) Shared Memory Pipelined Parareal. In: Lecture Notes in Computer Science. Euro-Par 2017: 23rd International European Conference on Parallel and Distributed Computing, 28 Aug - 01 Sep 2017, Santiago de Compostela, Spain. Springer Verlag , pp. 669-681. ISBN 978-3-319-64202-4

https://doi.org/10.1007/978-3-319-64203-1_48

© 2017, Springer International Publishing AG. This is an author produced version of a paper published in Lecture Notes in Computer Science. The final publication is available at Springer via https://doi.org/10.1007/978-3-319-64203-1_48. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Shared memory pipelined Parareal

Daniel Ruprecht^{1,2}

¹ School of Mechanical Engineering, LS2 9JT, Leeds, United Kingdom
d.ruprecht@leeds.ac.uk

² Institute of Computational Science, Università della Svizzera italiana, Via
Giuseppe Buffi 13, CH-6900 Lugano, Switzerland

Abstract. For the parallel-in-time integration method Parareal, pipelining can be used to hide some of the cost of the serial correction step and improve its efficiency. The paper introduces an OpenMP implementation of pipelined Parareal and compares it to a standard MPI-based variant. Both versions yield almost identical runtimes, but, depending on the compiler, the OpenMP variant consumes about 7% less energy and has a significantly smaller memory footprint. However, its higher implementation complexity might make it difficult to use in legacy codes and in combination with spatial parallelisation.

Keywords: Parareal, parallel-in-time integration, pipelining, OpenMP

1 Introduction

Computational science faces a variety of challenges stemming from the massive increase in parallelism in state-of-the-art high-performance computing systems. One important requirement is the development of new and inherently parallel numerical methods. Parallel-in-time integration has been identified as a promising direction of research for the parallelisation of the solution of initial value problems [5].

Probably the most widely studied and used parallel-in-time method is Parareal [11], but see M. Gander’s overview for a discussion of a variety of other methods [6]. Parareal iterates between an expensive fine integrator run in parallel and a cheap coarse method which runs in serial and propagates corrections forward in time. While the unavoidable serial part limits parallel efficiency according to Amdahl’s law, some of its cost can be hidden by using a so-called *pipelined* implementation [2,12]. Pipelining reduces the effective cost of the serial correction step in Parareal and therefore improves speedup. Even more optimisation is possible by using an event-based approach [4], but this requires a suitable execution framework that is not available on all machines.

Pipelining happens automatically when implementing Parareal in MPI but it is not straightforward in OpenMP and so far no shared memory version of Parareal with pipelining has been described. Previous studies almost exclusively used MPI to implement Parareal and the very few using OpenMP considered

only the non-pipelined version [9]. However, using shared memory can have advantages, since it avoids e.g. the need to allocate buffers for message passing. The disadvantage is that naturally OpenMP is limited to a shared memory unit. Since convergence of Parareal tends to deteriorate if too many parallel time slices are computed [7] and given the trend to “fat” compute nodes with large numbers of cores shared memory Parareal might nevertheless be an attractive choice. It could be useful, e.g., for simulations of power grids and other applications where comparatively small systems of differential(-algebraic) equations have to be solved faster than real-time [10]. This paper introduces an OpenMP-based version of pipelined Parareal and compares it to a standard MPI-based implementation.

2 The Parareal parallel-in-time integration method

The starting point for Parareal is an initial value problem

$$\dot{\mathbf{q}} = \mathbf{f}(\mathbf{q}(t), t), \quad \mathbf{q}(0) = \mathbf{q}_0, \quad t \in [0, T], \quad (1)$$

which, in the numerical example below, arises from the spatial discretisation of a PDE (“method-of-lines”) with $\mathbf{q} \in \mathbb{R}^{N_{\text{dof}}}$ being a vector containing all degrees-of-freedom. Let $\mathcal{F}_{\delta t}$ denote a numerical procedure for the approximate solution of (1), for example a Runge-Kutta method. Denote further by $\mathbf{q} = \mathcal{F}_{\delta t}(\tilde{\mathbf{q}}, t_2, t_1)$ the result of approximately integrating (1) forward in time from some starting value $\tilde{\mathbf{q}}$ at a time t_1 to a time $t_2 > t_1$ using $\mathcal{F}_{\delta t}$.

To parallelise the numerical integration of (1), decompose $[0, T]$ into time-slices $[t_p, t_{p+1}]$, $p = 0, \dots, P - 1$ where P is the number of cores, equal to the number of processes (MPI) or threads (OpenMP). For simplicity, assume here that all time slices have the same length and that the whole interval $[0, T]$ is covered with a single set of time slices so that no restarting or “sliding window” is required [15].

Parareal needs a second time integrator denoted $\mathcal{G}_{\Delta t}$, which has to be cheap to compute but can be much less accurate (commonly referred to as the “coarse propagator”). It begins with a prediction step, computing rough guesses of the starting value \mathbf{q}_p^0 for each time slice by running the coarse method once. Here, subscript p indicates an approximation of the solution at time t_p . It then computes the iteration

$$\mathbf{q}_{p+1}^k = \mathcal{G}_{\Delta t}(\mathbf{q}_p^k, t_{p+1}, t_p) + \mathcal{F}_{\delta t}(\mathbf{q}_p^{k-1}, t_{p+1}, t_p) - \mathcal{G}_{\Delta t}(\mathbf{q}_p^{k-1}, t_{p+1}, t_p) \quad (2)$$

concurrently on each time slice for $p = 0, \dots, P - 1$, $k = 1, \dots, K$. Because the computationally expensive evaluation of the fine propagator can be parallelised across time slices, iteration (2) can run in less wall clock time than running $\mathcal{F}_{\delta t}$ serially – provided the coarse method is cheap enough and the number of required iterations K is small.

The expected performance of Parareal can be described by a simple theoretical model [12]. Denote by c_c the cost of integrating over one time slice using

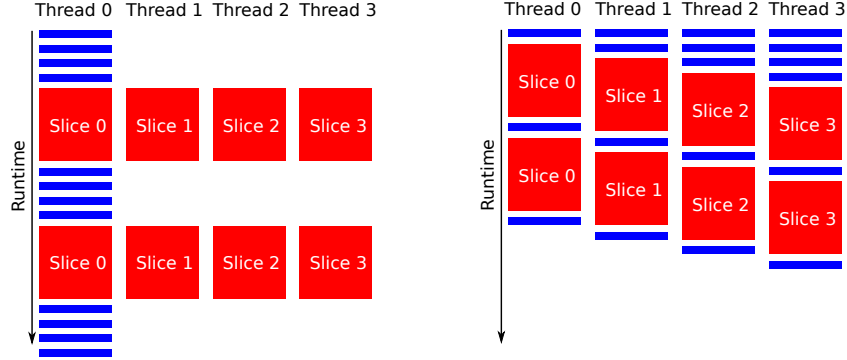


Fig. 1: Execution diagram for Parareal without (left) and with (right) pipelining. Red blocks correspond to c_f , the time needed to run the fine integrator $\mathcal{F}_{\delta t}$ over one time slice $[t_p, t_{p+1}]$ while blue blocks correspond to c_c , the time required by the coarse method $\mathcal{G}_{\Delta t}$. Pipelining allows to hide some of the cost of the coarse method. While it comes naturally when using MPI to implement Parareal (there, **Thread** would refer to a **Process**), using simple loop-based parallelism with OpenMP results in the non-pipelined version shown on the left.

$\mathcal{G}_{\Delta t}$ and by c_f the cost when using $\mathcal{F}_{\delta t}$. Because all time slices are assumed to consist of the same number of steps and an explicit method is used here, it can be assumed that c_f and c_c are identical for all time slices. Neglecting overhead, speedup of Parareal using K iterations against running the fine method in serial is approximately

$$s_{\text{np}}(P) = \frac{Pc_f}{(1+K)Pc_c + Kc_f} = \frac{1}{(1+K)\frac{c_c}{c_f} + \frac{K}{P}}. \quad (3)$$

It is possible to hide some of the cost of the coarse propagator and the name *pipelining* has been coined for this approach [12]. Figure 1 sketches the execution diagrams of both a non-pipelined (left) and pipelined (right) implementation for four time slices. As can be seen, pipelining reduces the effective cost of the coarse correction step in each iteration from $P \times c_c$ to c_c – but note that the initial prediction step still has cost $P \times c_c$ as before. For pipelined Parareal, estimate (3) changes to

$$s_{\text{p}}(P) = \frac{Pc_f}{Pc_c + Kc_c + Kc_f} = \frac{1}{(1 + \frac{K}{P})\frac{c_c}{c_f} + \frac{K}{P}}. \quad (4)$$

Because $K/P \ll K$, the pipelined version allows for better speedup, that is $s_{\text{np}}(P) \leq s_{\text{p}}(P)$. However, because pipelining only hides cost from the coarse integrator, the effect is smaller when the coarse method is very cheap and $c_c/c_s \ll 1$. In that case, the term K/P dominates the estimate which is not affected by pipelining.

3 Pipelined Parareal in OpenMP

The implementation of Parareal with pipelining in OpenMP is sketched in Algorithm 1. A description of how Parareal is implemented using MPI is available in the literature [1] and is therefore not repeated here.

Threads are spawned by the `OMP PARALLEL` directive in line 1.2 and terminated by `OMP END PARALLEL` in line 1.37. Manual synchronisation is required so that P OpenMP locks is created using `OMP_INIT_LOCK` in line 1.4, one for each thread. During the fine integrator and update step, these locks are set and unset using `OMP_SET_LOCK` and `OMP_UNSET_LOCK` to protect buffers during writes and avoid race conditions.

The algorithm consists of the following parts:

- **Prediction step:** lines 1.5–1.13. Each thread is computing its own coarse prediction of its starting value \mathbf{q}_p^0 in a parallelised loop. The coarse value $\mathcal{G}_{\Delta t}(\mathbf{q}_p^0, t_{p+1}, t_p)$ is also computed and stored for use in the first iteration. The later the time slice (indicated by a higher thread number p), the more steps the thread must compute and thus the larger its workload (cf. Figure 1). Therefore, at the end of the coarse prediction loop, the `NOWAIT` clause is required to avoid implicit synchronisation and enable pipelining.
- **Parareal iteration:** lines 1.14–1.36. Here, both the fine integrator and update step are performed inside a single loop over all time slices, parallelised by `OMP DO` directives. Because parts of the loop (the update step) have to be executed in serialised order, the `ORDERED` directive has to be used in line 1.15. Again, to avoid implicit synchronisation at the end of the loop, the `NOWAIT` clause is required in line 1.35.
 - **Fine integrator:** lines 1.17–1.20. Before the fine integrator is executed, an `OMP_LOCK` is set to indicate that the thread will start writing into buffer $q(p)$. Because thread $p-1$ accesses this buffer in its update step, locks are necessary to prevent race conditions and incorrect solutions. After the lock is set, the thread proceeds with the computation of $\mathcal{F}_{\delta t}(\mathbf{q}_p^k, t_{p+1}, t_p)$ and computation of the difference δq between coarse and fine value. Then, since $q(p)$ is now up to date and δq ready, the lock can be released.
 - **Update step:** lines 1.21–1.33. The update step has to be performed in proper order, from first to last time slice. Therefore, it is enclosed in `ORDERED` directives, indicating that this part of the loop is to be executed in serial order. Then, as in the two other versions, the update step is initialised with $\mathbf{q}_0^{k+1} = \mathbf{q}_0$. For every time slice, the coarse value of the updated initial guess is computed and the update performed. The updated end value is written into buffer $q(p+1)$ to serve as the new starting value for the following time slice. However, to prevent thread p from writing into $q(p+1)$ while thread $p+1$ is still running the fine integrator, thread p sets `OMP_LOCK` number $p+1$ while performing the update.

Algorithm 1: Parareal with pipelining using OpenMP

input: Initial value q_0 ; number of iterations K

```
1.1  $q(0) \leftarrow q_0$ 
1.2 OMP PARALLEL
1.3  $P = \text{OMP\_GET\_MAX\_THREADS}()$ 
1.4 OMP\_INIT\_LOCK( $P$ )
1.5 OMP DO
1.6 for  $p = 0, P - 1$  do
1.7    $q(p) \leftarrow q(0)$ 
1.8   if Thread not first then
1.9      $q(p) \leftarrow \mathcal{G}(q(p), t_p, 0)$ 
1.10  end
1.11   $g_c(p) \leftarrow \mathcal{G}(q(p), t_{p+1}, t_p)$ 
1.12 end
1.13 OMP END DO NOWAIT
1.14 for  $k = 1, K$  do
1.15   OMP DO ORDERED
1.16   for  $p = 0, P - 1$  do
1.17     OMP\_SET\_LOCK( $p$ )
1.18      $q(p) \leftarrow \mathcal{F}_{\delta t}(q(p), t_{p+1}, t_p)$ 
1.19      $\delta q(p) \leftarrow q(p) - q_c(p)$ 
1.20     OMP\_UNSET\_LOCK( $p$ )
1.21     OMP ORDERED
1.22     if Thread is first then
1.23       OMP\_SET\_LOCK( $0$ )
1.24        $q(0) \leftarrow q_0$ 
1.25       OMP\_UNSET\_LOCK( $0$ )
1.26     end
1.27      $q_c(p) \leftarrow \mathcal{G}_{\Delta t}(q(p), t_{p+1}, t_p)$ 
1.28     if Thread not last then
1.29       OMP\_SET\_LOCK( $p+1$ )
1.30        $q(p+1) \leftarrow q_c(p) + \delta q(p)$ 
1.31       OMP\_UNSET\_LOCK( $p+1$ )
1.32     end
1.33   end
1.34   OMP END ORDERED
1.35 end
1.36 end
1.37 OMP END PARALLEL
```

4 Numerical results

This section compares the pipelined OpenMP implementation to a straightforward MPI variant with respect to runtime, memory footprint and energy consumption. The code used here for benchmarking is written in Fortran 90 and available for download [14]. It is special-purpose and solves the nonlinear 3D

Burgers' equation

$$u_t + u \cdot \nabla u = \nu \Delta u \tag{5}$$

on $[0, 1]^3 \subset \mathbb{R}^3$ with periodic boundary conditions using finite differences. Both implementations of Parareal use the same modules to provide the coarse and fine integrator and spatial discretisation. Tests guarantee that the two implementations of Parareal produce results that are identical up to a tolerance of $\varepsilon = 10^{-14}$ and thus essentially to round-off error. To detect possible race conditions, the comparison is run multiple times. Up to 100 instances of the test were passed on both used architectures. Furthermore, both implementations of Parareal use three auxiliary buffers per time-slice: q to store the fine value and for communication, δq to store the difference $\mathcal{F}_{\delta t}(q) - \mathcal{G}_{\Delta t}(q)$ needed in the correction step and q_c to store the coarse value from the previous iteration.

A strong stability preserving Runge-Kutta method (RK3-SSP) [16] with a fifth order WENO finite difference discretisation [16] for the advection term and a fourth order centred difference for the diffusion term is used for $\mathcal{F}_{\delta t}$. For $\mathcal{G}_{\Delta t}$, a first order forward Euler with a simple first order upwind stencil for the advection term and a second order centred stencil for the diffusive term is used. Being a simplified version of the Navier-Stokes equations, (5) is a popular benchmark and finite difference stencils are a widely used motif in computational science, so that the results can be expected to hold for more general scenarios, at least qualitatively.

Parameters for the simulation are a viscosity parameter of $\nu = 0.02$ and a spatial discretisation on both levels with $N_x = N_y = N_z = 40$ grid points in every direction. The simulation is run until $T = 1.0$ with a coarse time step of $\Delta t = 1/192$ and a fine step of $\delta t = 1/240$. Because of the high computational cost of the WENO-5 method in comparison to a cheap first order upwind scheme and the fact that RK3SSP needs three evaluations of the right hand side per step while the Euler method needs only one, the coarse propagator is about a factor of forty faster, despite the fact that the coarse step is only a factor of 1.25 larger than the fine.

To fix the number of iterations to a meaningful value which guarantees comparable accuracy from Parareal and serial fine integrator, we estimate the discretisation error of $\mathcal{F}_{\delta t}$ by comparing against a reference solution with time step $\delta t/10$. This gives estimates for the fine relative error at $T = 1$ of about $e_{\text{fine}} \approx 5.9 \times 10^{-5}$ and for the coarse error of about $e_{\text{coarse}} \approx 7.3 \times 10^{-2}$. For $P = 24$ time slices, after three iterations, the defect between Parareal and the fine solution is approximately 1.4×10^{-4} , after four iterations 1.5×10^{-5} . We therefore fix the number of iterations to $K = 4$ so that for all values of P Parareal produces a solution with the same accuracy as the fine integrator.

Benchmarks are run on two systems. The first is a commodity work station with an 8-core Intel Xeon(R) E5-1660 CPU and 32 GigaByte of memory running CentOS Linux 7. Flags `-O3` and `-fopenmp` were used when compiling the code for maximum optimisation and to enable OpenMP. As the code is stand alone no external libraries have to be linked. The used MPI implementation is mpich-3.0.4, compiled with GCC-4.8.5.

The second system is one node of PIZ DORA at the Swiss National Supercomputing Centre CSCS.³ DORA is a Cray XC40 with a total of 1,256 compute nodes. Each node contains two 12-core Intel Broadwell CPUs and has 64 or 128 GigaByte of RAM and nodes are connected through a Cray Aries interconnect, using a dragonfly network topology. Two compilers are tested, the GCC-4.9.2 and the Cray Fortran compiler version 8.3.12. Both use the MPICH MPI library version 7.2.2. Compiler flags `-O3` and `-fopenmp` (GCC) or `omp` (Cray compiler) are used for compilation. Performance data for each completed job is generated using the Cray *Resource Utilisation Reporting* tool RUR [3]. RUR collects compute node statistics for each job and provides data on user and system time, maximum memory used, amount of I/O operations, consumed energy and other metrics. However, it only collects data for a full node and not for individual CPUs or cores.

4.1 Wall clock time and speedup

At first, runtime and speedup compared to the serial execution of the fine integrator are assessed. On both the Linux work station and DORA, five runs are performed for each variant of Parareal and each value of P and the average runtime is reported. Measured runtimes are quite stable across different runs: the largest relative standard deviation of all performed five-run ensembles is smaller than 0.05 on Linux and smaller than 0.01 on DORA. Therefore, plots show only the average values without error bars, because those are hardly recognisable and clutter the figure.

Figure 2 shows runtimes in seconds depending on the number of cores on DORA (left) and Linux (right). For DORA, only results from the Cray compiler are shown, which tends to generate slightly faster code than the GCC. The runtime of the serial fine integrator is indicated by a horizontal black line. The CPU in the Linux system has a slightly higher clock speed so that runtimes are a bit faster than on DORA. For $P = 8$ cores, for example, OpenMP-Parareal runs in slightly less than 5 s there while taking about 5.7 s on DORA. Differences between the OpenMP and MPI version are small on both systems, but for $P = 8$ OpenMP-Parareal is marginally faster than the MPI version on the work station.

In addition, Figure 3 shows the speedup relative to the fine integrator run serially. The black line indicates projected speedup according to (4). Both versions fall short of the theoretically possible speedup. Because of overheads, running P instances of $\mathcal{F}_{\delta t}$ on P cores does take longer than running a single instance. However, differences between MPI and OpenMP are small. As far as runtimes and speedup are concerned, there is no indication that using the more complex OpenMP version provides significant benefits.

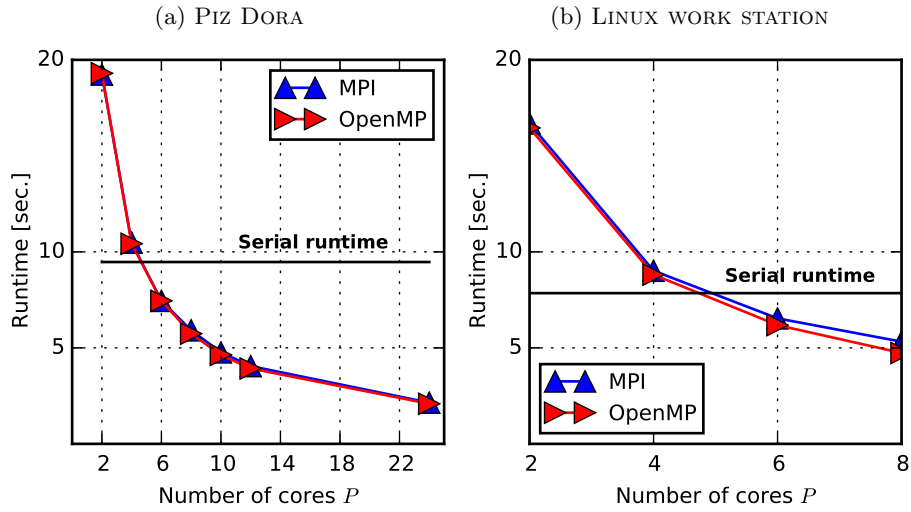


Fig. 2: Five run averages of runtime with relative standard deviation below 0.01 on DORA (2a) and 0.05 on a Linux work station (2b).

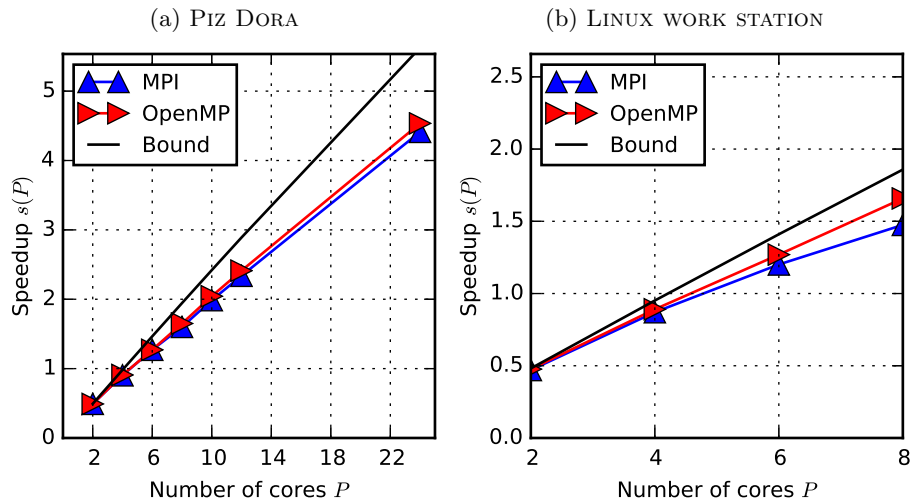


Fig. 3: Speedup computed from average runtimes shown in Figure 2 on DORA (3a) and a Linux work station (3b).

4.2 Memory footprint

The memory footprint of the code is measured only on DORA where RUR is available. In contrast to runtime and energy, the memory footprint, as expected, does not vary between runs. Therefore Figure 4 shows a visualisation of the data

³ <http://www.cscs.ch/>




Fig. 4: Maximum memory allocated in MegaByte for GCC (4a) and Cray compiler (4b) for the three different versions of Parareal depending on the number of used cores P . The black line indicates expected memory consumption computed as number of cores time memory footprint of serial fine integrator.

from a single run with no averaging. The bars indicate the maximum required memory in MegaByte (MB) while the black line indicates the expected memory consumption using P cores computed as

$$m(P) = P \times m_{\text{serial}} \quad (6)$$

where m_{serial} is the value measured for a reference run of the fine integrator. Because copies of the solution have to be stored for every time slice, the total memory required for Parareal can be expected to increase linearly with the number of cores in time. Note, however, that memory required *per core* stays constant if it follows (6).

For the OpenMP variant compiled with GCC, the memory footprint shown in Figure 4a exactly matches the expected values. The Cray compiler, shown in Figure 4b, leads to a smaller than expected memory footprint, but memory requirements still increase linearly with the number of time slices. For both compilers, the MPI version causes a noticeable overhead in terms of memory footprint, most likely because of internal allocation of additional buffers for sending and receiving [13]. For both OpenMP and MPI, the total memory footprint is *larger* for the Cray than for the GCC compiler, but the effect is much more pronounced for the MPI implementation (329 MB versus 261 MB) than for the OpenMP variant (200 MB versus 193 MB).

It is important to note that both implementations allocate three auxiliary buffers per core. The overhead in terms of memory in MPI does thus not simply

stem from allocating an additional buffer for communication but comes from within the MPI library. The OpenMP implementation avoids this overhead.

4.3 Energy-to-solution

The RUR tool reports the energy-to-solution for every completed job. Because RUR can only measure energy usage for a full node, results are reported here for runs using the full number of $P = 24$ cores available on a DORA node. In contrast to runtimes and memory footprint, energy measurements show significant variations between runs due to random fluctuations. Thus, the presented values are averages over ensembles of 50 runs for each version of Parareal. This number of runs has been sufficient to reduce the relative standard deviation to below 0.09 in both configurations and therefore gives a robust indication of actual energy requirements. Table 1 shows the results including 95% confidence intervals, assuming energy-to-solution is normally distributed. For comparison, energy-to-solution is also shown for a simple non-pipelined OpenMP implementation where only the loop around $\mathcal{F}_{\delta t}$ is parallelised using `OMP PARALLEL DO`.

GCC compiler			
	\emptyset energy (Joule)	confidence (Joule)	\emptyset runtime (sec.)
MPI	844.04	± 15.89	2.455
OpenMP	801.14	± 13.18	2.930
OpenMP(pipelined)	783.72	± 11.53	2.448
Cray compiler			
	\emptyset energy (Joule)	confidence (Joule)	\emptyset runtime (sec.)
MPI	784.24	± 11.93	2.146
OpenMP	833.12	± 20.99	2.400
OpenMP(pipelined)	784.72	± 12.18	2.088

Table 1: Average energy-to-solution for the three different variants run on $P = 24$ cores.

When using the GCC compiler, MPI-Parareal consumes more energy than both the pipelined and non-pipelined OpenMP versions. The averages are well outside the confidence interval for the MPI version, so this is very unlikely just a chance result. Moreover, because runtimes are almost identical for MPI and pipelined OpenMP, the differences in energy-to-solution cannot solely be attributed to differences in time-to-solution. This is supported by the fact that non-pipelined OpenMP, despite being significantly slower, still consumes less energy than the MPI variant. Tracking down the precise reason for the differences in energy-to-solution and power requirement will require detailed tracing of power uptake which is only possible on specially prepared machines [8] and thus left for future work.

For code generated with the Cray compiler, both OpenMP and MPI lead to almost identical energy requirements. Confidence intervals are 784.24 ± 11.93 J for

MPI and 784.24 ± 12.18 J for OpenMP. It seems likely that the compiler optimises the message passing to take advantage of the shared memory on the single node. Supposedly, the MPI version handles communication in a way that is similar to what is explicitly coded in the OpenMP version. However, as shown in Subsection 4.2, this automatic optimisation comes at the expense of a significantly larger memory footprint.

Note that the energy consumption of Parareal has previously been studied [1]. By comparing against a simple theoretical model, it has been shown that the energy overhead of Parareal (defined as energy-to-solution of Parareal divided by energy-to-solution of the fine serial integrator), is mostly due to Parareal's intrinsic suboptimal parallel efficiency. While improving parallel efficiency of parallel-in-time integration clearly remains the main avenue for improving energy efficiency, the results here suggest that in some cases a shared memory approach can provide non-trivial additional savings.

5 Summary

The paper introduces and analyses an OpenMP implementation of the parallel-in-time method Parareal with pipelining. Pipelining allows to hide some of the cost of the serial coarse correction step in Parareal and is important to optimise its efficiency (even though it cannot relax the inherent limit on parallel efficiency given by the inverse of the number of required iterations). Pipelining comes naturally in a distributed memory MPI implementation, but is not straightforward when using OpenMP. The new OpenMP implementation is compared to the standard MPI variant in terms of runtime, memory footprint and energy consumption for both a Cray compiler and the GCC. Both versions produce essentially identical runtimes. For both compilers, using OpenMP leads to reductions in memory footprint, but the effect is more pronounced for the Cray compiler. In terms of energy-to-solution, the results strongly depend on the compiler: while for GCC the OpenMP version is more energy efficient than the MPI version, there is no difference for the Cray compiler.

The results show that contemplating a shared memory strategy to implement “parallel-across-the-steps” methods like Parareal can be worthwhile. Even though it is more complicated, it can reduce memory requirements. A potential caveat is whether the benefits carry over to the full space-time parallel case, where a parallel-in-time method is combined with spatial decomposition. For Parareal without pipelining the potential of such a hybrid space-time parallel approach has been illustrated [9] but whether this applies to the pipelined version introduced here remains to be seen.

References

1. Arteaga, A., Ruprecht, D., Krause, R.: A stencil-based implementation of Parareal in the C++ domain specific embedded language STELLA. *Applied Mathematics and Computation* 267, 727–741 (2015), <http://dx.doi.org/10.1016/j.amc.2014.12.055>

2. Aubanel, E.: Scheduling of Tasks in the Parareal Algorithm. *Parallel Computing* 37, 172–182 (2011), <http://dx.doi.org/10.1016/j.parco.2010.10.004>
3. Barry, A.: Resource utilization reporting: Gathering and evaluating HPC system usage. In: CUG2013 Proceedings (2013), https://cug.org/proceedings/cug2013_proceedings/includes/files/pap103-file2.pdf
4. Berry, L.A., Elwasif, W.R., Reynolds-Barredo, J.M., Samaddar, D., Snchez, R.S., Newman, D.E.: Event-based parareal: A data-flow based implementation of parareal. *Journal of Computational Physics* 231(17), 5945–5954 (2012), <http://dx.doi.org/10.1016/j.jcp.2012.05.016>
5. Dongarra, J., al.: Applied Mathematics Research for Exascale Computing. Tech. Rep. LLNL-TR-651000, Lawrence Livermore National Laboratory (2014), <http://science.energy.gov/%7E/media/ascr/pdf/research/am/docs/EMWGreport.pdf>
6. Gander, M.J.: 50 years of Time Parallel Time Integration. In: Multiple Shooting and Time Domain Decomposition. Springer (2015), http://dx.doi.org/10.1007/978-3-319-23321-5_3
7. Gander, M.J., Vandewalle, S.: Analysis of the Parareal Time-Parallel Time-Integration Method. *SIAM Journal on Scientific Computing* 29(2), 556–578 (2007), <http://dx.doi.org/10.1137/05064607X>
8. Isci, C., Martonosi, M.: Runtime power monitoring in high-end processors: Methodology and empirical data. In: Proceedings of the 36th Annual IEEE/ACM International Symposium on Microarchitecture. pp. 93–. MICRO 36, IEEE Computer Society, Washington, DC, USA (2003), <http://dl.acm.org/citation.cfm?id=956417.956567>
9. Krause, R., Ruprecht, D.: Hybrid Space-Time Parallel Solution of Burgers' Equation. In: Domain Decomposition Methods in Science and Engineering XXI. Lecture Notes in Computational Science and Engineering, vol. 98, pp. 647–655. Springer International Publishing (2014), http://dx.doi.org/10.1007/978-3-319-05789-7_62
10. Lecouvez, M., Falgout, R., Woodward, C., Top, P.: A parallel multigrid reduction in time method for power systems (2016), <http://www.osti.gov/scitech/biblio/1281664>
11. Lions, J.L., Maday, Y., Turinici, G.: A "parareal" in time discretization of PDE's. *Comptes Rendus de l'Academie des Sciences - Series I - Mathematics* 332, 661–668 (2001), [http://dx.doi.org/10.1016/S0764-4442\(00\)01793-6](http://dx.doi.org/10.1016/S0764-4442(00)01793-6)
12. Minion, M.L.: A Hybrid Parareal Spectral Deferred Corrections Method. *Communications in Applied Mathematics and Computational Science* 5(2), 265–301 (2010), <http://dx.doi.org/10.2140/camcos.2010.5.265>
13. Rabenseifner, R., Hager, G., Jost, G.: Hybrid MPI/OpenMP parallel programming on clusters of multi-core SMP nodes. In: 17th Euromicro International Conference on Parallel, Distributed and Network-based processing. pp. 427–436 (2009)
14. Ruprecht, D.: PararealF90: Shared memory pipelined parareal (2017), <http://doi.org/10.5281/zenodo.260095>
15. Schreiber, M., Peddle, A., Haut, T., Wingate, B.: A decentralized parallelization-in-time approach with parareal (2015), <http://arxiv.org/abs/1506.05157>
16. Shu, C.W., Osher, S.: Efficient implementation of essentially non-oscillatory shock-capturing schemes II. *Journal of Computational Physics* 83, 32–78 (1989)