



UNIVERSITY OF LEEDS

This is a repository copy of *A stencil-based implementation of Parareal in the C++ domain specific embedded language STELLA*.

White Rose Research Online URL for this paper:  
<http://eprints.whiterose.ac.uk/90543/>

Version: Accepted Version

---

**Article:**

Arteaga, A, Ruprecht, D [orcid.org/0000-0003-1904-2473](http://orcid.org/0000-0003-1904-2473) and Krause, R (2015) A stencil-based implementation of Parareal in the C++ domain specific embedded language STELLA. *Applied Mathematics and Computation*, 267. pp. 727-741. ISSN 0096-3003

<https://doi.org/10.1016/j.amc.2014.12.055>

---

(c) 2014, Elsevier. Licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International  
<http://creativecommons.org/licenses/by-nc-nd/4.0/>

**Reuse**

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

**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](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# A stencil-based implementation of Parareal in the C++ domain specific embedded language STELLA

Andrea Arteaga<sup>a</sup>, Daniel Ruprecht<sup>b,\*</sup>, Rolf Krause<sup>b</sup>

<sup>a</sup>Center for Climate System and Modeling, ETH Zürich, Universitätstrasse 16, CH-8092 Zurich, Switzerland

<sup>b</sup>Institute of Computational Science, Università della Svizzera italiana, Via Giuseppe Buffi 13, CH-6900 Lugano, Switzerland

---

## Abstract

In view of the rapid rise of the number of cores in modern supercomputers, time-parallel methods that introduce concurrency along the temporal axis are becoming increasingly popular. For the solution of time-dependent partial differential equations, these methods can add another direction for concurrency on top of spatial parallelization. The paper presents an implementation of the time-parallel Parareal method in a C++ domain specific language for stencil computations (STELLA). STELLA provides both an OpenMP and a CUDA backend for a shared memory parallelization, using the CPU or GPU inside a node for the spatial stencils. Here, we intertwine this node-wise spatial parallelism with the time-parallel Parareal. This is done by adding an MPI-based implementation of Parareal, which allows us to parallelize in time across nodes. The performance of Parareal with both backends is analyzed in terms of speedup, parallel efficiency and energy-to-solution for an advection-diffusion problem with a time-dependent diffusion coefficient.

*Keywords:* Parareal, STELLA, GPU, CPU, parallel-in-time, stencil computation, speedup, energy consumption

---

## 1. Introduction

For the numerical solution of initial value problems on massively parallel computers, concurrency is typically introduced by distributing the spatial degrees of freedom of the solution over multiple nodes, processors or cores. In the case of time-dependent partial differential equations, this decomposition is usually induced by a decomposition of the underlying spatial computational domain and therefore this approach is referred to as “spatial parallelism”. The approximate solution is then computed by advancing forward in time from some given initial value by time marching methods, e.g. Runge-Kutta or multi-step schemes, which compute time-step after time-step in a serial fashion.

As the number of cores and processors in state-of-the-art supercomputers continues to rise quickly, the time dimension is more and more becoming a serial bottleneck in at least the following sense: If, for a given problem, the resolution is increased, the increase in computational cost *per time-step* can be compensated by using more cores for the parallelization of the spatial problem—at least, as long as the used method and its implementation show good weak scaling behavior. When the temporal resolution is increased in order to match, e.g., stability or accuracy constraints, however, obviously more time-steps have to be computed to reach the same final time. Here, the resulting increase in computational cost cannot be reduced by spatial parallelization. As discussed in [1], the parallel complexity of time-stepping is always at least  $\Theta(N_t)$ , with  $N_t$  being the number of time-steps. Following [1],  $\Theta(x)$  denotes “a positive quantity whose leading order term is proportional to  $x$ ”. Because of the eventual saturation of spatial strong scaling and the resulting minimum time required per time-step, the lower constant in the estimate is bounded away from zero as the number of processors increases and the minimum computing time per step is approached.

Hence, methods that introduce some degree of concurrency along the time axis are becoming increasingly popular. Very early ideas go back to the 1960’s [2, 3]. In the 1980’s and 1990’s, parabolic and space-time multigrid methods [4,

---

\*Corresponding author; Phone: +41 58 666 4975

Email address: daniel.ruprecht@usi.ch (Daniel Ruprecht)

5] and parallel multiple shooting for initial value problems [6] were studied. A more recent and today very widely studied method is Parareal [7], which is also the method this paper is concerned with. However, a number of other methods exist, e.g. the "parallel implicit time algorithm" PITA [8], "revisionist deferred corrections" or RIDC [9, 10], the "parallel full approximation scheme in space and time" PFASST [11, 12], space-time parallel solvers for periodic problems [13, 14] and newer approaches to space-time multigrid [15, 16]. A recent overview can be found in [17].

There is literature available on the performance of time-parallel methods for a wide range of benchmark problems from different applications, e.g. finance [18], quantum chemistry [19] or plasma-physics [20]. Also, particularly for Parareal, there are a number of papers concerned with analysis of the method [21, 22, 23, 24, 25] or the development of improvements or modifications [26, 27, 28, 29, 30, 31].

Not much literature exists studying, let alone comparing, different implementation strategies for time-parallel methods: There are obviously many different possibilities to map space-time parallelism to a computer consisting of a large number of multi-core nodes, probably also equipped with accelerators. Small-scale benchmarks for an all-MPI implementation of Parareal with spatial parallelization are discussed in [32], but with somewhat inconclusive performance results. Good performance on close to half a million cores of an implementation of PFASST and parallel multigrid in space using only MPI has been demonstrated in [33]. An MPI-based implementation of PFASST combined with a hybrid tree-code on more than two hundred thousand cores is studied in [34]. Smaller scale benchmarks for a hybrid space-time parallel approach, combining OpenMP in time and MPI in space, are reported in [35] for RIDC and in [36] for Parareal. In [37], the original method by Nievergelt [2] is implemented on GPUs. As performance metric, all publications so far consider only speedup, other metrics like energy-to-solution, memory footprint etc. have not yet been analyzed for time parallelization. Some concerns about the memory requirements of time-parallel methods on possible Exascale computers have been voiced in [38].

In this paper, we introduce a stencil-based implementation of Parareal in the recently developed C++ domain specific embedded language STELLA (STencil Loop LAnguage) [39]. STELLA provides functionality for the application of finite difference stencil operators on structured grids, including an OpenMP and CUDA backend to run the computations on a multi-core CPU or a GPU. Stencil computation on structured grids is one of the "seven dwarfs" of high-performance computing defined by Collela in 2004, see e.g. [40]. The backends parallelize the stencil computation in a node across either its cores or by using an attached GPU. The implementation of Parareal in STELLA is then employed to parallelize in time across nodes: Each node holds one "timeslice" in Parareal and performs the corresponding computations exploiting the spatial concurrency from the corresponding STELLA backend while MPI is used to communicate between nodes. In order to optimize data transfer between GPUs, the CUDA-aware MPI implementation MPICH by Cray was used, which supports *GPUDirect*. This technology avoids the need for copying the data from the GPU memory into a temporary buffer of the host memory and transfers instead the data directly to the network interface, thus reducing the latency and increasing the bandwidth. In summary, we employ a paradigm combining a MPI-based distributed memory parallelization in time with an OpenMP or CUDA based shared memory parallelization in space. Performance of Parareal based on both the CPU and the GPU backend is investigated in terms of speedup, parallel efficiency and energy-to-solution and compared to theoretically predicted values. To our knowledge, this is the first time that the performance of a parallel-in-time method in terms of energy-to-solution is analyzed.

## 2. Methods and implementation

Below, we comment briefly on serial time-stepping and then describe Parareal, including a theoretical model for expected speedup and energy overhead. A short introduction of the STELLA language and the corresponding stencil-based implementation of Parareal is given.

### 2.1. Time-marching

Consider an initial value problem of the form

$$u_t = f(u(t), t), \quad u(0) = u_0 \in \mathbb{R}^d, \quad 0 \leq t \leq T, \quad (1)$$

where, in the examples below, the right hand side  $f$  stems from the spatial discretization of a partial differential equation. To fix notation for the introduction of Parareal, let the time-interval  $[0, T]$  be decomposed into  $N_p$  so-called

“time-slices”  $[t_n, t_{n+1}]$ ,  $n = 0, \dots, N_p - 1$  and let  $u_n$  be an approximation of the solution  $u$  of (1) at  $t_n$ , that is  $u_n \approx u(t_n)$ . Note that for Parareal,  $N_p$  is identical to the number of processors used *in time*. It should not be confused with the number of processors *in space* or the total number of processors.

We denote by  $\mathcal{F}_{\delta t}$  a so-called “fine integrator”, typically a higher order method with a fine time-step size  $\delta t$ . That is, for some given approximation  $u_n$  at  $t_n$ , denote by

$$u_{n+1} = \mathcal{F}_{\delta t}(u_n, t_{n+1}, t_n) \quad (2)$$

the approximation provided by running the fine method over the time-slice  $[t_n, t_{n+1}]$  using  $u_n$  as starting value. For convenience, we assume here that  $\delta t$  is such that integration over a time-slice can be done in an integer number of steps. Otherwise, one, e.g., would have to perform a single smaller step at the end to provide  $u_{n+1}$ . Evaluating (2) step-by-step for  $n = 0, \dots, N_p - 1$  corresponds to the “classical” application of  $\mathcal{F}_{\delta t}$  as a time-marching scheme.

### 2.1.1. Serial runtime

We denote by  $N_t$  the total number of fine steps, by  $N_f = N_t/N_p$  the fine steps per time-slice and by  $\tau_f$  the runtime for a single step of the fine propagator. Then, the total runtime for the serial computation of  $u_{N_p}$  using  $\mathcal{F}_{\delta t}$  is

$$C_f = N_t \tau_f = N_p N_f \tau_f. \quad (3)$$

Parallelization in space, for a problem of fixed size, corresponds to reducing the computational time per time-step  $\tau_f$ . However, because strong scaling eventually saturates as communication time becomes dominant, there is a minimum value  $\tau_f^- \leq \tau_f$ . The upper bound  $\tau_f^+$  corresponds to a serial run using only a single core. Note that due to memory requirements such a run might often not be feasible in practice. Hence, for any number of processors used to parallelize in space, it is

$$\tau_f^- \leq \tau_f \leq \tau_f^+ \quad (4)$$

and therefore the parallel complexity of serial time-stepping is  $\Theta(N_t)$ , cf. [1].

## 2.2. Parareal

Parareal replaces serial time-marching by an iterative scheme that offers some concurrency along the time axis. To this end, a so-called “coarse integrator” denoted as  $\mathcal{G}_{\Delta t}$  is introduced, using a time-step  $\Delta t$ . Typically,  $\mathcal{G}_{\Delta t}$  will be of lower order than  $\mathcal{F}_{\delta t}$  and  $\Delta t \gg \delta t$ . Parareal now replaces (2) by the iteration

$$u_{n+1}^{k+1} = \mathcal{G}_{\Delta t}(u_n^{k+1}, t_{n+1}, t_n) + \mathcal{F}_{\delta t}(u_n^k, t_{n+1}, t_n) - \mathcal{G}_{\Delta t}(u_n^k, t_{n+1}, t_n) \quad (5)$$

with  $n = 0, \dots, N_p - 1$  and  $k$  the iteration counter until some convergence criterion is reached. The initial guesses  $u_n^0$  are typically generated by one start-up run of  $\mathcal{G}_{\Delta t}$ . The key point is that when computing the new iterates  $u_n^{k+1}$  out of given values  $u_n^k$  in (5), the computationally expensive evaluations of  $\mathcal{F}_{\delta t}$  can be parallelized: The time required to compute all  $N_p$  values  $\mathcal{F}_{\delta t}(u_n^k, t_{n+1}, t_n)$  is then effectively equal to only a single evaluation of  $\mathcal{F}_{\delta t}$ . Before computing  $\mathcal{G}_{\Delta t}(u_n^{k+1}, t_{n+1}, t_n)$ , a process has to wait until  $u_n^{k+1}$  is available from the processor computing the previous time-slice. However, if the execution of the processes is pipelined properly, the effective wall clock time required for the coarse method in each iteration is only the cost of one call to  $\mathcal{G}_{\Delta t}$ , see [11].

An MPI implementation of the full algorithm is sketched in Algorithm 1. First, every processor computes its own initial guess  $u_n^0$  by repeatedly propagating the initial value  $u_0$  with the coarse propagator. Here, processors handling later time slices have to perform more runs with the coarse method, leading to the “pipelined” execution discussed e.g. in [11] or [41]. After the initialization phase, a fixed number of Parareal iterations is performed, but the corresponding loop could alternatively be terminated by checking some convergence criterion. In each iteration, the processor first runs the fine propagator over its assigned time slice to produce  $\tilde{u}_{p+1}^k$ . Then, it receives the updated initial value  $u_p^{k+1}$  from its predecessor which, at this point, should already have completed its correction in iteration  $k$  and reached the send routine. The processor then applies the coarse propagator to the just received new initial value and computes its own updated final value  $u_{p+1}^{k+1}$  according to (5). This is then sent to the next processor, allowing it to perform its own correction and so on. Note that process  $p = 0$  does not call the MPI library to receive but uses the initial value  $u_0$  as starting value  $u_p^{k+1}$  in all iterations. Accordingly, process  $p = N_p - 1$  does not send but e.g. writes the final value to disk.

---

**Algorithm 1:** Parareal on a time slice  $[t_p, t_{p+1}]$  handled by an MPI process with rank  $p$ .

---

**Data:** Initial condition  $u_0$   
**Result:** Output  
*# Initialization for Parareal: Perform  $p - 1$  runs of the coarse method to produce an initial guess  $u_p^0$*

```

1  $u_p^0 = u_0$ 
2 for  $n = 0 \dots p - 1$  do
3   |  $u_p^0 = \mathcal{G}_{\Delta t}(u_p^0, t_{n+1}, t_n)$ 
4 end
5  $\tilde{u}_{p+1}^0 = \mathcal{G}_{\Delta t}(u_p^0, t_{p+1}, t_p)$ 

# Parareal iteration on process  $p$ .
6 for  $k = 0 \dots k_{max} - 1$  do
7   | # Fine propagator
8   |  $\hat{u}_{p+1}^{k+1} = \mathcal{F}_{\delta t}(u_p^k, t_{p+1}, t_p)$ 
9   | # Receive  $u_p^{k+1}$  from predecessor; for  $p = 0$ , set  $u_0^{k+1} = u_0$ .
10  | if  $p = 0$  then
11  |   |  $u_0^{k+1} = u_0$ 
12  | else
13  |   |  $\text{MPI\_Recv}(u_p^{k+1}, p - 1)$ 
14  | end
15  | # Coarse propagator
16  |  $\tilde{u}_{p+1}^{k+1} = \mathcal{G}_{\Delta t}(u_p^{k+1}, t_{p+1}, t_p)$ 
17  | # Correction.
18  |  $u_{p+1}^{k+1} = \tilde{u}_{p+1}^{k+1} + \hat{u}_{p+1}^{k+1} - \tilde{u}_{p+1}^k$ 
19  | # Send updated value forward; last process does not send.
20  | if  $p < N_p - 1$  then
21  |   |  $\text{MPI\_Send}(u_{p+1}^{k+1}, p + 1)$ 
22  | end
23 end

```

---

### 2.2.1. Parallel runtime

Assume that  $\Delta t$  is chosen such that  $\mathcal{G}_{\Delta t}$  integrates over one time-slice in an integer number of steps  $N_c$  and denote by  $\tau_c$  the cost of a single coarse time-step. The initialization phase then needs  $N_p N_c$  many time steps on the last processor, where it takes the longest. After that, each Parareal iteration costs  $N_c$  coarse steps and  $N_f$  fine steps, where, for the sake of simplicity, we ignore overhead from e.g. communication. In total, the expected time-to-solution for the pipelined version of Parareal with  $K$  iterations and one serial run of  $\mathcal{G}_{\Delta t}$  for start-up is then

$$C_p = N_p N_c \tau_c + K (N_c \tau_c + N_f \tau_f) = (N_p + K) N_c \tau_c + K N_f \tau_f. \quad (6)$$

Note that if the total number of fine time steps  $N_t = N_p N_f$  is increased together with the number of processors  $N_p$ , so that the number of fine steps per slice  $N_f$  remains constant, the effective cost for the serial method in Parareal, i.e. the second term in (6), stays the same – at least as long as the number of iterations  $K$  does not increase. This, however, is typically not the case, but as shown in [23], for diffusive problems the increase in iterations with increasing  $N_p$  is small. Therefore, if a cheap enough coarse method can be found that still leads to good convergence, the cost from having to compute more time steps increases only slowly through the first term in (6) and a mild increase of  $K$ . Note that recent modifications of Parareal towards a full-fledged multigrid in time show no increase in iteration count, at least for simple diffusive problems [15].

### 2.2.2. Speedup

For a fixed resolution, Parareal provides an additional direction of concurrency and can extend strong scaling beyond the saturation of the spatial parallelization: Combining (6) and (3) provides a theoretical bound for the speedup for Parareal

$$S_{\text{bound}}(N_p) = \frac{C_f}{C_p} = \frac{1}{\left(1 + \frac{K}{N_p}\right) \frac{N_c \tau_c}{N_f \tau_f} + \frac{K}{N_p}}, \quad (7)$$

see again also [11]. As the number of processors  $N_p$  increases, while  $N_f$  and  $N_c$  remain fixed, two different bounds on the maximum speedup obtainable by Parareal can be derived from (7):

$$S_{\text{bound}}(N_p) \leq \frac{N_p}{K} \quad \text{and} \quad S_{\text{bound}}(N_p) \leq \frac{N_f \tau_f}{N_c \tau_c}. \quad (8)$$

Neglecting the first term in the denominator in (7) yields the first, neglecting the second term yields the second bound. The first bound illustrates that good performance of Parareal requires rapid convergence in only a few iterations: The maximum parallel efficiency is bounded by  $1/K$ . The second bound stems from the ratio of the runtime of the fine to the coarse method: The latter has to be significantly cheaper in order to allow for reasonable speedup. Using time steps  $\Delta t \gg \delta t$  improves the ratio  $N_f/N_c$ . The ratio  $\tau_f/\tau_c$  can be improved e.g. by using a lower order method for  $\mathcal{G}_{\Delta t}$ , a lower order spatial discretization [28] or even a coarser spatial mesh [42]. The latter approach, however, requires interpolation and restriction operators between the two spatial meshes and the order of interpolation can have a significant influence on convergence [43].

*Parallel Efficiency.* Ideally, a parallel code using  $N_p$  processors would run  $N_p$  times faster than its serial counterpart, therefore providing ideal speedup

$$S_{\text{ideal}}(N_p) = N_p. \quad (9)$$

For Parareal, however, ideal speedup is by design not possible and the achievable speedup is limited by  $S_{\text{bound}}(N_p)$ . Given some measured speedup  $S_{\text{measured}}(N_p)$ , we can therefore consider and compare the following two parallel efficiencies

$$E_{\text{bound}}(N_p) = \frac{S_{\text{bound}}(N_p)}{S_{\text{ideal}}(N_p)} \quad \text{and} \quad E_{\text{measured}}(N_p) = \frac{S_{\text{measured}}(N_p)}{S_{\text{ideal}}(N_p)}. \quad (10)$$

Here,  $E_{\text{bound}}$  is the maximum theoretically possible efficiency for Parareal while  $E_{\text{measured}}$  is the measured efficiency. Ideally, we want  $E_{\text{measured}}$  to closely match  $E_{\text{bound}}$ .

*Energy consumption.* Denote by  $Q_s$  the energy in Joule consumed by a serial run on a single node with runtime  $T_s$  (in seconds) and by  $Q_p$  a parallel run on  $N_p$  nodes with runtime  $T_p$ . The corresponding powers in Watt, equal to Joule per second, that is energy per unit time, can approximately be computed from

$$P_p = \frac{Q_p}{T_p}, \quad \text{and} \quad P_s = \frac{Q_s}{T_s}. \quad (11)$$

In an ideal case, operating  $N_p$  nodes would require exactly  $N_p$  times as much power as operating a single node, so that

$$P_p = N_p P_s. \quad (12)$$

Therefore, the ideal ratio of overhead in terms of consumed energy would be

$$\gamma_{\text{ideal}} = \frac{Q_p}{Q_s} = \frac{N_p}{S_p} = E_p^{-1} \quad (13)$$

where

$$S_p = \frac{T_s}{T_p}, \quad \text{and} \quad E_p = \frac{S_p}{N_p} \quad (14)$$




Figure 1: Diagram of the access pattern of a stencil. In this example the neighboring points around the central point  $u_{i,j,k}$  are accessed for the computation. This pattern is used for instance in the computation of the Laplacian of a three-dimensional field.

denote speedup provided by  $N_p$  nodes for the parallel run and the resulting parallel efficiency. We thus expect the energy ratio of the parallel to the serial run to be approximately the inverse of the parallel efficiency. Note that for a method with 100% efficiency, that is  $E_p = 1$ , there will be no overhead so that  $\gamma_{\text{ideal}} = 1$ . Parareal, however, cannot provide optimal efficiency by design. In order to distinguish between energy overhead as a result of intrinsic suboptimal scaling and additional overhead e.g. from communication, define the intrinsic minimal energy overhead for Parareal as

$$\gamma_{\text{bound}} = \frac{N_p}{S_{\text{bound}}} \quad (15)$$

according to the maximum theoretically possible speedup  $S_{\text{bound}}$ . In Section 3 we will demonstrate that measured energy overhead of Parareal matches  $\gamma_{\text{bound}}$  reasonably well.

### 2.2.3. Convergence

To assess convergence of Parareal, we use the relative defect between the solution at the final time  $T = t_{N_p}$  provided by Parareal after  $k$  iterations  $u_{N_p}^k$  and the solution  $u_{\text{fine}}$  provided by running the fine method serially

$$d^k := \frac{\|u_{N_p}^k - u_{\text{fine}}\|_{\infty}}{\|u_{\text{fine}}\|_{\infty}}. \quad (16)$$

Note that  $u_{\text{fine}}$  is obtained by evaluating (2) step-by-step for  $n = 0, \dots, N_p - 1$ . The relative discretization error of the solution provided by Parareal can be estimated as

$$\varepsilon_{\text{parareal}} := \frac{\|u_{N_p}^k - u_{\text{exact}}\|_{\infty}}{\|u_{\text{exact}}\|_{\infty}} \leq \frac{\|u_{N_p}^k - u_{\text{fine}}\|_{\infty}}{\|u_{\text{exact}}\|_{\infty}} + \frac{\|u_{\text{fine}} - u_{\text{exact}}\|_{\infty}}{\|u_{\text{exact}}\|_{\infty}} = d^k \frac{\|u_{\text{fine}}\|_{\infty}}{\|u_{\text{exact}}\|_{\infty}} + \varepsilon_{\text{fine}} \quad (17)$$

where  $\varepsilon_{\text{fine}}$  is the relative discretization error of the serial fine method. Because  $\|u_{\text{fine}}\| / \|u_{\text{exact}}\| \approx 1$ , if  $d^k \ll \varepsilon_{\text{fine}}$ , we get  $\varepsilon_{\text{parareal}} \approx \varepsilon_{\text{fine}}$ , that is the Parareal solution has essentially the same discretization error as the serial fine integrator.

Note that (16) can of course only be computed in benchmark simulations, where the fine solution is computed, too. In production runs, where one replaces the fine propagator by Parareal, other means have to be used to monitor Parareal's convergence: Possible choices would be the difference between two consecutive iterates, i.e.  $\|u_p^{k+1} - u_p^k\|$  or the residual  $\|\mathcal{F}_{\delta t}(y_{p-1}^k) - y_p^k\|$  at each time slice, cf. [43].

### 2.3. Stencil language

STELLA (STencil Loop Language) is a C++ domain specific embedded language designed and written by T. Gysi et al. [39] within the Swiss HP2C programme (High Performance and High Productivity Computing). It allows

the user to define abstract formulations for stencil formulas and deploys the computation in a high-performance environment by using multiple threads or graphical processors (GPUs). The framework optimizes the data layout, the memory accesses, the concurrency schemes and many other factors for optimal performance. An extensive description can be found in STELLA’s documentation [44]. Although the main motivation behind the development of STELLA is the easy and efficient implementation of finite-difference stencils, its application domain is more general. A stencil computation is defined as any operation applied identically on all grid cells of a multi-dimensional array that requires a local set of neighboring values. The access pattern can be represented in a diagram where every point corresponds to a value used in the computation. Figure 1, for instance, shows the access pattern of a computation that requires all the nearest neighboring values of the point where the operation is applied. One example of such an operator is the second-order discretization of the Laplacian of a three-dimensional field

$$\Delta u_{i,j,k} \approx \frac{u_{i+1,j,k} + u_{i-1,j,k} + u_{i,j+1,k} + u_{i,j-1,k} + u_{i,j,k+1} + u_{i,j,k-1} - 6u_{i,j,k}}{\Delta x^2}. \quad (18)$$

Again for the sake of simplicity, we assume here that step size in all dimensions are identical, i.e.  $\Delta x = \Delta y = \Delta z$ , but this is not an intrinsic limitation of STELLA. However, not just finite-difference operators can be considered stencils, but also other computations fall into this category. For instance, the simple sum of two fields with a multiplicative factor can be considered a stencil and thus executed by STELLA:

$$u_{i,j,k} += \alpha \cdot v_{i,j,k}, \quad (19)$$

where  $(a += b) := (a = a + b)$  following C++ syntax. The access pattern diagram of such an operation would be a simple point at the position  $(i, j, k)$ .

Here, we introduce a stencil-based implementation of Parareal, using STELLA’s backends to parallelize loops over the vector of degrees-of-freedom and benchmark it for the three-dimensional advection diffusion equation (20) introduced below in Section 3. In particular, STELLA parallelizes the evaluation of the finite difference stencils in the spatial discretization, providing us with easy access to shared-memory spatial parallelism. We employ STELLA for the following tasks in Parareal, sketched in Algorithm 1:

- Coarse propagation (Algorithm 1, Lines 3, 5 and 13). Here, we use for  $\mathcal{G}_{\Delta t}$  a forward Euler method with a first order upwind discretization of the convective term  $\mathbf{c} \cdot \nabla u(\mathbf{x}, t)$  in (20) and a second order centered discretization of the diffusive term  $\nu(t)\Delta u(\mathbf{x}, t)$ . Algorithm 2 sketches one call of  $\mathcal{G}_{\Delta t}$ , performing  $N_c$  steps with the coarse method. STELLA performs two different stencil computations in each step: The computation of the right hand side (Algorithm 2, Lines 3-18), that is the evaluation of the two finite difference stencils, is combined into a single complex STELLA stencil (for readability, both stencils are shown independently in Algorithm 2, though). The for-loop at Line 2 is performed and parallelized by STELLA. After the computation of the right hand side, the forward Euler step in Line 21 is again performed and parallelized by STELLA.
- Fine propagation (Algorithm 1, Line 7). Analogously to the coarse propagation, the fine propagation scheme consists of a sequence of time steps. It uses centered fourth order differences for both advection and diffusion and a fourth order accurate explicit Runge-Kutta method (the classical Runge-Kutta-4 scheme). Both the spatial finite difference stencils as well as the summation of the stages in the Runge-Kutta method are done with STELLA, just as for the coarse method.
- Correction (Algorithm 1, Line 14). At the end of the parareal iteration, after having computed the fine solution, having received the updated solution from the previous process, and having computed the coarse solution, the algorithm must apply the update formula (5), which is a sum of three fields. This sum is also computed and parallelized by STELLA.

The spatial stencils sketched in Algorithm 2 correspond to the discretization of a linear advection-diffusion equations as used in Section 3 for benchmarking. However, defining different stencils would easily allow to treat other PDEs as well. Modifying the stencils for the stages would also allow to use other time stepping schemes.



---

**Algorithm 2:** Explicit Euler as coarse propagator for the convection-diffusion problem (20) implemented in STELLA. A first order upwind stencil is used for the convective term, a second order centered stencil for the diffusive term. The fine method consists of a similar implementation of a fourth order Runge-Kutta method with fourth-order space discretization (not shown here).

---

**Data:** Initial value  $u$   
**Result:**  $\tilde{u} = \mathcal{G}_{\Delta t}(u)$

```

1 for  $n = 0 \dots N_c$  do
2   for  $i, j, k \in \mathbb{F}$  do
3      $rhs_{i,j,k} = \nu(t) \cdot (u_{i+1,j,k} + u_{i-1,j,k} + u_{i,j+1,k} + u_{i,j-1,k} + u_{i,j,k+1} + u_{i,j,k-1} - 6 \cdot u_{i,j,k}) / (\Delta x \cdot \Delta x)$ 
4     if  $c_x > 0$  then
5       |  $rhs_{i,j,k} -= c_x \cdot (u_{i,j,k} - u_{i-1,j,k}) / \Delta x$ 
6     else
7       |  $rhs_{i,j,k} -= c_x \cdot (u_{i+1,j,k} - u_{i,j,k}) / \Delta x$ 
8     end
9     if  $c_y > 0$  then
10      |  $rhs_{i,j,k} -= c_y \cdot (u_{i,j,k} - u_{i,j-1,k}) / \Delta x$ 
11    else
12      |  $rhs_{i,j,k} -= c_y \cdot (u_{i,j+1,k} - u_{i,j,k}) / \Delta x$ 
13    end
14    if  $c_z > 0$  then
15      |  $rhs_{i,j,k} -= c_z \cdot (u_{i,j,k} - u_{i,j,k-1}) / \Delta x$ 
16    else
17      |  $rhs_{i,j,k} -= c_z \cdot (u_{i,j,k+1} - u_{i,j,k}) / \Delta x$ 
18    end
19  end
20  for  $i, j, k \in \mathbb{F}$  do
21    |  $\tilde{u}_{i,j,k} = u_{i,j,k} + \Delta t \cdot rhs_{i,j,k}$ 
22  end
23 end
24 Return  $\tilde{u}$ 

```

---

#### 2.4. Mapping to hardware

There are many possible ways to map Parareal (or any other time-parallel method for that matter) to a specific parallel computer. Here, we assign each time-slice in Parareal to one compute node and let STELLA make use of either the node's CPUs through its OpenMP backend or the node's GPU through its CUDA backend to parallelize the stencil computations required to run  $\mathcal{G}_{\Delta t}$  and  $\mathcal{F}_{\delta t}$ . Communication in time in Parareal (lines 11 and 16 in Algorithm 1) between different nodes is done using the MPI library. The setup is sketched in Figure 2: The time slice  $[t_p, t_{p+1}]$  is handled by MPI process  $p$ , which is assigned one full compute node. The resources of the node (multi-core CPU or GPU) are then used for the iteration sketched in Algorithm 1, cf. the discussion in 2.3. After a completed Parareal iteration, the update is sent to the next node/process using MPI. For the CUDA backend, we use the *GPUDirect mode* MPI extension, which allows for a direct communication between GPUs on different nodes without the need for temporary buffers in the system memory. For the OpenMP backend, non-blocking communication (i.e. MPI\_IRecv, MPI\_Isend and corresponding waits) could be used and has been tested, in order to overlap communication and computation. However, the benefits turned out to be negligible. Because non-blocking communication was unavailable for direct GPU communication at the time of writing the software, we decided to make the comparison between CPU and GPU fairer by using a blocking communication scheme for Parareal for both backends.

In summary, we here combine a distributed memory parallelization in time with shared memory parallelization in space using a multi-core CPU or a GPU. Other approaches have been studied as well, e.g. a fully MPI-based space-time parallelization [33], a MPI-in-time parallelization with a hybrid spatial parallelization [34], a hybrid space-time




Figure 2: Mapping of Parareal-STELLA to multiple compute nodes on Piz Daint: Each of the  $N_p$  time-slices is assigned to a node. The stencil computations required to advance the solution over this time-slices are carried out by STELLA, represented by the blue cube, using the coarse integrator  $\mathcal{G}_{\Delta t}$  or using the fine integrator  $\mathcal{F}_{\delta t}$  (solid arrows). The dashed arrows represent the communications from one node to the other, which occur once per Parareal iteration.

approach combining OpenMP in time with MPI in space [36, 35] or a GPU-based implementation [37] Note that while feasibility and efficiency has been demonstrated for all these different strategies, so far there are no comparisons of their relative strengths and weaknesses. Some implementation strategies like PGAS languages or one-sided communication have, to our knowledge, not yet been explored at all for parallel-in-time methods.

### 3. Numerical results

For the benchmarks reported below, we use the implementation of Parareal in STELLA described in Section 2 to solve the three-dimensional advection-diffusion equation with a time-dependent diffusion coefficient

$$u_t(\mathbf{x}, t) + \mathbf{c} \cdot \nabla u(\mathbf{x}, t) = v(t)\Delta u(\mathbf{x}, t), \quad 0 \leq t \leq T, \quad (20)$$

with  $\mathbf{c} \in \mathbb{R}^3$  and  $v(t) \in \mathbb{R}^+$ . The spatial domain is  $\Omega = [0, 1]^3$  with periodic boundary conditions. The initial value is

$$u(\mathbf{x}, 0) = u_0(\mathbf{x}) = \sin(2\pi x) \sin(2\pi y) \sin(2\pi z). \quad (21)$$

For now, let  $\mathbf{c} = 0$  and assume a solution of the form

$$u(\mathbf{x}, t) = a(t)u_0(\mathbf{x}). \quad (22)$$

Because  $\Delta u_0(\mathbf{x}) = -12\pi^2 u_0(\mathbf{x})$ , this results in the initial value problem

$$a'(t) = -12\pi^2 v(t)a(t), \quad a(0) = 1 \quad (23)$$

for the coefficient  $a$  in (22). Variation of parameters gives the solution

$$a(t) = \exp\left(-\int_0^t 12\pi^2 v(s) ds\right). \quad (24)$$

Note that for  $v \equiv 1$ , this reduces to  $a(t) = \exp(-12\pi^2 t)$ . For a simple profile of the form

$$v(t) = v_0 + \frac{v_0}{2} \sin(\omega t), \quad v_0 > 0 \quad (25)$$

we get

$$a(t) = -12\pi^2 \left( v_0 t - \frac{v_0}{2\omega} (\cos(\omega t) - 1) \right). \quad (26)$$

For an arbitrary  $\mathbf{c} \in \mathbb{R}$ , the full solution then is obtained as

$$u(\mathbf{x}, t) = a(t)u_0(\mathbf{x} - \mathbf{c}t). \quad (27)$$

In the tests reported below, we use  $v_0 = 0.1$ , with  $\omega = 0$  (no time-dependance in  $v$ ) and  $\omega = 100$  (rapidly oscillating  $v$ ). Furthermore, the advection velocity is set to  $\mathbf{c} = (1, 1, 1)$  and the simulation is run until  $T = 0.1$ .

Here, to allow for precise benchmarking of STELLA-Parareal, we choose a rather simple linear problem where an analytic solution is available. This allows to compute and compare the defect of Parareal with the discretization error of the coarse and fine propagator, cf. the discussion in 3.3.1. In production runs, this solution won't be known and other means have to be used to monitor convergence of Parareal and to decide when to stop iterating, see also 2.2.3.

### 3.1. Discretization

The fine propagator  $\mathcal{F}_{\delta t}$  uses fourth order centered finite differences for both advection and diffusion on a spatial mesh with  $128^3$  points. A Runge-Kutta-4 method with a time-step  $\delta t = \frac{T}{215}$  is used for integration, resulting in an overall relative error at the end of the simulation of  $\varepsilon_{\text{fine}} \approx 1.9 \times 10^{-7}$  for  $\omega = 0$  and of  $4.8 \times 10^{-6}$  for  $\omega = 100$ . The coarse method  $\mathcal{G}_{\Delta t}$  uses a first order upwind stencil for advection and a second order centered stencil for diffusion, again on a spatial mesh with  $128^3$  points. A first order forward Euler method with time-step  $\Delta t = \frac{T}{211}$  is used, resulting in an overall relative error at  $T = 0.1$  of about  $\varepsilon_{\text{coarse}} \approx 4.5 \times 10^{-2}$  for both  $\omega = 0$  and  $\omega = 100$ . The time-steps are chosen such that the temporal discretization error is about the same as the spatial discretization error, i.e., further refinement in time without simultaneously refining the spatial mesh does no longer improve the accuracy of the solution.

Note that if the time step of  $\mathcal{F}_{\delta}$  is close to the stability limit (and not, as in the case considered here, determined by accuracy considerations), the coarse propagator will very likely have to be an implicit method in order to allow for a coarse time step  $\Delta t \gg \delta t$ . Alternatively, it is possible to use similar time steps for both coarse and fine method and reduce the computational cost of  $\mathcal{G}_{\Delta t}$  by other means, e.g. aggressive coarsening of the spatial discretization.

### 3.2. Hardware

All runs are performed on the Cray XC30 "Piz Daint" at the Swiss National Supercomputing Centre (CSCS). Piz Daint consists of  $5 \cdot 272$  hybrid compute nodes, each with an octo-core Intel Xeon Processor E5-2670 and an NVIDIA Tesla K20X GPU, connected through the high-bandwidth, low-latency Aries routing and communications ASIC. The system provides a rich set of programming environments and software, including the GNU, Cray, Intel and PGI compilers and the Cray derivative of the MPICH2 implementation of the MPI library. In our experiments we used the GNU compiler in its version 4.8 and applied the standard optimization features provided by the switch `-O3`.

### 3.3. Results

In order to give a fair and meaningful assessment of the speedup provided by Parareal, it is important to make sure that it provides comparable accuracy as the time-serial fine solution. Therefore, in 3.3.1, some precursory tests are run to ensure that a configuration is found for Parareal for which reported speedups do actually compare solutions of comparable accuracy, that is  $\varepsilon_{\text{fine}} \approx \varepsilon_{\text{parareal}}$ , cf. the discussion in 2.2.3. This aspect is sometimes not considered, but it should be noted that if the time-parallel solution does not provide comparable accuracy as running the fine method serially, reporting speedups by comparing their runtimes becomes essentially meaningless. Sufficiently many iterations of Parareal have to be performed in order to reduce the defect  $d^k$  defined in (16) below the relative discretization error  $\varepsilon_{\text{fine}}$  of the fine serial solution. Here, we have an analytical solution at hand to compute  $\varepsilon_{\text{fine}}$ .

Moreover, for a fair comparison of STELLA's CPU and GPU backends, the optimal number of threads to be used on one node by the CPU version is determined. Performance results for the time parallelization are shown in 3.3.3 while in 3.3.4 Parareal's efficiency in terms of energy is analyzed.

#### 3.3.1. Convergence of Parareal

Figure 3 shows  $d^k$  versus the number of iterations for  $N = 8, 32, 128$  concurrently computed time-slices and for  $\omega = 0$  and  $\omega = 100$ . Parareal converges rapidly in all cases and neither the number of time-slices nor the time-dependent diffusion coefficient has a significant impact. This matches the analysis for constant-coefficient diffusive problems in [23]. The marginal affect of a time-dependent diffusion coefficient is in line with the study of Parareal's convergence for space- and time-dependent diffusion coefficients in [45]. See also the comments in 2.2.1. Note that both the results presented here as well as the papers mentioned above are concerned with linear problems: For complex nonlinear problems, less favorable convergence behavior is possible.

After three iterations, the defect between Parareal and  $\mathcal{F}_{\delta t}$  is in all cases significantly smaller than the discretization error of  $\varepsilon_{\text{fine}} \approx 4.8 \times 10^{-6}$  for  $\omega = 100$  and about the order of magnitude of the fine error for  $\omega = 0$ . Hence, for  $k = 3$  the parallel and serial solution have essentially the same error and we thus report below speedups for Parareal with  $\omega = 100$  and  $k = 3$  iterations.

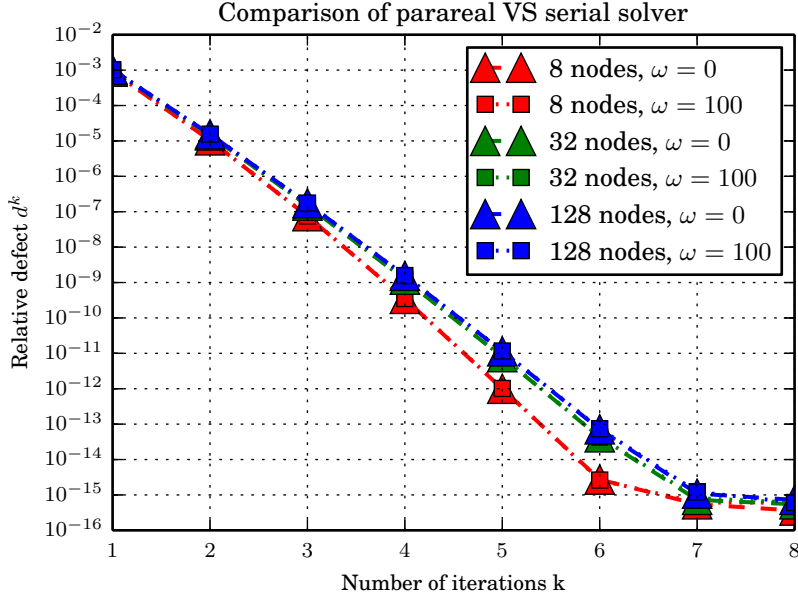


Figure 3: Relative defect  $d^k$ , see (16), between Parareal and the serial fine solution versus number of iterations. Shown are results for different numbers of nodes, that is different numbers  $N_p$  of concurrently computed time-slices, and for values of  $\omega = 0$  (no time-dependance in the diffusion coefficient) and  $\omega = 100$  (rapidly oscillating diffusion coefficient).

### 3.3.2. Optimal number of threads for STELLA on CPUs

To provide a fair serial reference for the CPU version of STELLA, we have to find the optimal number of threads to be used for the stencil computation on a node. Therefore, we run Parareal and keep the number of time-slices, and thus the number of nodes, fixed and vary the number of OpenMP threads used for the spatial parallelization in each node from 1 to 8 (because Piz Daint has 8 cores per node).

Figure 4 shows the speedup provided by using more threads for STELLA, measured against Parareal using only a single thread on each node for computation of the spatial stencils. Although speedup from using more threads is less than optimal, it increases up to all eight cores in a node and therefore, in the CPU benchmarks reported below, STELLA’s CPU backend is always using 8 threads on each node. It is interesting to note that the speedup from STELLA seems to be hardly affected by the number of nodes used for Parareal: The speedup curves for Parareal across 8 nodes and across 128 nodes are nearly identical. This suggests that here spatial and temporal parallelization function more or less independently and do not interfere with each other. The reason is probably that Parareal essentially uses  $\mathcal{F}_{\delta t}$  and  $\mathcal{G}_{\Delta t}$  as black-boxes with very little interweaving of space- and time solver. In case of more closely intertwined space-time parallelization, better overall speedup can be achieved, but a more complex interplay of the efficiencies of both approaches can be encountered, see [46].

### 3.3.3. Speedup and runtimes for Parareal

Figure 5a shows the actual runtimes for the CPU and GPU versions of Parareal, depending on the number of nodes. The horizontal blue and red line indicate the runtimes of the serial reference runs. The GPU version is constantly about a factor of 4.5 faster than the CPU version, but both curves are more or less parallel to each other, suggesting that achieved speedups should be similar for both backends. This is confirmed by Figure 5b, which shows the speedup provided by Parareal depending on the number of nodes, i.e. the number of time slices. Red triangles denote GPU runs, blue triangles CPU runs, and the lines show the theoretical maximum speedup according to (7). Speedup is measured against a time-serial run of the fine method using a single node and STELLA’s CPU or GPU backend, respectively. Performance of Parareal in terms of speedup is nearly oblivious against the choice of the STELLA backend: Speedups are almost identical for the CPU and GPU version.

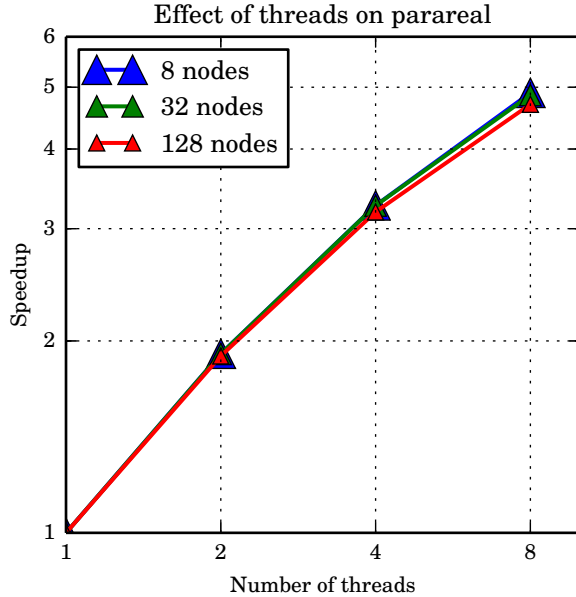


Figure 4: Speedup from the spatial parallelization using STELLA’s CPU backend for multi-threaded stencil computations across different numbers of nodes used for Parareal.

Table 1 summarizes speedups  $S_{\text{bound}}$  and  $S_{\text{measured}}$  as well as efficiencies  $E_{\text{bound}}$  and  $E_{\text{measured}}$  defined in (10) for the CPU and GPU version of Parareal-STELLA depending on the number of nodes. Expected speedup is computed from (7), with the coarse-to-fine ratio  $\tau_c/\tau_f$  measured from serial runs and therefore slightly different for both versions. Up to 128 nodes, the CPU version provides speedups and efficiencies very close to the theoretical bounds. The slightly super-optimal value for 64 nodes comes from the inaccuracy in the formula for  $S_{\text{bound}}$  by computing the ratio  $\tau_c/\tau_f$  experimentally. The GPU version also provides speedups and efficiencies very close to the theoretical bound, with a little drop off for large node numbers. This is probably due to the fact that actual runtimes here are much smaller so that small overheads have a larger effect.

### 3.3.4. Energy consumption

The energy consumption of a supercomputer constitutes a large part of its effective cost. Moreover, there is an increasing concern about the impact of scientific computing on the environment and about the sustainability of the growth of the performance of supercomputers. For these reasons, in the last decade, energy consumption has become a major issue in high-performance computing. A prominent sign of the increasing importance is that, besides the Top500 ranking of the most powerful supercomputers [47], a new ranking called Green500 publishes a list of the most energy-efficient supercomputers [48]. Piz Daint, the machine of CSCS where we carried our experiments, entered this ranking at position four, being the most energy-efficient petaflop machine. On this machine, a comprehensive set of power management facilities has been installed by Cray [49] and these tools have been tested in a study to assess the energy efficiency of multiple applications [50]. Assessing the performance of algorithms not only in terms of time-to-solution but also energy-to-solution is important, but for time-parallel methods there seem to be no such studies yet.

Figure 6 shows the total energy consumed by the supercomputer Piz Daint to perform the time-serial integration on a single node and the Parareal algorithm on multiple nodes. Energy consumed by the node, i.e. mainly main memory and CPU, is marked in red. Blue indicates energy consumed by the network while purple is energy consumed by the cooling system. Finally, for the GPU backend, the energy consumed by the GPU itself is marked in green, and labeled as *Device* energy. The GPU backend has a factor of three to four lighter energy consumption throughout (note that differently scaled y axes). Going from a single node (time serial) to four or more nodes (Parareal) results in a significant increase in consumed energy (about a factor of three) for both the CPU and the GPU version. After

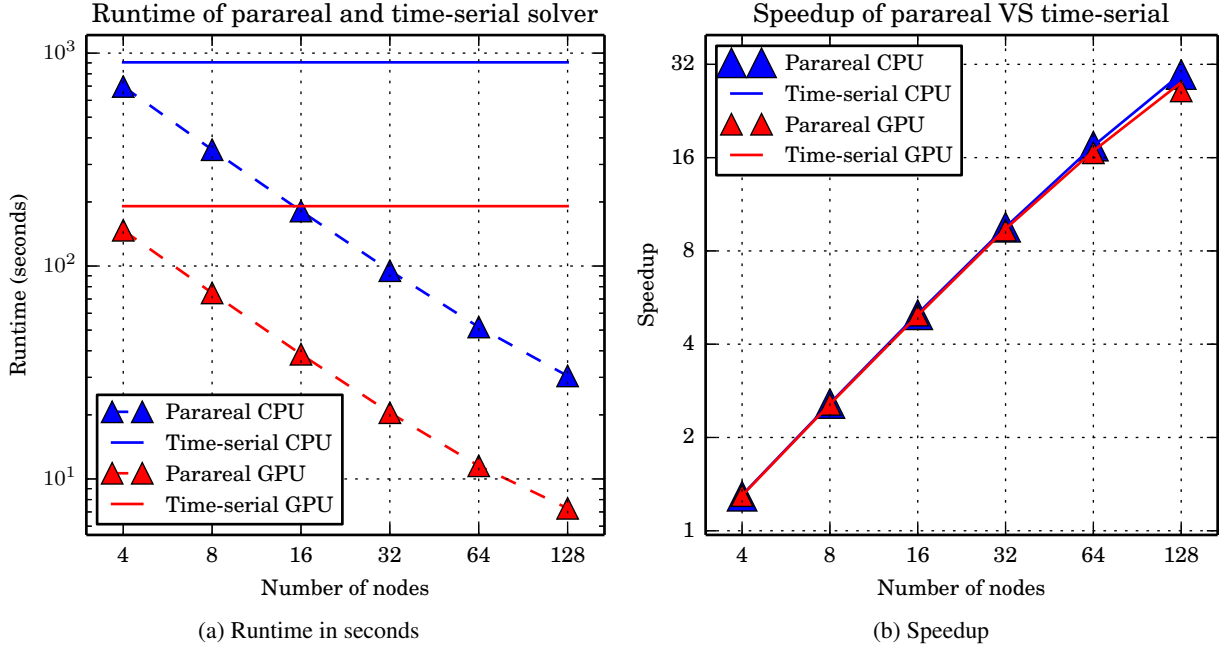


Figure 5: Runtime in seconds (left) and speedup (right) of Parareal depending on the number of nodes using the nodes’ GPU (red) or CPU (blue) for stencil evaluations. Speedup is measured against  $\mathcal{F}_{\delta t}$  run serially on a single node, using the CPU or the GPU, respectively. The solid lines in (b) show the maximum theoretical speedup  $S_{bound}(N_p)$ .

that, the energy consumption from using more nodes for Parareal increases at a much slower pace: The reason is that e.g. from Parareal across four to Parareal across eight nodes, the runtime is reduced by about a factor of two (see e.g. Table 1, the speedup on eight nodes is double the speedup on four nodes). Therefore, while running Parareal on eight nodes doubles power consumption, because it also about halves runtime the energy-to-solution stays about the same. Thus, the difference in energy-to-solution between Parareal runs using different numbers of nodes is governed by Parareal’s intrinsic parallel efficiency while the difference between Parareal and the serial propagator is naturally governed by Parareal’s parallel efficiency in comparison to  $\mathcal{F}_{\delta t}$ . For larger numbers of concurrently computed time slices, Parareal’s intrinsic efficiency start to go down, too, and energy-to-solution increases again when compared with Parareal on four nodes.

In Table 2 we summarize the power consumption per node of different components of the machine when performing the simulations. Power consumption depends heavily on the kind of computation performed, on its computational intensity, on the usage of memory and on other factors. Nevertheless, the runs with the time-serial integration scheme and those with Parareal showed very similar per-node power consumption, with variations on the order of a few percent. This supports the relatively coarse assumptions made for the derivation of  $\gamma_{bound}$  above and suggests that Parareal’s overhead in terms of energy should, as expected, be closely linked to its parallel efficiency with little impact from additional overhead e.g. from communication.

For both backends, the major portion of power (three quarters or more) is required by the computing devices, i.e. the node for the OpenMP and the node and the device for the CUDA backend. The power required by the network is fixed at a rate of 25 W/node, while the cooling system constantly consumes about 14 W/node. It is interesting to note that the GPU backend actually has a significantly higher power consumption than the CPU version, so that its lower energy-to-solution is only due to its noticeably smaller runtimes.

Figure 7 plots the measured energy overhead of Parareal, that is

$$\gamma_{\text{measured}} = \frac{Q_p}{Q_s}, \quad (28)$$

cf. 2.2.2. In addition, the value  $\gamma_{bound}$  defined in (15) is shown, that is the overhead that is to be expected from the

CPU version					GPU version				
# Nodes	$S_{\text{bound}}$	$S_{\text{measured}}$	$E_{\text{bound}}$	$E_{\text{measured}}$	# Nodes	$S_{\text{bound}}$	$S_{\text{measured}}$	$E_{\text{bound}}$	$E_{\text{measured}}$
4	1.3	1.3	32.6	32.3	4	1.3	1.3	32.5	32.4
8	2.6	2.6	32.2	32.2	8	2.6	2.6	32.0	32.0
16	5.0	5.0	31.4	31.1	16	5.0	5.0	31.0	31.0
32	9.5	9.5	29.8	29.8	32	9.4	9.3	29.4	29.2
64	17.4	17.6	27.3	27.5	64	16.8	16.6	26.3	25.9
128	29.8	29.7	23.3	23.2	128	28.0	26.3	21.9	20.5

Table 1: Speedup and parallel efficiency in percent of the CPU and GPU version of Parareal-STELLA depending on the number of nodes used for the time-parallelization. The theoretical bounds for speedup and efficiency according to (7) are denoted as  $S_{\text{bound}}$  and  $E_{\text{bound}}$ . The measured values are  $S_{\text{measured}}$  and  $E_{\text{measured}}$ . The coarse-to-fine runtime ratios  $\tau_c/\tau_f$  in (7) are determined experimentally by running coarse and fine method on a single node.

	CPU		GPU	
	Consumption	Percentage	Consumption	Percentage
Node	133 W	77 %	70 W	29 %
Network	25 W	15 %	25 W	10 %
Blower	14 W	8 %	14 W	6 %
Device	—	—	135 W	55 %
<b>Total</b>	172 W	100 %	245 W	100 %

Table 2: Power consumption per node of the considered simulations with different backends. For each component the power consumption and the its percentage of the total is shown. These values are the result of empirical measurements performed on Piz Daint. The Node and Device (GPU) power consumption fluctuates depending on the intensity of computations, while the Network and Blower power consumption is constant. The difference in power consumption between the time-serial integration and Parareal with different numbers of nodes is negligible and thus only one set of values is shown.

sub-optimal scaling of Parareal. For both the CPU and the GPU backends, the actual energy overhead is for some reason slightly better than the theoretical minimum, but gives a reasonable estimate. The reasons most likely are the inaccuracies in both the energy measurement as well as the experimental computation of the term  $\tau_f/\tau_c$  in the speedup estimate and thus the theoretically expected parallel efficiency that enters into  $\gamma_{\text{bound}}$ . A more comprehensive analysis would have to sample over a large ensemble of runs and invoke statistical measures here, in order to provide a more detailed estimate. This is, however, beyond the scope of the current paper. Nevertheless, Figure 7 strongly suggests that the significant overhead of Parareal in terms of energy depicted in Figure 6 can to a large degree be attributed to its rather harsh limit on parallel efficiency: Therefore, improving the parallel efficiency of Parareal in particular but also related parallel-in-time methods in general will simultaneously significantly improve efficiency in terms of energy consumption.

#### 4. Conclusions and outlook

The paper presents a stencil-based implementation of the time-parallel Parareal method in the C++ domain specific embedded language STELLA. STELLA is the acronym for STencil Loop Language and is designed for easy and efficient implementation of stencil-based computations on different platforms. It provides different backends, e.g. using OpenMP or CUDA, for the parallelized evaluation of stencils on structured grids. In the here presented combination, STELLA is used for the spatial intra-node parallelization, using either the node’s GPU or multi-core CPU, depending on the backend. Parareal is then added to parallelize in time across nodes, using MPI for communication of volume data within the Parareal iteration. This corresponds to a paradigm using distributed memory parallelization in time and shared memory parallelization in space.

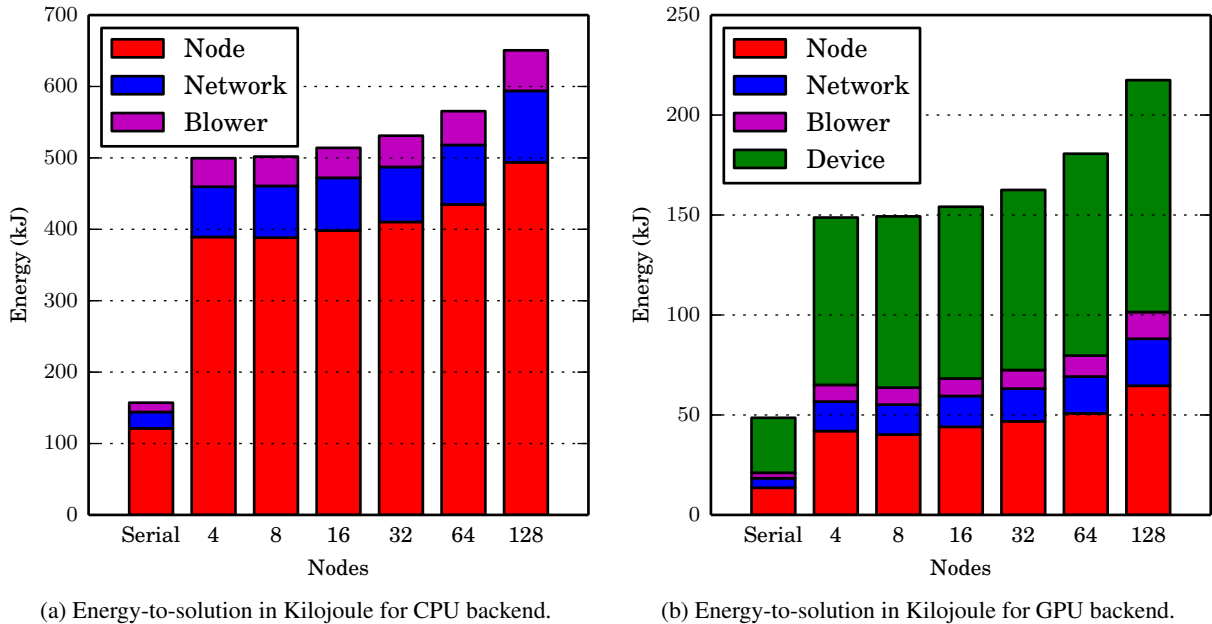


Figure 6: Total energy consumed by Parareal-STELLA with CPU backend (left) and GPU backend (right) depending on the number  $N_p$  of nodes or concurrently computed time-slices. The red portion is the energy consumption of the CPU, the main memory and all other components of the compute node. The blue portion corresponds to the energy consumption of the interconnect network. The purple portion represents the consumption of the cooling system of the machine. The green portion is the consumption of the GPU, which is shown only in the GPU version.

Performance of the stencil-based Parareal implementation is assessed in terms of runtimes, speedup and energy consumption. A three-dimensional advection-diffusion equation with a time-dependent diffusion coefficient is used as a benchmark problem. First, convergence of Parareal is analyzed and a setup is defined where the time-parallel solution provides accuracy comparable to the time-serial reference. In line with previous results, it is found that a time-dependent diffusion coefficient has only marginal influence on the convergence of Parareal. Then, the performance of both the OpenMP and CUDA backend are investigated and compared. For both versions, measured speedups closely match theoretically predicted values, illustrating the efficiency of the approach. Speedup from Parareal is found to be roughly identical for both backends, with the GPU backend being consistently about a factor of 4.5 faster in terms of runtimes. In terms of energy-to-solution, Parareal is found to cause significant overhead, with the GPU version requiring significantly less energy than the CPU version for both time-serial and time-parallel runs. To distinguish between energy overhead caused by the intrinsically sub-optimal parallel efficiency of Parareal and additional overhead e.g. from communication, an estimate for the energy overhead is derived from the theoretical bound on speedup and compared to the measured overhead. It is found that the significant rise in energy-to-solution from Parareal can be attributed to a very large degree to its harsh bound on parallel efficiency. This result is of interest for other "parallel across the time-steps" approaches as well (e.g. PITA [8], PFASST [12] or time multigrid [15] or [16]) because it suggests that improving parallel efficiency of such methods will also be key to improve their efficiency in terms of energy. Derivation of a more detailed model for the energy consumption of Parareal or other time-parallel methods would be a very interesting next step.

Moreover, in order to have an analytic solution available, the benchmark considered here is a relatively simple linear advection-diffusion problem. Having been designed to be used in the dynamical core of the numerical weather prediction and climate model COSMO, STELLA can be employed for complex nonlinear equations, too. Studying the performance of Parareal-STELLA for such a problem is planned for future work.

This paper considered spatial parallelism across one single node/GPU combined with time-parallelism across



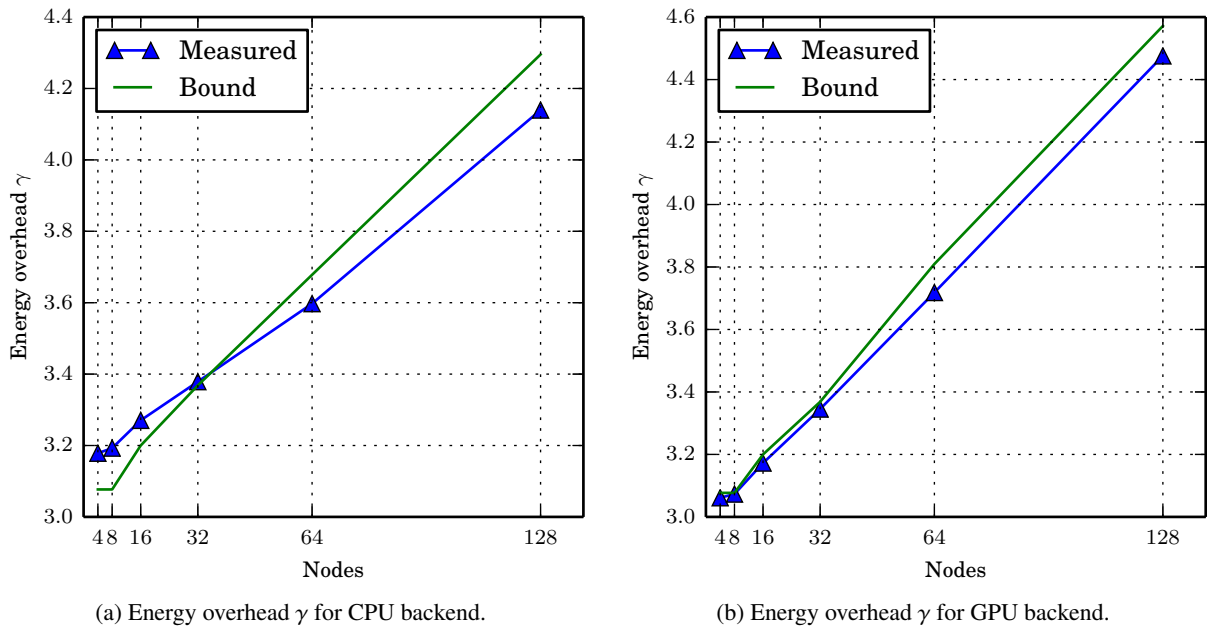


Figure 7: Energy overhead of Parareal as defined in (28), i.e. the energy-to-solution required by Parareal divided by energy-to-solution for the time-serial fine method. The expected overhead according to (15), represented by the green line without markers, indicates the overhead that is unavoidable due to Parareal’s intrinsically less than optimal parallel scaling. This bound contains experimentally measured terms and is thus not infinitely accurate, which is the reason for the apparent super-optimal behavior.

nodes. For larger problems, where a time-slice does not fit on a single node, the Parareal-STELLA benchmarks can be extended to use distributed memory in time and a hybrid parallelization in space, similar to what has been explored for a particle discretization in [34]. In the other direction, for problems too small to fully utilize a single GPU or many-core CPU with only spatial parallelism, a shared-memory space-time parallel approach could extend the degree of concurrency and be used to improve utilization of a single node. So far, this approach has not been explored, but the framework studied here would provide a good starting point for such a study. Other studies exist that use shared memory in time on a single node [28] or in combination with a distributed memory parallelization in space [35, 36]. Moreover, pure MPI based approaches have also been studied e.g. in [32] or [33]. A meticulous comparison of different implementation strategies for space-time parallelization in terms of important metrics like speedup, energy consumption, memory requirements, etc. would be a very interesting future research direction.

## Acknowledgments

The authors thankfully acknowledge the support of the Swiss PASC (“Platform for Advanced Scientific Computing”) initiative. They would also like to thankfully acknowledge the many helpful discussions with Oliver Fuhrer from MeteoSwiss, Carlos Osuna from C2SM and Will Sawyer, Benjamin Cumming and Gilles Fourestey from the Swiss National Supercomputing Centre CSCS.

## References

- [1] G. Horton, S. Vandewalle, P. Worley, An algorithm with polylog parallel complexity for solving parabolic partial differential equations, *SIAM Journal on Scientific Computing* 16 (3) (1995) 531–541. URL <http://dx.doi.org/10.1137/0916034>
- [2] J. Nievergelt, Parallel methods for integrating ordinary differential equations, *Commun. ACM* 7 (12) (1964) 731–733. URL <http://dx.doi.org/10.1145/355588.365137>

- [3] W. L. Miranker, W. Liniger, Parallel methods for the numerical integration of ordinary differential equations, *Mathematics of Computation* 21 (99) (1967) 303–320.  
URL <http://dx.doi.org/10.1090/S0025-5718-1967-0223106-8>
- [4] W. Hackbusch, Parabolic multi-grid methods, *Computing Methods in Applied Sciences and Engineering*, VI (1984) 189–197.  
URL <http://dl.acm.org/citation.cfm?id=4673.4714>
- [5] G. Horton, The time-parallel multigrid method, *Communications in Applied Numerical Methods* 8 (9) (1992) 585–595.  
URL <http://dx.doi.org/10.1002/cnm.1630080906>
- [6] M. Kiehl, Parallel multiple shooting for the solution of initial value problems, *Parallel Computing* 20 (3) (1994) 275 – 295.  
URL [http://dx.doi.org/10.1016/S0167-8191\(06\)80013-X](http://dx.doi.org/10.1016/S0167-8191(06)80013-X)
- [7] J.-L. Lions, Y. Maday, G. Turinici, A "parareal" in time discretization of PDE's, *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics* 332 (2001) 661–668.  
URL [http://dx.doi.org/10.1016/S0764-4442\(00\)01793-6](http://dx.doi.org/10.1016/S0764-4442(00)01793-6)
- [8] C. Farhat, M. Chandesris, Time-decomposed parallel time-integrators: theory and feasibility studies for fluid, structure, and fluid-structure applications, *International Journal for Numerical Methods in Engineering* 58 (9) (2003) 1397–1434.  
URL <http://dx.doi.org/10.1002/nme.860>
- [9] A. J. Christlieb, C. B. Macdonald, B. W. Ong, Parallel high-order integrators, *SIAM Journal on Scientific Computing* 32 (2) (2010) 818–835.  
URL <http://dx.doi.org/10.1137/09075740X>
- [10] A. J. Christlieb, R. D. Haynes, B. W. Ong, A parallel space-time algorithm, *SIAM Journal on Scientific Computing* 34 (5) (2012) C233–C248.  
URL <http://dx.doi.org/10.1137/110843484>
- [11] M. L. Minion, A hybrid parareal spectral deferred corrections method, *Communications in Applied Mathematics and Computational Science* 5 (2) (2010) 265–301.  
URL <http://dx.doi.org/10.2140/camcos.2010.5.265>
- [12] M. Emmett, M. L. Minion, Toward an efficient parallel in time method for partial differential equations, *Communications in Applied Mathematics and Computational Science* 7 (2012) 105–132.  
URL <http://dx.doi.org/10.2140/camcos.2012.7.105>
- [13] P. Arbenz, A. Hildebrand, D. Obrist, A parallel space-time finite difference solver for periodic solutions of the shallow-water equation, in: R. Wyrzykowski, J. Dongarra, K. Karczewski, J. Waśniewski (Eds.), *Parallel Processing and Applied Mathematics*, Vol. 7204 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2012, pp. 302–312.  
URL [http://dx.doi.org/10.1007/978-3-642-31500-8\\_31](http://dx.doi.org/10.1007/978-3-642-31500-8_31)
- [14] P. Arbenz, D. Hupp, D. Obrist, A parallel solver for the time-periodic Navier-Stokes equations, in: R. Wyrzykowski, J. Dongarra, K. Karczewski, J. Waśniewski (Eds.), *Parallel Processing and Applied Mathematics*, Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2014, pp. 291–300.  
URL [http://dx.doi.org/10.1007/978-3-642-55195-6\\_27](http://dx.doi.org/10.1007/978-3-642-55195-6_27)
- [15] S. Friedhoff, R. D. Falgout, T. V. Kolev, S. MacLachlan, 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.  
URL <http://www.osti.gov/scitech/servlets/purl/1073108>
- [16] M. J. Gander, M. Neumüller, Analysis of a time multigrid algorithm for dg-discretizations in time (2014).  
URL <http://arxiv.org/abs/1409.5254>
- [17] M. J. Gander, 50 years of time parallel time integration, submitted (2014).  
URL <http://www.unige.ch/~gander/Preprints/50YearsTimeParallel.pdf>
- [18] G. Bal, Y. Maday, A "parareal" time discretization for non-linear PDE's with application to the pricing of an american put, in: L. Pavarino, A. Toselli (Eds.), *Recent Developments in Domain Decomposition Methods*, Vol. 23 of Lecture Notes in Computational Science and Engineering, Springer Berlin, 2002, pp. 189–202.  
URL [http://dx.doi.org/10.1007/978-3-642-56118-4\\_12](http://dx.doi.org/10.1007/978-3-642-56118-4_12)
- [19] Y. Maday, G. Turinici, Parallel in time algorithms for quantum control: Parareal time discretization scheme, *Int. J. Quant. Chem.* 93 (3) (2003) 223–228.  
URL <http://dx.doi.org/10.1002/qua.10554>
- [20] D. Samaddar, D. E. Newman, R. S. Sánchez, Parallelization in time of numerical simulations of fully-developed plasma turbulence using the parareal algorithm, *Journal of Computational Physics* 229 (2010) 6558–6573.  
URL <http://dx.doi.org/10.1016/j.jcp.2010.05.012>
- [21] G. Bal, On the convergence and the stability of the parareal algorithm to solve partial differential equations, in: R. Kornhuber, et al. (Eds.), *Domain Decomposition Methods in Science and Engineering*, Vol. 40 of Lecture Notes in Computational Science and Engineering, Springer, Berlin, 2005, pp. 426–432.  
URL [http://dx.doi.org/10.1007/3-540-26825-1\\_43](http://dx.doi.org/10.1007/3-540-26825-1_43)
- [22] G. A. Staff, E. M. Rønquist, Stability of the parareal algorithm, in: R. Kornhuber, et al. (Eds.), *Domain Decomposition Methods in Science and Engineering*, Vol. 40 of Lecture Notes in Computational Science and Engineering, Springer, Berlin, 2005, pp. 449–456.  
URL [http://dx.doi.org/10.1007/3-540-26825-1\\_46](http://dx.doi.org/10.1007/3-540-26825-1_46)
- [23] M. J. Gander, S. Vandewalle, Analysis of the parareal time-parallel time-integration method, *SIAM Journal on Scientific Computing* 29 (2) (2007) 556–578.  
URL <http://dx.doi.org/10.1137/05064607X>
- [24] J. Reynolds-Barredo, D. Newman, R. Sanchez, D. Samaddar, L. Berry, W. Elwasif, Mechanisms for the convergence of time-parallelized, parareal turbulent plasma simulations, *Journal of Computational Physics* 231 (23) (2012) 7851 – 7867.  
URL <http://dx.doi.org/10.1016/j.jcp.2012.07.028>
- [25] M. J. Gander, E. Hairer, Analysis for parareal algorithms applied to Hamiltonian differential equations, *Journal of Computational and Applied*

- Mathematics 259, Part A (0) (2014) 2 – 13, proceedings of the Sixteenth International Congress on Computational and Applied Mathematics (ICCAM-2012), Ghent, Belgium, 9-13 July, 2012.  
URL <http://dx.doi.org/10.1016/j.cam.2013.01.011>
- [26] C. Farhat, J. Cortial, C. Dastillung, H. Bavestrello, Time-parallel implicit integrators for the near-real-time prediction of linear structural dynamic responses, *International Journal for Numerical Methods in Engineering* 67 (2006) 697–724.  
URL <http://dx.doi.org/10.1002/nme.1653>
- [27] M. Gander, M. Petcu, Analysis of a Krylov subspace enhanced parareal algorithm for linear problems, *ESAIM: Proc.* 25 (2008) 114–129.  
URL <http://dx.doi.org/10.1051/proc:082508>
- [28] D. Ruprecht, R. Krause, Explicit parallel-in-time integration of a linear acoustic-advection system, *Computers & Fluids* 59 (0) (2012) 72 – 83.  
URL <http://dx.doi.org/10.1016/j.compfluid.2012.02.015>
- [29] X. Dai, Y. Maday, Stable parareal in time method for first- and second-order hyperbolic systems, *SIAM Journal on Scientific Computing* 35 (1) (2013) A52–A78.  
URL <http://dx.doi.org/10.1137/110861002>
- [30] M. Gander, S. Güttel, Pararexp: A parallel integrator for linear initial-value problems, *SIAM Journal on Scientific Computing* 35 (2) (2013) C123–C142.  
URL <http://dx.doi.org/10.1137/110856137>
- [31] T. Haut, B. Wingate, An asymptotic parallel-in-time method for highly oscillatory PDEs, *SIAM Journal on Scientific Computing* In press.
- [32] J. M. F. Trindade, J. C. F. Pereira, Parallel-in-time simulation of two-dimensional, unsteady, incompressible laminar flows, *Numerical Heat Transfer, Part B: Fundamentals* 50 (1) (2006) 25–40.  
URL <http://dx.doi.org/10.1080/10407790500459379>
- [33] D. Ruprecht, R. Speck, M. Emmett, M. Bolten, 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.  
URL [http://sc13.supercomputing.org/sites/default/files/PostersArchive/tech\\_posters/post148s2-file3.pdf](http://sc13.supercomputing.org/sites/default/files/PostersArchive/tech_posters/post148s2-file3.pdf)
- [34] R. Speck, D. Ruprecht, R. Krause, M. Emmett, M. Minion, M. Winkel, 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*, IEEE Computer Society Press, Los Alamitos, CA, USA, 2012, pp. 92:1–92:11.  
URL <http://dx.doi.org/10.1109/SC.2012.6>
- [35] R. D. Haynes, B. W. Ong, MPI-OpenMP algorithms for the parallel space-time solution of time dependent PDEs, in: *Domain Decomposition Methods in Science and Engineering XXI*, Vol. 98 of *Lecture Notes in Computational Science and Engineering*, Springer International Publishing, 2014, pp. 179–187.  
URL [http://dx.doi.org/10.1007/978-3-319-05789-7\\_14](http://dx.doi.org/10.1007/978-3-319-05789-7_14)
- [36] R. Krause, D. Ruprecht, Hybrid space-time parallel solution of Burgers' equation, in: *Domain Decomposition Methods in Science and Engineering XXI*, Vol. 98 of *Lecture Notes in Computational Science and Engineering*, Springer International Publishing, 2014, pp. 647–655.  
URL [http://dx.doi.org/10.1007/978-3-319-05789-7\\_62](http://dx.doi.org/10.1007/978-3-319-05789-7_62)
- [37] A. T. Barker, A minimal communication approach to parallel time integration, *International Journal of Computer Mathematics* (2013) 1–15.  
URL <http://dx.doi.org/10.1080/00207160.2013.800193>
- [38] D. E. Keyes, Exaflop/s: The why and the how, *Comptes Rendus Mécanique* 339 (23) (2011) 70 – 77. doi:10.1016/j.crme.2010.11.002.
- [39] O. Fuhrer, C. Osuna, X. Lapillonne, T. Gysi, B. Cumming, M. Bianco, A. Arteaga, T. C. Schulthess, Towards a performance portable, architecture agnostic implementation strategy for weather and climate models, *Supercomputing frontiers and innovations I* (1).
- [40] K. Asanovic, R. Bodik, J. Demmel, T. Keaveny, K. Keutzer, J. Kubiatowicz, N. Morgan, D. Patterson, K. Sen, J. Wawrzynek, D. Wessel, K. Yelick, A view of the parallel computing landscape, *Commun. ACM* 52 (2009) 56–67.
- [41] L. Berry, W. Elwasif, J. Reynolds-Barredo, D. Samaddar, R. Sanchez, D. Newman, Event-based parareal: A data-flow based implementation of parareal, *Journal of Computational Physics* 231 (17) (2012) 5945 – 5954.  
URL <http://dx.doi.org/10.1016/j.jcp.2012.05.016>
- [42] P. F. Fischer, F. Hecht, Y. Maday, A parareal in time semi-implicit approximation of the Navier-Stokes equations, in: R. Kornhuber, et al. (Eds.), *Domain Decomposition Methods in Science and Engineering*, Vol. 40 of *Lecture Notes in Computational Science and Engineering*, Springer, Berlin, 2005, pp. 433–440.  
URL [http://dx.doi.org/10.1007/3-540-26825-1\\_44](http://dx.doi.org/10.1007/3-540-26825-1_44)
- [43] D. Ruprecht, Convergence of parareal with spatial coarsening, submitted to PAMM (2014).
- [44] T. Gysi, O. Fuhrer, C. Osuna, B. Cumming, T. Schulthess, Stella: A domain-specific embedded language for stencil codes on structured grids, in: *EGU General Assembly Conference Abstracts*, Vol. 16, 2014, p. 8464.
- [45] D. Ruprecht, R. Speck, R. Krause, Parareal for diffusion problems with space- and time-dependent coefficients, arXiv:1401.7829 [math.NA], submitted to *Lecture Notes in Computational Science and Engineering* (2014).  
URL <http://arxiv.org/abs/1401.7829>
- [46] R. Speck, D. Ruprecht, M. Emmett, M. Bolten, 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.  
URL <http://dx.doi.org/10.3233/978-1-61499-381-0-263>
- [47] J. J. Dongarra, H. W. Meuer, E. Strohmaier, et al., Top500 supercomputer sites, *Supercomputer* 13 (1997) 89–111.
- [48] W. Feng, K. Cameron, The green500 list: Encouraging sustainable supercomputing, *Computer* 40 (12) (2007) 50–55. doi:10.1109/MC.2007.445.
- [49] G. Fourestey, B. Cumming, L. Gilly, T. C. Schulthess, First experiences with validating and using the cray power management database tool, arXiv preprint arXiv:1408.2657.

- [50] B. Cumming, G. Fourestey, O. Fuhrer, T. Gysi, M. Fatica, T. C. Schulthess, Application centric energy-efficiency study of distributed multi-core and hybrid cpu-gpu systems, accepted SC 14.