

# Exa-Dune—Flexible PDE Solvers, Numerical Methods and Applications



**Peter Bastian, Mirco Altenbernd, Nils-Arne Dreier, Christian Engwer, Jorrit Fahlke, René Fritze, Markus Geveler, Dominik Göddeke, Oleg Iliev, Olaf Ippisch, Jan Mohring, Steffen Müthing, Mario Ohlberger, Dirk Ribbrock, Nikolay Shegunov, and Stefan Turek**

**Abstract** In the EXA-DUNE project we have developed, implemented and optimised numerical algorithms and software for the scalable solution of partial differential equations (PDEs) on future exascale systems exhibiting a heterogeneous massively parallel architecture. In order to cope with the increased probability of hardware failures, one aim of the project was to add flexible, application-oriented resilience capabilities into the framework. Continuous improvement of the underlying hardware-oriented numerical methods have included GPU-based sparse approximate inverses, matrix-free sum-factorisation for high-order discontinuous Galerkin discretisations as well as partially matrix-free preconditioners. On top of that, additional scalability is facilitated by exploiting massive coarse grained parallelism offered by multiscale and uncertainty quantification methods where we have focused on the adaptive choice of the coarse/fine scale and the overlap region as well as the combination of local reduced basis multiscale methods and the multilevel Monte-Carlo algorithm. Finally, some of the concepts are applied in a land-surface model including subsurface flow and surface runoff.

---

P. Bastian (✉) · S. Müthing  
Heidelberg University, Heidelberg, Germany  
e-mail: [peter.bastian@iwr.uni-heidelberg.de](mailto:peter.bastian@iwr.uni-heidelberg.de)

N.-A. Dreier · C. Engwer · J. Fahlke · R. Fritze · M. Ohlberger  
University of Münster, Münster, Germany

M. Altenbernd · D. Göddeke  
University of Stuttgart, Stuttgart, Germany

M. Geveler · D. Ribbrock · S. Turek  
TU Dortmund, Dortmund, Germany

O. Iliev · J. Mohring · N. Shegunov  
Fraunhofer ITWM, Kaiserslautern, Germany

O. Ippisch  
TU Clausthal-Zellerfeld, Clausthal-Zellerfeld, Germany

## 1 Introduction

In the EXA-DUNE project we extend the Distributed and Unified Numerics Environment (DUNE)<sup>1</sup> [6, 7] by hardware-oriented numerical methods and hardware-aware implementation techniques developed in the (now) FEAT3<sup>2</sup> [55] project to provide an exascale-ready software framework for the numerical solution of a large variety of partial differential equation (PDE) systems with state-of-the-art numerical methods including higher-order discretisation schemes, multi-level iterative solvers, unstructured and locally-refined meshes, multiscale methods and uncertainty quantification, while achieving close-to-peak performance and exploiting the underlying hardware.

In the first funding period we concentrated on the node-level performance as the framework and in particular its algebraic multigrid solver already show very good scalability in MPI-only mode as documented by the inclusion of DUNE's solver library in the High-Q-Club, the codes scaling to the full machine in Jülich at the time, with close to half a million cores. Improving the node-level performance in light of future exascale hardware involved multithreading ("MPI+X") and in particular exploiting SIMD parallelism (vector extensions of modern CPUs and accelerator architectures). These aspects were addressed within the finite element assembly and iterative solution phases. Matrix-free methods evaluate the discrete operator without storing a matrix, as the name implies, and promise to be able to achieve a substantial fraction of peak performance. Matrix-based approaches on the other hand are limited by memory bandwidth (at least) in the solution phase and thus typically exhibit only a small fraction of the peak (GFLOP/s) performance of a node, but decades of research have led to robust and efficient (in terms of number of iterations) iterative linear solvers for practically relevant systems. Importantly, a consideration of matrix-free and matrix-based methods needs to take the order of the method into account. For low-order methods it is imperative that a matrix entry can be recomputed in less time than it takes to read it from memory, to counteract the memory wall problem. This requires to exploit the problem structure as much as possible, i.e., to rely on constant coefficients, (locally) regular mesh structure and linear element transformations [28, 37]. In these cases it is even possible to apply stencil type techniques, like developed in the EXA-STENCIL project [40]. On the other hand, for high-order methods with tensor-product structure the complexity of matrix-free operator evaluation can be much less than that of matrix-vector multiplication, meaning that less floating-point operations have to be performed which at the same time can be executed at a higher rate due to reduced memory pressure and better suitability for vectorization [12, 39, 50]. This makes high-order methods extremely attractive for exascale machines [48, 51].

---

<sup>1</sup><http://www.dune-project.org/>.

<sup>2</sup><http://feast.tu-dortmund.de/>.

In the second funding phase we have mostly concentrated on the following aspects:

1. **Asynchronicity and fault tolerance:** High-level C++ abstractions form the basis of transparent error handling using exceptions in a parallel environment, fault-tolerant multigrid solvers as well as communication hiding Krylov methods.
2. **Hardware-aware solvers for PDEs:** We investigated matrix-based sparse-approximate inverse preconditioners including novel machine-learning approaches, vectorization through multiple right-hand sides as well as matrix-free high-order Discontinuous Galerkin (DG) methods and partially matrix-free robust preconditioners based on algebraic multigrid (AMG).
3. **Multiscale (MS) and uncertainty quantification (UQ) methods:** These methods provide an additional layer of embarrassingly parallel tasks on top of the efficiently parallelized forward solvers. A challenge here is load balancing of the asynchronous tasks which has been investigated in the context of the localized reduced basis multiscale method and multilevel Monte Carlo methods.
4. **Applications:** We have considered large-scale water transport in the subsurface coupled to surface flow as an application where the discretization and solver components can be applied.

In the community, there is broad consensus on the assumptions about exascale systems that did not change much during the course of this 6 year project. A report by the Exascale Mathematics Working Group to the U.S. Department of Energy's Advanced Scientific Computing Research Program [16] summarises these challenges as follows, in line with [35] and more recently the Exascale Computing Project.<sup>3</sup> (1) The anticipated power envelope of 20 MW implies strong limitations on the amount and organisation of the hardware components, an even stronger necessity to fully exploit them, and eventually even power-awareness in algorithms and software. (2) The main performance difference from peta- to exascale will be through a 100–1000 fold increase in parallelism at the node level, leading to extreme levels of concurrency and increasing heterogeneity through specialised accelerator cores and wide vector instructions. (3) The amount of memory per 'core' and the memory and interconnect bandwidth/latency will only increase at a much smaller rate, hence increasing the demand for lower memory footprints and higher data locality. (4) Finally, hardware failures, and thus the mean-time-between-failure (MTBF), were expected to increase proportionally (or worse) corresponding to the increasing number of components. Recent studies have indeed confirmed this expectation [30], although not at the projected rate. First exascale systems are scheduled for 2020 in China [42], 2021 in the US and 2023 [25] in Europe. Although the details are not yet fully disclosed, it seems that the number of nodes will not be larger than  $10^5$  and will thus remain in the range of previous machines such as the

---

<sup>3</sup><https://www.exascaleproject.org/>.

BlueGene. The major challenge will thus be to exploit the node level performance of more than 10 TFLOP/s.

The rest of this paper is organized as follows. In Sect. 2 we lay the foundations of asynchronicity and resilience, while Sect. 3 discusses several aspects of hardware-aware and scalable iterative linear solvers. These building blocks will then be used in Sects. 4 and 5 to drive localized reduced basis and multilevel Monte-Carlo methods. Finally, Sect. 6 covers our surface-subsurface flow application.

## 2 Asynchronicity and Fault Tolerance

As predicted in the first funding period, latency has indeed become a major issue, both within a single node as well as between different MPI ranks. The core concept underlying all latency- and communication-hiding techniques is asynchronicity. This is also crucial to efficiently implement certain local-failure local-recovery methods. Following the DUNE philosophy, we have designed a generic layer that abstracts the use of asynchronicity in MPI from the user. In the following, we first describe this layer and its implementation, followed by representative examples on how to build middleware infrastructure on it, and on its use for s-step Krylov methods and fault tolerance beyond global checkpoint-restart techniques.

### 2.1 Abstract Layer for Asynchronicity

We first introduce a general abstraction for asynchronicity in parallel MPI applications, which we developed for DUNE. While we integrated these abstractions with the DUNE framework, most of the code can easily be imported into other applications, and is available as a standalone library.

The C++ API for MPI was dropped from MPI-3 since it offered no real advantage over the C bindings, beyond being a simple wrapper layer. Most MPI users coding in C++ are still using the C bindings, writing their own C++ interface/layer, in particular in more generic software frameworks. At the same time the C++11 standard introduced high-level concurrency concepts, in particular the future/promise construct to enable an asynchronous program flow while maintaining value semantics. We adopt this approach as a first principle in our MPI layer to handle asynchronous MPI operations and propose a high-level C++ MPI interface, which we provide in DUNE under the generic interface of `Dune::Communication` and a specific implementation `Dune::MPICommunication`.

An additional issue of the concrete MPI library in conjunction with C++ is the error handling concept. In C++, exceptions are the advocated approach to handle error propagation. As exceptions change the local code path on the, e.g., failing process in a hard fault scenario, exceptions can easily lead to a deadlock. As we

discuss later, the introduction of our asynchronous abstraction layer enables global error handling in an exception friendly manner.

In concurrent environments a C++ future decouples values from the actual computation (promise). The program flow can continue while a thread is computing the actual result and promotes this via promise to the future. The MPI C and Fortran interfaces offer asynchronous operations, but in contrast to thread parallel, the user does not specify the operation within the concurrent operation. Actually, MPI on its own does not offer any real concurrency at all, and provides instead a handle-based programming interface to avoid certain cases of deadlocks: the control flow is allowed to continue without finishing the communication, while the communication usually only proceeds when calls into the MPI library are executed.

We developed a C++ layer on top of the asynchronous MPI operations, which follows the design of the C++11 future. Note that the actual `std::future` class cannot be used for this purpose.

```
template<typename T>
class Future{
    void wait();
    bool ready() const;
    bool valid() const;
    T get();
};
```

As different implementations like thread-based `std::future`, task-based `TBB::future`, and our new `MPIFuture` are available, usability greatly benefits from a dynamically typed interface. This is a reasonable approach, as `std::future` is using a dynamical interface already and also the MPI operations are coarse grained, so that the additional overhead of virtual function calls is negligible. At the same time the user expects a future to offer value semantics, which contradicts the usual pointer semantics used for dynamic polymorphism. In EXA-DUNE we decided to implement type-erasure to offer a clean and still flexible user interface. An `MPIFuture` is responsible for handling all states associated with an MPI operation.

```
class MPIFuture{
private:
    mutable MPI_Request req_;
    mutable MPI_Status status_;
    impl::Buffer<R> data_;
    impl::Buffer<S> send_data_;
public:
    ...
};
```

The future holds a mutable `MPI_Request` and `MPI_Status` to access information on the current operation and it holds buffer objects, which manage the actual data. These buffers offer a great additional value, as we do not access the raw data directly, but can include data transformation and varying ownership. For example

it is now possible to directly send an `std::vector<double>`, where the receiver automatically resizes the `std::vector` according to the incoming data stream.

This abstraction layer enables different use cases, highlighted below:

1. **Parallel C++ exception handling:** Exceptions are the recommended way to handle faults in C++ programs. As exceptions alter the execution path of a single node, they are not suitable for parallel programs. As asynchronicity allows for moderately diverging execution paths, we can use it to implement parallel error propagation using exceptions.
2. **Solvers and preconditioners tolerant to hard and soft faults:** This functionality is used for failure propagation, restoration of MPI in case of a hard fault, and asynchronous in-memory checkpointing.
3. **Asynchronous Krylov solvers:** Scalar products in Krylov methods require global communication. Asynchronicity can be used to hide the latency and improve strong scalability.
4. **Asynchronous parallel IO:** The layer allows to transform any non-blocking MPI operation into a really asynchronous operation. This allows also to support asynchronous IO, to hide the latency of write operations and overlap with the computation of the next iteration or time step.
5. **Parallel localized reduced basis methods:** Asynchronicity will be used to mitigate the load-imbalance inherent in the error estimator guided adaptive online enrichment of local reduced bases.

## 2.2 *Parallel C++ Exception Handling*

In parallel numerical algorithms, unexpected behaviour can occur quite frequently: a solver could diverge, the input of a component (e.g., the mesher) could be inappropriate for another component (e.g., the discretiser), etc. A well-written code should detect unexpected behaviour and provide users with a possibility to react appropriately in their own programs, instead of simply terminating with some error code. For C++, exceptions are the recommended method to handle this. With well placed exceptions and corresponding try-catch blocks, it is possible to accomplish a more robust program behaviour. However, the current MPI specification [44] does not define any way to propagate exceptions from one rank (process) to another. In the case of unexpected behaviour within the MPI layer itself, MPI programs simply terminate, maybe after a time-out. This is a design decision that unfortunately implies a severe disadvantage in C++, when combined with the ideally asynchronous progress of computation and communication: an exception that is thrown locally by some rank can currently lead to a communication deadlock, or ultimately even to undesired program termination. Even though exceptions are technically an illegal use of the MPI standard (a peer no longer participates in a communication), it undesirably conflicts with the C++ concept of error handling.

Building on top of the asynchronicity layer, we have developed an approach to enable parallel C++ exceptions. We follow C++11 techniques, e.g., use future-like abstractions to handle asynchronous communication. Our currently implemented

interface requires ULFM [11], an MPI extension to restore communicators after rank losses, which is scheduled for inclusion into MPI-4. We also provide a fallback solution for non-ULFM MPI installations, that employs an additional communicator for propagation and can, by construction, not handle hard faults, i.e., the loss of a node resulting in the loss of rank(s) in some communicator.

To detect exceptions in the code we have extended the `Dune::MPIGuard`, that previously only implemented the scope guard concept to detect and react on local exceptions. Our extension revokes the MPI communicator using the ULFM functionality if an exception is detected, so that it is now possible to use communication inside a block with scope guard. This makes it superfluous to call the `finalize` and `reactivate` methods of the `MPIGuard` before and after each communication.

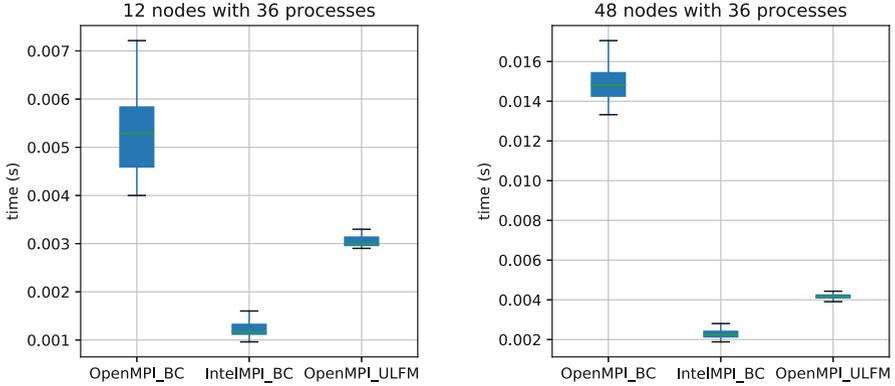
```
try{
    MPIGuard guard(comm);
    do_something();
    communicate(comm);
}catch(...){
    comm.shrink();
    recover(comm);
}
```

#### Listing 1 MPIGuard

Listing 1 shows an example how to use the `MPIGuard` and recover the communicator in a node loss scenario. In this example, an exception that is thrown only on a few ranks in `do_something()` will not lead to a deadlock, since the `MPIGuard` would revoke the communicator. Details of the implementation and further descriptions are available in a previous publication [18]. We provide the “black-channel” fallback implementation as a standalone version.<sup>4</sup> This library uses the P-interface of the MPI standard, which makes it possible to redefine MPI functions. At the initialization of the MPI setting the library creates an opaque communicator, called `blackchannel`, on which a pending `MPI_Irecv` request is waiting. Once a communicator is revoked, the revoking rank sends messages to the pending `blackchannel` request. To avoid deadlocks, we use `MPI_Waitany` to wait for a request, which listens also for the `blackchannel` request. All blocking communication is redirected to non-blocking calls using the P-interface. The library is linked via `LD_PRELOAD` which makes it usable without recompilation and could be removed easily once a proper ULFM implementation is available in MPI.

Figure 1 shows a benchmark comparing the time which is used for duplicating a communicator, revoking it and restore a valid state. The benchmark was performed on PALMA2, the HPC cluster of the University of Muenster. Three implementations are compared; `OpenMPI_BC` and `IntelMPI_BC` are using the blackchannel library based on OpenMPI and IntelMPI, respectively. `OpenMPI_ULFM` uses the ULFM

<sup>4</sup><https://gitlab.dune-project.org/exadune/blackchannel-ulfm>.



**Fig. 1** Benchmark of different MPI implementations: 12 nodes with 36 processes (left), 48 nodes with 36 processes (right), cf. [18]

implementation provided by [fault-tolerance.org](http://fault-tolerance.org), which is based on OpenMPI. We performed 100 measurements for each implementation. The blackchannel implementation is competitive to the ULFM implementation. As OpenMPI is in this configuration not optimized and does not use the RDMA capabilities of the interconnect, it is slower than the IntelMPI implementation. The speed up of the *OpenMPI\_ULFM* version compared to the *OpenMPI\_BC* version is due to the better communication strategy.

### 2.3 Compressed in-Memory Checkpointing for Linear Solvers

The previously described parallel exception propagation, rank loss detection and communicator restoration by using the ULFM extension, allow us to implement a flexible in-memory checkpointing technique which has the potential to recover from hard faults on-the-fly without any user interaction. Our implementation establishes a backup and recovery strategy which in part is based on a local-failure local-recovery (LFLR) [54] approach, and involves lossy compression techniques to reduce the memory footprint as well as bandwidth pressure. The contents of this subsection have not been published previously.

**Modified Solver Interface** To enable the use of exception propagation as illustrated in the previous section and to implement different backup recovery approaches we kept all necessary modifications to DUNE-ISTL, the linear solver library. We embed the solver initialisation and the iterative loop in a try-catch block, and provide additional entry and execution points for recovery and backup, see Listing 2 for details. Default settings are provided on the user level, i.e., DUNE-PDELAB.

```

1  init_variables();
2  done = false;
3  while (!done) try {
4      MPIGuard guard(comm);
5      if (this->processRecovery(...))
6          reinit_execution();
7      } else {
8          init_execution();
9      }
10     for (i=0 ; i<=maxit; i++ ) {
11         do_iteration();
12         if (converged) {
13             done = true;
14             break;
15         }
16         this->processBackup(...);
17     }
18 } catch(Exception & e) {
19     done = false;
20     comm.reconstitute();
21     if (!this->processOnException(...))
22         throw;
23 }

```

**Listing 2** Solver modifications

This implementation ensures that the iterative solving process is active until the convergence criterion is reached. An exception inside the try-block on any rank is detected by the `MPIGuard` and propagated to all other ranks, so that all ranks will jump to the catch-block.

This catch-block can be specialised for different kind of exceptions, e.g., if a solver has diverged and a corresponding exception is thrown it could define some specific routine to define a modified restart with a possibly more robust setting and/or initial guess. The catch-block in Listing 2 exemplarily shows a possible solution in the scenario of a communicator failure, e.g., a node loss which is detected by using the ULFM extension to MPI, encapsulated by our wrapper for MPI exceptions. Following the detection and propagation, all still valid ranks end up in the catch-block and the communicator must be re-established in some way (Listing 2, line 20). This can be done by shrinking the communicator or replacing lost nodes by some previously allocated spare ones. After the communicator reconstitution a user-provided stack of functions can be executed (Listing 2, line 21) to react on the exception. If there is no on-exception-function or neither of them returns true the exception is re-thrown to the next higher level, e.g., from the linear solver to the application level, or in case of nested solvers, e.g. in optimisation or uncertainty quantification.

Furthermore, there are two additional entry points for user provided function stacks: In line 5 of Listing 2 a stack of recovery functions is executed and if it returns true, the solver expects that some modification, i.e., recovery, has been done.

In this case it could be necessary that the other participating ranks have to update some data, like resetting their local right hand side to the initial values. The backup function stack in line 16 allows the user to provide functions for backup creation etc., after an iteration finished successfully.

**Recovery Approaches** First, regardless of these solver modifications, we describe the recovery concepts which are implemented into an exemplary recovery interface class providing functions that can be passed to the entry points within the modified solver. The interoperability of these components and the available backup techniques are described later. Our recovery class supports three different methods to recover from a data loss. The first approach is a global rollback to a backup, potentially involving lossy compression: progress on non-faulty ranks may be lost but the restored data originate from the same state, i.e., iteration. This means there is no asynchronous progression in the recovered iterative process but possibly just an error introduced through the used backup technique, e.g., through lossy compression. This compression error can reduce the quality of the recovery and lead to additional iterations of the solver, but is still superior to a restart, as seen later. For the second and third approaches, we follow the local-failure local-recovery strategy and re-initialize the data which are lost on the faulty rank by using a backup. The second, slightly simpler strategy uses these data to continue with solver iterations. The third method additionally smoothes out the probably deteriorated (because of compression) data by solving a local auxiliary problem [29, 31]. This problem is set up by restricting the global operator to its purely local degrees of freedom with indices  $\mathcal{F} \subset \mathbb{N}$  and a Dirichlet boundary layer. The boundary layer can be obtained by extending  $\mathcal{F}$  to some set  $\mathcal{J}$  using the ghost layer, or possibly the connectivity pattern of the operator  $\mathbf{A}$ . The Dirichlet values on the boundary layer are set to their corresponding values  $x_N$  on the neighbouring ranks and thus additional communication is necessary:

$$\begin{aligned} \mathbf{A}(\mathcal{F}, \mathcal{F})\tilde{x}(\mathcal{F}) &= b(\mathcal{F}) && \text{in } \mathcal{F} \\ \tilde{x} &= x_N && \text{on } \mathcal{J} \setminus \mathcal{F} \end{aligned}$$

If this problem is solved iteratively and backup data are available, the computation speed can be improved by initializing  $\tilde{x}$  with the data from the backup.

**Backup Techniques** Our current implementation provides two different techniques for compressed backups as well as a basic class which allows ‘zero’-recovery (zeroing of lost data) if the user wants to use the auxiliary solver in case of data loss without storing any additional data during the iterative procedure.

The next backup class uses a multigrid hierarchy for lossy data compression. Thus it should only be used if a multigrid operator is already in use within the solving process because otherwise the hierarchy has to be built beforehand and introduces additional overhead. Compressing the iterative vector with the multigrid hierarchy currently involves a global communication. In addition there is no adaptive control of the compression depth (i.e., hierarchy level where the

backup is stored), but it has to be specified by the user, see a previous publication for details [29].

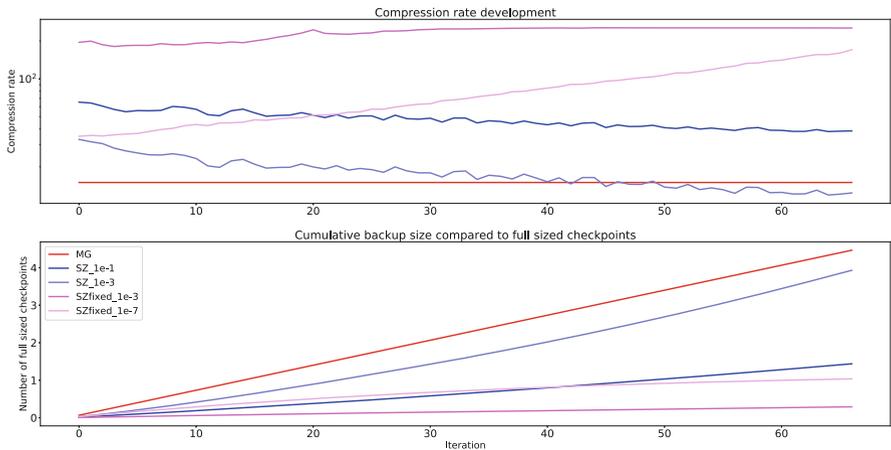
We also implemented a compressed backup technique based on SZ compression [41]. SZ allows compression to a specified accuracy target and can yield better compression rates than multigrid compression. The compression itself is purely local and does not involve any additional communication. We provide an SZ backup with a fixed user-specified compression target as well as a fully adaptive one which couples the compression target to the residual norm within the iterative solver. For the first we achieve an increased rate while we approach the approximate solution, as seen in Fig. 2 (top, pink lines), at the price of an increased overhead in case of a data loss (cf. Fig. 3). The backup with adaptive compression target (blue lines) gives more constant compression rates, and a better recovery in case of faults in particular in the second half of the iterative procedure of the solver.

The increased compression rate for the fixed SZ backup is obtained because, during the iterative process, the solution gets more smooth and thus can be compressed better by the algorithm. For the adaptive method this gain is counteracted by the demand of a higher compression accuracy.

All backup techniques require to communicate a data volume smaller than the volume of four full checkpoints, see Fig. 2 (bottom). Furthermore this bandwidth requirement is distributed over all 68 iterations (in the fault-free scenario) and could be decreased further by a lower checkpoint frequency.

The chosen backup technique is initiated before the recovery class and passed to it. Further backup techniques can be implemented by using the provided base class and overloading the virtual functions.

**Bringing the Approaches Together** The recovery class provides three functions which are added to the function stacks within the modified solver interface. The



**Fig. 2** Compression rate in the iterative solution for an anisotropic Poisson problem on 52 cores with approximately 480 K DOF per core

backup routine is added to the stack of backup functions of the specified iterative solver and generates backups of the current iterative solution by using the provided backup class.

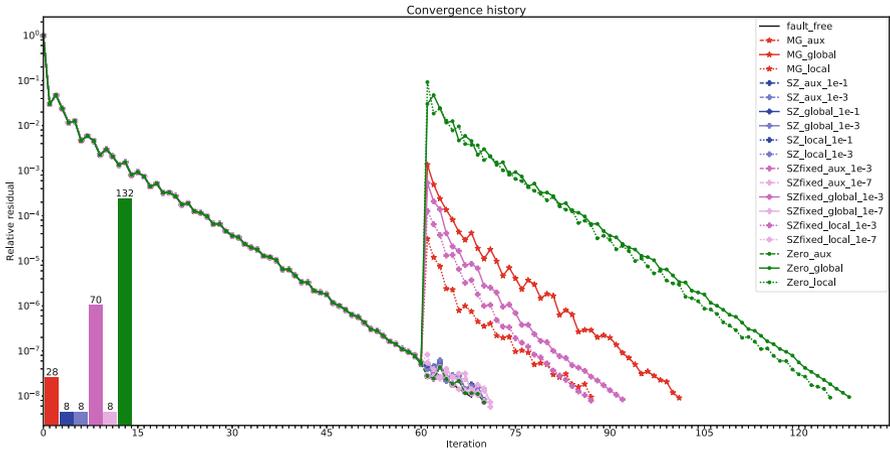
```
SomeSolver solver;
SomeBackup backup;
Recovery recovery(backup);
solver.addBackupFunction(&Recovery::backup, &recovery);
```

To adapt numerical as well as communication overhead for different fault scenarios and machine characteristics, the backup creation frequency can be varied. After the creation of the backup it is sent to a remote rank where it is kept in memory but never written to disk. In the following this is called ‘remote backup’. Currently the backup propagation happens circular by rank. It is also possible to trigger writing a backup to disk.

In the near future we will implement an on-the-fly recovery if an exception is thrown. These will be provided to the other two function stacks and will differ depending on the availability of the ULFM extensions: if the extension is not available we can only detect and propagate exceptions but not recover a communicator in case of hard faults, i.e., node losses (cf. Sect. 2.2). In this scenario the function provided to the on-exception stack will only write out the global state. Fault-free nodes will write the data of the current iterative vector, whereas for faulty nodes the corresponding remote backup is written. In the following the user will be able to provide a flag to the executable which modifies the backup object initiation to read in the stored checkpoint data. Afterwards the recovery function of our interface will overwrite the initial values of the solver with the checkpointed and possibly smoothed data like described above. If the ULFM extensions are available, the recovery can be realised without any user interaction: during the backup class initiation a global communication ensures that it is the first and therefore fault-free start of the parallel execution. If the process is a respawned one which replaces a lost rank, this communication is matched by a send communication created from the rank which holds the corresponding remote backup. This communication will be initiated by the on-exception function. In addition to this message the remote backup rank sends the stored compressed backup so that the respawned rank can use this backup to recover the lost data.

So far, we have not fully implemented rebuilding the solver and preconditioner hierarchy, and the re-assembly of the local systems, in case of a node loss. This can be done with, e.g., message logging [13], or similar techniques which allow recomputing the individual data on the respawned rank without additional communication.

Figure 3 shows the effect of various combinations of different backup and recovery techniques in case of a data loss on one rank after iteration 60. The problem is an anisotropic Poisson problem with zero Dirichlet boundary conditions which reaches the convergence criterion after 68 iterations in a fault-free scenario (black



**Fig. 3** Convergence history in case of data loss and recovery on one rank, same setting as in Fig. 2. Bottom left: number of iterations to solve the auxiliary problem when using the backups as initial guess. Note that the groups of the same colour are important, not the individual graphs

line). It is executed in parallel on 52 ranks with approximately 480,000 degrees of freedom per rank. Thus one rank loss corresponds to a loss of around 2% of data. For solving a conjugate gradient solver with an algebraic multigrid preconditioner is applied. In addition to the residual norm we show the number of iterations which are needed to solve the auxiliary problem when using different backups as initial guess at the bottom left.

The different backup techniques are colour-coded (multigrid: red; adaptive SZ compression: blue; fixed SZ compression: pink; no backup: green). For the SZ techniques we consider two cases, each with a different compression accuracy (fixed compression), respectively a different additional scaling coefficient (SZ). Recovery techniques are coded with different line styles: global roll-back recovery is indicated by straight lines; simple local recovery is shown with dotted lines and if an auxiliary problem is solved to improve the quality of the recovery it is drawn with a dashed line style. We observe that a zero recovery, multigrid compression and a fixed SZ backup with a low accuracy target are not competitive if no auxiliary problem is solved. The number of iterations needed until convergence then increases significantly. By applying an auxiliary solver the convergence can be almost fully restored (one additional global iteration) but the auxiliary solver needs a high amount of iterations (multigrid: 28; sz: 70; no backup: 132). Other backup techniques only need 8 auxiliary solver iterations. When using adaptive or very accurate fixed SZ compression the convergence behaviour can be nearly preserved even when only a local recovery or a global roll-back is applied. The adaptive compression technique has similar data overhead as the fixed SZ compression

(cf. Fig. 2, bottom) but gives slightly better results: both adaptive SZ compression approaches introduce only one additional iteration for all recovery approaches. For the accurate fixed SZ compression (SZfixed\*\_1e-7) we have two additional iterations when using local or global recovery but if we apply the auxiliary solver we also have only one additional iteration until convergence.

## 2.4 Communication Aware Krylov Solvers

In Krylov methods multiple scalar products per iteration must be computed. This involves global sums in a parallel setting. As a first improvement we merged the evaluation of the convergence criterion to the computation of a scalar product. Obviously this does not effect the computed values, but the iteration terminates one iteration later. However this reduces the number of global reductions per iteration from 3 to 2 and thus already saves communication overhead.

As a second step we modify the algorithm, such that only one global communication is performed per iteration. This algorithm can also be found in the paper of Chronopoulos and Gear [15]. Another optimization is to overlap the two scalar products with the application of the operator and preconditioner, respectively. This algorithm was first proposed by Gropp [27]. A fully elaborate version was then presented by Ghysels and Vanroose [27]. This version only needs one global reduction per iteration, which is overlapped with both the application of the preconditioner and operator. This algorithm is shown in Algorithm 2.

---

### Algorithm 1 PCG

---

$$r_0 = b - Ax_0$$

$$p_1 = Mr_0$$

$$\rho_1 = \langle p_1, r_0 \rangle$$

**for**  $i = 1, \dots$  **do**

$$q_i = Ap_i$$

$$\alpha_i = \langle p_i, q_i \rangle$$

$$x_i = x_{i-1} + \frac{\rho_i}{\alpha_i} p_i$$

$$r_i = r_{i-1} - \frac{\rho_i}{\alpha_i} q_i$$

$$z_{i+1} = Mr_i$$

**break** if  $\|r_i\| < \varepsilon$

$$\rho_{i+1} = \langle z_{i+1}, r_i \rangle$$

$$p_{i+1} = \frac{\rho_{i+1}}{\rho_i} p_i + z_{i+1}$$


---

---

**Algorithm 2** Pipelined CG

---

```

r0 = b - Ax0
p1 = Mr0
q1 = Ap1
ρ1 = ⟨p1, r0⟩
α1 = ⟨p1, q1⟩
s1 = Mq1
t1 = As1
for i = 1, ... do
    xi = xi-1 +  $\frac{\rho_i}{\alpha_i}$  pi
    ri = ri-1 -  $\frac{\rho_i}{\alpha_i}$  qi
    break if ||ri|| < ε
    zi+1 = zi -  $\frac{\rho_i}{\alpha_i}$  si
    wi+1 = wi -  $\frac{\rho_i}{\alpha_i}$  ti
    ρi+1 = ⟨zi+1, ri⟩
    α̃i+1 = ⟨zi+1, wi+1⟩
    αi+1 =  $\frac{\alpha_i \rho_{i+1}^2}{\rho_i^2} + \tilde{\alpha}_{i+1}$ 
    vi+1 = Mwi+1
    ui+1 = Avi+1
    si+1 =  $\frac{\rho_{i+1}}{\rho_i}$  si + vi+1
    ti+1 =  $\frac{\rho_{i+1}}{\rho_i}$  ti + ui+1
    pi+1 =  $\frac{\rho_{i+1}}{\rho_i}$  pi + zi+1
    qi+1 =  $\frac{\rho_{i+1}}{\rho_i}$  qi + wi+1

```

---

With the new communication interface, described above, we are able to compute multiple sums in one reduction pattern and overlap the communication with computation. To apply these improvements in Krylov solvers the algorithm must be adapted, such that the communication is independent of the overlapping computation. For this adaption we extend the `ScalarProduct` interface by a function which can be passed multiple pairs of vectors for which the scalar product should be computed. The function returns a `Future` which contains a `std::vector<field_type>`, once it has finished.

```

Future<vector<field_type>>
dots(initializer_list<tuple<X&, X&>> pairs);

```

The function can be used in the Krylov methods like this:

```

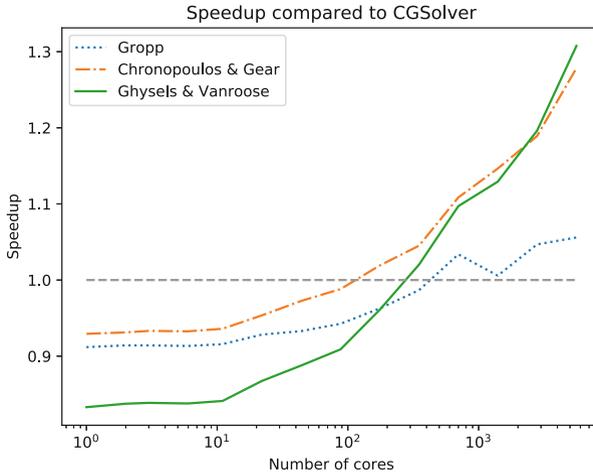
scalarproduct_future = sp.dot_norm({{p,q}, {z, b}, {b,b}});
// compute while communicate
auto result = scalarproduct_future.get();
field_type p_dot_q = result[0];
field_type z_dot_b = result[1];
field_type norm_b = std::sqrt(result[2]);

```

The runtime improvement of the algorithm strongly depends on the problem size and on the hardware. On large systems the communication overhead makes up a

**Table 1** Memory requirement, computational effort and global reductions per iteration for different versions of the preconditioned conjugate gradients method

	Required memory	Additional computational effort	Global reductions
PCG	$4N$	–	2
Chronopoulos and Gear	$6N$	$1N$	1
Gropp	$6N$	$2N$	2 overlapped
Ghysels and Vanroose	$10N$	$5N$	1 overlapped

**Fig. 4** Strong scaling for (pipelined) Krylov subspace methods

large part of the runtime. However, the maximum speedup is 3 for reducing the number of global reductions and 2 for overlapping communication and computation, compared to the standard version, so that a maximum speedup of 6 is possible. The optimization also increases the memory requirements and vector operations per iteration. An overview of runtime and memory requirements of the methods can be found in Table 1.

Figure 4 shows strong scaling for different methods. The shown speedup is per iteration and with respect to the `Dune::CGSolver`, which is the current CG implementation in DUNE. We use an SSOR preconditioner in an additive overlapping Schwarz setup. The problem matrix is generated from a 5-star Finite Difference model problem. With less cores the current implementation is faster than our optimized one. But with higher core count our optimized version outperforms it. The test was executed on the helics3 cluster of the University on Heidelberg, with 5600 cores on 350 nodes. We expect that on larger systems the speedup will further increase, since the communication is more expensive. The overlap of communication and computation does not really come into play, since the currently used MPI version does not support it completely.

### 3 Hardware-Aware, Robust and Scalable Linear Solvers

In this section we highlight improved concepts for high-performance iterative solvers. We provide matrix-based robust solvers on GPUs using sparse approximate inverses and optimize algorithm parameters using machine learning. On CPUs we significantly improve the node-level performance by using optimal matrix-free operators for Discontinuous Galerkin methods, specialized partially matrix-free preconditioners as well as vectorized linear solvers.

#### 3.1 *Strong Smoothers on the GPU: Fast Approximate Inverses with Conventional and Machine Learning Approaches*

In continuation of the first project phase, we enhanced the assembly of sparse approximate inverses (SPAI), a kind of preconditioner that we had shown to be very effective within the DUNE solver before [9, 26]. Concerning the assembly of such matrices we have investigated three strategies regarding their numerical efficacy (that is their quality in approximating  $A^{-1}$ ), the computational complexity of the actual assembly and ultimately, the total efficiency of the amortised assembly combined with all applications during a system solution. For both strategies, this includes a decisive performance engineering for different hardware architectures with focus on the exploitation of GPUs.

**SPAI-1** As a starting point we have developed, implemented and tuned a fast SPAI-1 assembly routine based on MKL/LAPACK routines (CPU) and on the cuBlas/cuSparse libraries, performing up to four times faster on the GPU. This implementation is based on the batched solution of QR decompositions that arise in Householder transformations during the SPAI minimisation process. In many cases, we observe that the resulting preconditioner features a high quality comparable to Gauss–Seidel methods. Most importantly, this result still holds true when taking into account the total time-to-solution, which includes the assembly time of the SPAI, even on a single core where the advantages of SPAI preconditioning over forward/backward substitution during the iterative solution process are not yet exploited. More systematic experiments with respect to these statements as well as their extension to larger test architectures are currently being conducted.

**SAINV** This preconditioner creates an approximation of the factorised inverse  $A^{-1} = ZDR$  of a matrix  $A \in \mathbb{R}^{N \times N}$  with  $D$  being a diagonal,  $Z$  an upper triangular and  $R$  a lower triangular Matrix.

To describe our new GPU implementation, we write the row-wise updates in the right-looking, outer product form of the  $A$ -biconjugation-process of the SAINV factorisation as follows: The assembly of the preconditioner is based on a loop over the existing rows  $i \in \{1, \dots, N\}$  of  $Z$  (initialised as unit matrix  $I_N$ ), where in every iteration the loop generally calls three operations, namely a sparse-matrix vector

---

**Algorithm 3** Algorithm of the row-wise updates
 

---

```

for ( $j = i + 1, \dots, N$ ) do
  if  $D_{jj} \neq 0$  then ▷ check if the fraction is unequal to zero
     $\alpha \leftarrow -\frac{D_{jj}}{D_{ii}} z_i^{(i-1)}$ 
    for  $n = 1, \dots, \text{nnz}(\alpha)$  do
      if  $\alpha_n > \varepsilon * \max_{i,j}(A_{ij})$  then ▷ here  $\alpha_n$  is the n-th entry of the vector  $\alpha$ 
        if  $\text{check}(z_j^n, z_j^n)$  then ▷ Has  $z_j$  already an entry at the columnindex of the n-th entry of  $\alpha$  ?
           $\text{add}(z_j^n, \alpha_n)$ 
           $\text{update\_minimum}(z_j)$  ▷ get new min. value of j-th row
        else if  $\text{nnz}(z_j) < \omega \times \frac{\text{nnz}(A)}{\text{dim}(A)}$  then ▷ maximum number of rowentries already reached?
           $\text{insert}(z_j, \alpha_n)$  ▷ insert the value  $\alpha_n$  at the fitting position
           $\text{update\_minimum}(z_j)$ 
        else if  $\alpha_n > \text{min}(z_j)$  then ▷ check if the value of  $\alpha_n$  is bigger than the minimum of  $z_j$ 
           $\text{replace}(\text{min}(z_j), \alpha_n)$  ▷ replace the old minimum with the value of  $\alpha_n$ 
           $\text{update\_minimum}(z_j)$ 
  
```

---

multiplication, a dot product and an update of the remaining rows  $i + 1, \dots, N$  based on a drop-parameter  $\varepsilon$ .

In our implementation we use the ELLPACK and CSR formats, pre-allocating a fixed amount of nonzeros of the matrix  $Z$  using  $\omega$  times the average number of nonzeros per row of  $A$ . Having a fixed row size, no reallocation of the arrays of the matrix format is needed and the row-wise update can be computed in parallel. This idea is based on the observation that while the density  $\omega$  for typical drop tolerances is not strictly limited, it generally falls into the interval ]0, 3[. As the SpMV and the dot kernels are well established, we take a closer look at the row-wise update, which is described more detailed in Algorithm 3. We first compute the values to be added and store them in a variable  $\alpha$ . Then we iterate over all nonzero entries of  $\alpha$  (which of course has the same sparsity pattern as  $z_i$ ) and check if the computed value exceeds a certain drop-tolerance. If this condition is met, we have three conditions for an insertion into the matrix  $Z$ :

1. Check if there is already an existing nonzero value in the  $j$ -th row at the column index of the value  $\alpha_n$  and search for the new minimal entry of this row.
2. Else check if there is still place in the  $j$ -th row, so we can simply insert the value  $\alpha_n$  into that row and search for the new minimal entry of this row.
3. Else check if the value  $\alpha_n$  is greater than the current minimum. If this condition is satisfied, then switch the old minimal value with  $\alpha_n$  and search for the new minimal entry of this row.

If none of these conditions is met, we drop the computed value without updating the current column and repeat these steps for the next values unequal to zero of the current row. This cap of values for each row also has the following disadvantages: by having a too small maximum of nonzeros per row, a qualitative  $A$ -orthogonalization cannot be performed. To avoid this case we only take values of  $\omega$  greater than one, which seems to be sufficient. Also, if a row has already reached the maximum number of nonzeros, additional but relatively small values may be dropped. This

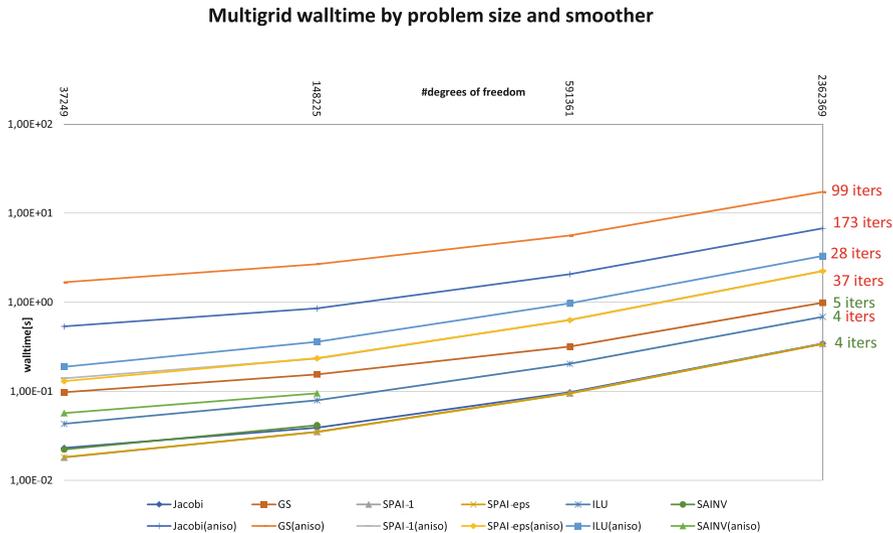


Fig. 5 GPU smoother comparison, isotropic and anisotropic Poisson benchmarks

can become an issue if the sum of these small numbers leads to a relevant entry in a later iteration. For a comparison, Fig. 5 depicts the time-to-solution for V-cycle multigrid using different strong smoothers on a P100 GPU. All smoothers are constructed using 8 Richardson iterations with (reasonably damped if necessary) preconditioners such as Jacobi, Gauss–Seidel, ILU-0, SPAI-1, SPAI- $\epsilon$  and SAINV. We set up the benchmark case from a 2D Poisson problem in the isotropic case and with two-sided anisotropies in the grid to harden the problem even for well-ordered ILU approaches. The SPAI approaches are the best choice for the smoother on the GPU.

**Machine Learning** Finally we started investigating how to construct approximate inverses using methods from Machine Learning [53]. The basic idea here is to treat  $A^{-1}$  as a discrete function in the course of a function regression process. The neural network therefore learns how to deduct (extrapolate) an approximation of the inverse. Once trained with many data pairs of matrices and their inverse (a sparse representation of it) a neural network like a multilayer perceptron can be able to approximate inverses rapidly. As a starting point we have employed the finite element method for the Poisson equation on different domains with linear basis functions and have used it to generate expedient systems of equations to solve. Problems of this kind are usually based on sparse M-matrices with characteristics that can be used to reduce the calculation time and effort of the neural network training and evaluation. Our results show that given the pre-defined quality of the preconditioner (equivalent to the  $\epsilon$  in a SPAI- $\epsilon$  method), we can by far numerically outperform even Gauss–Seidel. Using Tensorflow [1] and numpy [4], the learning algorithm can even be performed on the GPU. Here we have used a three-layered

fully-connected perceptron with fifty neurons in each layer plus input and output layers, and employed the resulting preconditioners in a Richardson method to solve the mentioned problem on a three times refined L-domain with a fixed number of degrees of freedom. The numerical effort of each evaluation of the neural network is basically the effort of a matrix-vector-multiplication for each layer in which the matrix size depends on the number of neurons per layer ( $M$ ) and the non zero entries ( $N$ ) of the input matrix, like  $O(NM)$  for the first layer. The inner layers' effort, without input and output layer, just depends on the number of neurons. The crucial task now is to balance the quality of the resulting approximation and the effort to evaluate the network. We use fully connected feed-forward multilayer perceptrons as a starting point. Fully connected means that every neuron in the network is connected to each neuron of the next layer. Moreover there are no backward connections between the different layers (feed-forward). The evaluation of such neural networks is a sequence of chained matrix-vector products.

The entries of the system matrix are represented vector-wise in the input layer (cf. Fig. 6). In the same way, our output layer contains the entries of the approximate inverse. Between these layers we can add a number of hidden layers consisting of hidden neurons. How many hidden neurons we need to create strong approximate inverses is a key design decision and we discuss this below. In general our supervised training algorithm is a backward propagation with random initialisation. Alongside a linear propagation function  $i_{total} = \mathbf{W} \cdot o_{total} + b$  with the total (layer) net input  $i_{total}$ , the weight matrix  $\mathbf{W}$ , the vector for the bias weights  $b$  and the total output of the previous layer  $o_{total}$ , we use the rectified linear unit (ReLu) function as activation function  $\alpha(x)$  and thus we can calculate the output  $y$  of each neuron as  $y := \alpha(\sum_j o_j \cdot w_{ij})$ . Here  $o_j$  is the output of the preceding sending units and  $w_{ij}$  are the corresponding weights between the neurons.

For the optimization we use the L2 error function and update the weights with  $w_{ij}^{(t+1)} = w_{ij}^{(t)} + \gamma \cdot o_i \cdot \delta_j$ , with the output  $o_i$  of the sending unit and learning rate  $\gamma$ .  $\delta_j$  symbolises the gradient decent method:

$$\delta_j = \begin{cases} f'(i_j) \cdot (\hat{o}_j - o_j) & \text{if neuron } j \text{ is an output neuron} \\ f'(i_j) \cdot \sum_{k \in S} (\delta_k \cdot w_{kj}) & \text{if neuron } j \text{ is a hidden neuron.} \end{cases}$$

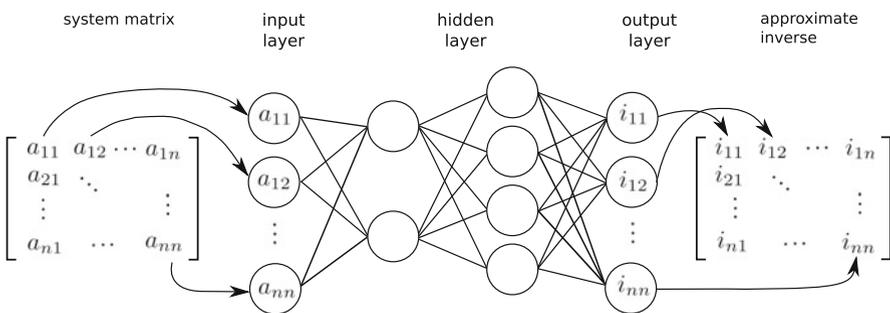


Fig. 6 Model of a neural network for matrix inversion, cf. [53]

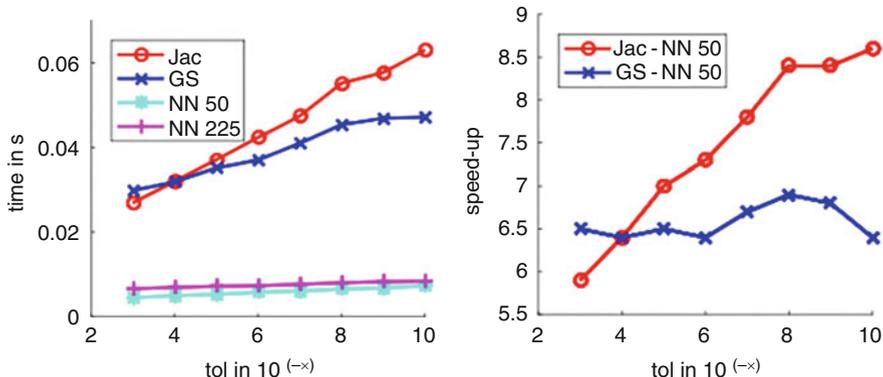


Fig. 7 Results for the defect correction with the neural network, cf. [53]

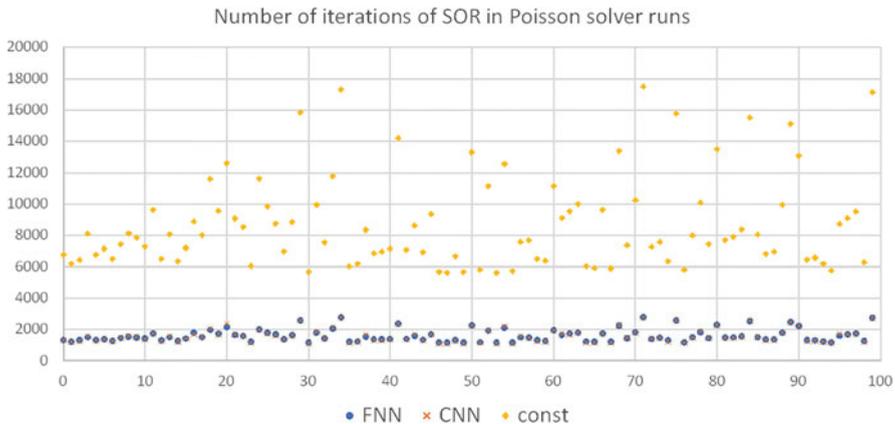
For details concerning the test/training algorithm we refer to a previous publication [53]. For the defect correction prototype, we find a significant speedup for a moderately anisotropic Poisson problem, see Fig. 7.

### 3.2 Autotuning with Artificial Neural Networks

Inspired by our usage of Approximate Inverses generated by artificial neural networks (ANNs), we exploit (Feed Forward-) neural networks (FNN) for the automatic tuning of solver parameters. We were able to show that it is possible to use such an approach to provide much better a-priori choices for the parametrisation of iterative linear solvers. In detailed studies for 2D Poisson problems we conducted benchmarks for many test matrices and autotuning systems using FNNs as well as convolutionary neural networks (CNNs) to predict the  $\omega$  parameter in a SOR solver. In Fig. 8 we depict 100 randomly chosen samples of this study. It can be seen that even for good a-priori choices of  $\omega$  the NN-driven system can compete whilst ‘bad’ choices (labeled constant) might lead to a stalling solver.

### 3.3 Further Development of Sum-Factorized Matrix-Free DG Methods

While we were able to achieve good node-level performance with our matrix-free DG methods in the first funding period, our initial implementations still did not utilize more than about 10% of the theoretical peak FLOP throughput. In the second funding period, we systematically improved on those results by focusing on several aspects of our implementation:



**Fig. 8** Result for 100 samples of the FNN-based autotuning system for the  $\omega$  parameter in SOR

**Introduction of Block-Based DOF Processing** Our implementation is based on DUNE-PDELAB, a very flexible discretization framework for both continuous and discontinuous discretizations of PDEs. In order to support a wide range of discretizations, PDELab has a powerful system for mapping DOFs to vector and matrix entries. Due to this flexibility, the mapping process is rather expensive. On the other hand, Discontinuous Galerkin values will always be blocked in a cell-wise manner. This can be exploited by only ever mapping the first degree of freedom associated with each cell and then assuming that all subsequent values for this cell are directly adjacent to the first entry. We have added a special ‘DG codepath’ to DUNE-PDELAB which implements this optimization.

**Avoiding Unnecessary Memory Transfers** As all of the values for each cell are stored in consecutive locations in memory, we can further optimize the framework behavior by skipping the customary gather/scatter steps before and after the assembly of each cell and facet integral. This is implemented by replacing the data buffer normally passed to the integration kernels with a dummy buffer that stores a pointer to the first entry in the global vector/matrix and directly operates on the global values. This is completely transparent to the integration kernels, as they only ever access the global data through a well-defined interface on these buffer objects. Together with the previous optimization, these two changes have allowed us to reduce the overhead of the framework infrastructure on assembly times from more than 100% to less than 5%.

**Explicit Vectorization** The DG implementation used in the first phase of the project was written as scalar code and relied on the compiler’s auto vectorization support to utilize the SIMD instruction set of the processor, which we tried to facilitate by providing compile time loop bounds and aligned data structures. In the second phase, we have switched to explicit vectorization with a focus on AVX2, which is a common foundation instruction set across all current  $\times 86$ -based HPC

processors. We exploit the possibilities of our C++ code base and use a well-abstracted library which wraps the underlying compiler intrinsic calls [23]. In a separate project [34], we are extending this functionality to other SIMD instruction sets like AVX512.

**Loop Reordering and Fusion** While vectorization is required to fully utilize modern CPU architectures, it is not sufficient. We also have to feed the execution units with a substantial number of mutually independent chains of computation ( $\approx 40$ – $50$  on current CPUs). This amount of parallelism can only be extracted from typical DG integration kernels by fusing and reordering computational loops. In contrast to other implementations of matrix-free DG assembly [22, 43], we do not group computations across multiple cells or facets, but instead across quadrature points and multiple input/output variables. In 3D, this works very well for scalar PDEs that contain both the solution itself and its gradient, which adds up to four quantities that exactly fit into an AVX2 register.

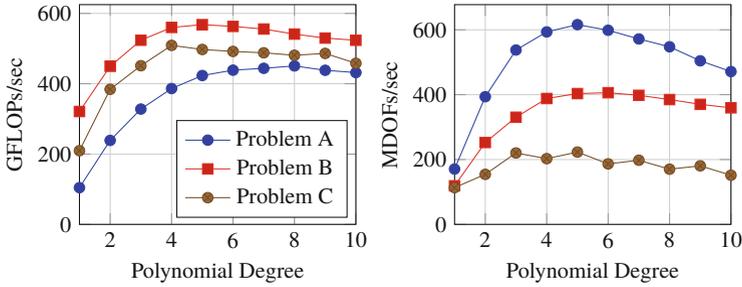
**Results** Table 2 compares the throughput and the hardware efficiency of our matrix-free code for two diffusion-reaction problems A (axis-parallel grid, constant coefficients per cell) and B (affine geometries, variable coefficients per cell) with a matrix-based implementation. Figure 9 compares throughput and floating point performance of our implementation for these problems as well as an additional problem C with multi-linear geometries, demonstrating that we are able to achieve more than 50% of theoretical peak FLOP rate on this machine as well as a good computational processing rate as measured in DOFs/s.

While our work in this project was mostly focused on scalar diffusion-advection-reaction problems, we have also applied the techniques shown here to projection-based Navier–Stokes solvers [51]. One important lesson learned was the unsustainable amount of work required to extend our approach to different problems and/or hardware architectures. This led us to develop a Python-based code generator in a new project [34], which provides powerful abstractions for the building blocks listed above. This toolbox can be extended and combined in new ways to achieve performance comparable to hand-optimized code. Especially for more complex problems involving systems of equations, there are a large number of possible ways to group variables and their derivatives into sum factorization kernels due to our approach of vectorizing over multiple quantities within a single cell. The resulting search space is too large for manual exploration, which the above project solved by the addition of benchmark-driven automatic comparison of those variants. Finally, initial results show good scalability of our code as shown by the strong scaling results in Fig. 10. Our implementation shows good scalability until we reach a local problem size of just 18 cells, where we still need to improve the asynchronicity of ghost data communication and assembly.

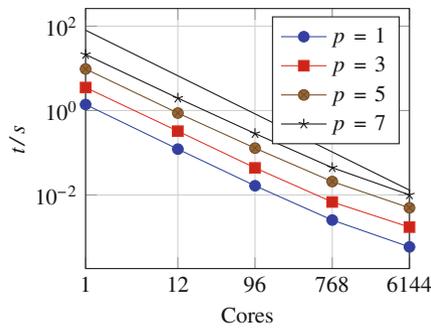
**Table 2** Full operator application, 2× Intel Xeon E5-2698v3 2.3 GHz (32 cores), for two problems of different complexity, and for comparison matrix assembly for the simpler problem and matrix-based operator application

p	Matrix-free A		Matrix-free B		Matrix-based		Matrix assembly	
	$\frac{\text{DOF}}{s}$	$\frac{\text{GFLOP}}{s}$	$\frac{\text{DOF}}{s}$	$\frac{\text{GFLOP}}{s}$	$\frac{\text{DOF}}{s}$	$\frac{\text{GFLOP}}{s}$	$\frac{\text{DOF}}{s}$	$\frac{\text{GFLOP}}{s}$
1	$1.70 \times 10^8$	104	$1.19 \times 10^8$	321	$2.06 \times 10^8$	24.3	$1.60 \times 10^7$	345
2	$3.93 \times 10^8$	238	$2.52 \times 10^8$	450	$6.42 \times 10^7$	27.3	$8.71 \times 10^6$	371
3	$5.38 \times 10^8$	328	$3.30 \times 10^8$	524	$2.69 \times 10^7$	29.8	$4.66 \times 10^6$	368
4	$5.95 \times 10^8$	387	$3.88 \times 10^8$	560	$9.54 \times 10^6$	23.5	$2.57 \times 10^6$	301
5	$6.17 \times 10^8$	424	$4.03 \times 10^8$	568	$4.58 \times 10^6$	23.3	$1.93 \times 10^6$	307
6	$5.99 \times 10^8$	439	$4.06 \times 10^8$	563	$2.31 \times 10^6$	23.5	$1.21 \times 10^6$	231
7	$5.70 \times 10^8$	442	$3.98 \times 10^8$	556	$1.46 \times 10^6$	30.3	$6.65 \times 10^5$	143
8	$5.41 \times 10^8$	445	$3.85 \times 10^8$	541	$6.14 \times 10^5$	24.1	$8.74 \times 10^5$	198
9	$5.05 \times 10^8$	439	$3.70 \times 10^8$	530	$2.98 \times 10^5$	26.2	$7.01 \times 10^5$	133
10	$4.71 \times 10^8$	432	$3.59 \times 10^8$	524	—	—	—	—

Note that the matrix-based computations use significantly smaller problem sizes due to memory constraints.  $p$  denotes the polynomial degree of the DG space



**Fig. 9** Floating point performance in GFLOPs/s and throughput in MDOFs/s for full operator application, 2× Intel Xeon E5-2698v3 2.3 GHz for all model problems



**Fig. 10** Runtimes for strong scalability on IWR compute cluster (416 nodes with 2× E5-2630 v3 each, 64 GiB/node, QDR Infiniband)

### 3.4 Hybrid Solvers for Discontinuous Galerkin Schemes

In Sect. 3.3 we concentrated on the performance of matrix-free *operator application*. This is sufficient for instationary problems with explicit time integration, but in case of stationary problems or implicit time integration, (linear) algebraic systems need to be solved. This requires operator application and robust, scalable preconditioners.

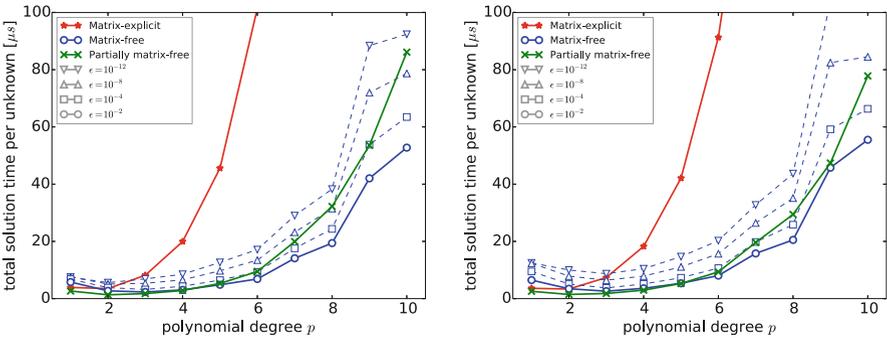
For this we extended hybrid AMG-DG preconditioners [8] in a joint work with Eike Müller from Bath University, UK, [10]. In a solver for matrices arising from higher order DG discretizations the basic idea is to perform all computations on the DG system in a matrix-free fashion and to explicitly assemble only a matrix in a low-order subspace which is significantly smaller. In the sense of subspace correction methods [58] we employ a splitting

$$V_{DG}^p = \sum_{T \in \mathcal{T}_h} V_T^p + V_c$$

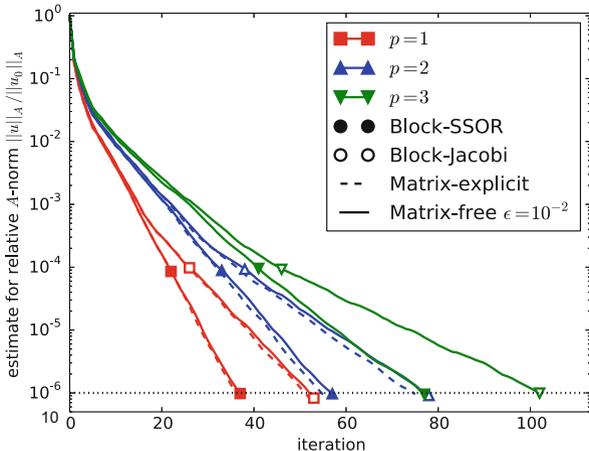
where  $V_T^p$  is the finite element space of polynomial degree  $p$  on element  $T$  and the coarse space  $V_c$  is either the lowest-order conforming finite element space  $V_h^1$  on the mesh  $\mathcal{T}_h$ , or the space of piecewise constants  $V_h^0$ . Note that the symmetric weighted interior penalty DG method from [21] reduces to the cell-centered finite volume method with two-point flux approximation on  $V_h^0$ . Note also, that the system on  $V_c$  can be assembled without assembling the large DG system.

For solving the blocks related to  $V_T^p$ , two approaches have been implemented. In the first (named *partially matrix-free*), these diagonal blocks are factorized using LAPACK and each iteration uses a backsolve. In the second approach the diagonal blocks are solved iteratively to low accuracy using matrix-free sum factorization. Both variants can be used in additive and multiplicative fashion. Figure 11 shows that the partially matrix-free variant is optimal for polynomial degree  $p \leq 5$ , but starting from  $p = 6$ , the fully matrix-free version starts to outperform all other options.

In order to demonstrate the robustness of our hybrid AMG-DG method we use the permeability field of the SPE10 benchmark problem [14] within a heterogeneous elliptic problem. This is considered to be a hard test problem in the porous media community. The DG method from [21] is employed. Figure 12 depicts results for different variants and polynomial degrees run in parallel on 20 cores. A moderate increase with the polynomial degree can be observed. With respect to time-to-solution (not reported) the additive (block Jacobi) partially matrix-free variant is to be preferred for polynomial degree larger than one.



**Fig. 11** Total solution time for different implementations and a range of block-solver tolerances  $\epsilon$  for the Poisson problem (left) and the diffusion problem with spatially varying coefficients (right), cf. [10]



**Fig. 12** Convergence history for SPE10 benchmark. The relative energy norm is shown for polynomial degrees 1 (red squares), 2 (blue upward triangles) and 3 (green downward triangles). Results for the block-SSOR smoother are marked by filled symbols and results for the block-Jacobi smoother by empty symbols. cf. [10]

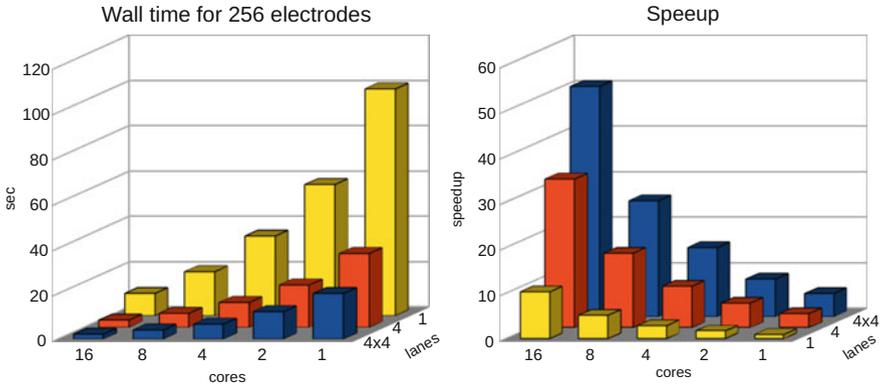
### 3.5 Horizontal Vectorization of Block Krylov Methods

Methods like Multiscale FEM (see Sect. 4), optimization and inverse problems need to invert the same operator for many right-hand-side vectors. This leads to a block problem, by the following conceptual reformulation:

$$\text{foreach } i \in [0, N] : \text{solve } Ax_i = b_i \quad \rightarrow \quad \text{solve } AX = B,$$

with matrices  $X = (x_0, \dots, x_N)$ ,  $B = (b_0, \dots, b_N)$ . Such problems can be solved using Block Krylov solvers. The benefit is that the approximation space can grow faster, as the solver orthogonalizes the updates for all right-hand-sides. Even for a single right-hand-side Block Krylov based enriched Krylov methods can be used to accelerate the solution process.

Preconditioners and the actual Krylov solver can be sped up using horizontal vectorization. Assuming  $k$  right-hand-sides we observe that the scalar product yields a  $k \times k$  dense matrix and has  $O(k^2)$  complexity. While the mentioned larger approximation space should improve the convergence rate, this is only true for weaker preconditioners, therefore we pursued a different strategy and approximate the scalar product matrix by a sparse matrix, so that we again retain  $O(k)$  complexity. In particular we consider the case of a diagonal or block-diagonal matrix. The diagonal matrix basically results in  $k$  independent solvers running in parallel, so that the performance gain is solely based on SIMD vectorization and the associated favorable memory layout.



**Fig. 13** Horizontal vectorization of a linear solver for 256 right-hand-side vectors. Timings on a Haswell-EP (E5-2698v3, 16 cores, AVX2, 4 lanes). Comparison with 1–16 cores and no SIMD, AVX (4 lanes), AVX (4 × 4 lanes)

For the implementation in DUNE-ISTL we use platform portable C++ abstractions of SIMD intrinsics, building on the VC library[38] and some DUNE specific extensions. We use this to exchange the underlying data type of the right-hand-side and the solution vector, so that we no longer store scalars, but SIMD vectors. This is possible when using generic programming techniques, like C++ templates, and yields a row-wise storage of the dense matrices  $X$  and  $B$ . This row-wise storage is optimal and ensures a high arithmetic intensity. The implementations of the Krylov solvers have to be adapted to the SIMD data types, since some operations, like casts and branches, are not available generically for SIMD data types. As a side effect, all preconditioners, including the AMG, are now fully vectorized.

Performance tests using 256 right-hand-side vectors for a 3D Poisson problem show nearly optimal speedup on a 64 core system (see Fig. 13). The tests are carried out on a Haswell-EP (E5-2698v3, 16 cores, AVX2, 4 lanes). We observe a speedup of 50, while the theoretical speedup is 64.

## 4 Adaptive Multiscale Methods

The main goal in the second funding phase was a distributed adaptive multilevel implementation of the localized reduced basis multi-scale method (LRBMS [49]). Like Multiscale FEM (MsFEM), LRBMS is designed to work on heterogenous multiscale or large scale problems. It performs particularly well for problems that exhibit scale separation with effects on both a fine and a coarse scale contributing to the global behavior. Unlike MsFEM, LRBMS is best applied in multi-query settings in which a parameterized PDE needs to be solved many times for different parameters. As an amalgam of domain decomposition and model order reduction

techniques, the computational domain is partitioned into a coarse grid with each macroscopic grid cell representing a subdomain for which, in an offline pre-compute stage, local reduced bases are constructed. Appropriate coupling is then applied to produce a global solution approximation from localized data. For increased approximation fidelity we can integrate localized global solution snapshots into the bases, or the local bases can adaptively be enriched in the online stage, controlled by a localized a-posteriori error estimator.

### 4.1 Continuous Problem and Discretization

We consider elliptic parametric multi-scale problems on a domain  $\Omega \subset \mathbb{R}^d$  where we look for  $p(\boldsymbol{\mu}) \in Q$  that satisfy

$$b(p(\boldsymbol{\mu}), q; \boldsymbol{\mu}) = l(q) \quad \text{for all } q \in H_0^1(\Omega), \tag{1}$$

$\boldsymbol{\mu}$  are parameters with  $\boldsymbol{\mu} \in \mathcal{P} \subset \mathbb{R}^p$ ,  $p \in \mathbb{N}$ . We let  $\epsilon > 0$  be the multi-scale parameter associated with the fine scale. For demonstration purposes we consider a particular linear elliptic problem setup in  $\Omega \subset \mathbb{R}^d$  ( $d = 2, 3$ ) that exhibits a multiplicative splitting in the quantities affected by the multi-scale parameter  $\epsilon$ . It is a model for the so called global pressure  $p(\boldsymbol{\mu}) \in H_0^1(\Omega)$  in two phase flow in porous media, where the total scalar mobility  $\lambda(\boldsymbol{\mu})$  is parameterized.  $\kappa_\epsilon$  denotes the heterogenous permeability tensor and  $f$  the external forces. Hence, we seek  $p$  that satisfies weakly in  $H_0^1(\Omega)$ ,

$$-\nabla \cdot (\lambda(\boldsymbol{\mu})\kappa_\epsilon \nabla p(\boldsymbol{\mu})) = f \quad \text{in } \Omega. \tag{2}$$

With  $A(x; \boldsymbol{\mu}) := \lambda(\boldsymbol{\mu})\kappa_\epsilon(x)$  this gives rise to the following definition of the forms in (1)

$$b(p(\boldsymbol{\mu}), q; \boldsymbol{\mu}) := \int_{\Omega} A(\boldsymbol{\mu}) \nabla p \cdot \nabla q, \quad l(q) := \int_{\Omega} f q.$$

For the discretization we first require a triangulation  $\mathcal{T}_H$  of  $\Omega$  for the macro level. We call the elements  $T \in \mathcal{T}_H$  subdomains. We then require each subdomain be covered with a fine partition  $\tau_h(T)$  in a way that  $\mathcal{T}_H$  and  $\tau_h := \sum_{T \in \mathcal{T}_H} \tau_h(T)$  are nested. We denote by  $\mathcal{F}_H$  the faces of the coarse triangulation and by  $\mathcal{F}_h$  the faces of the fine triangulation.

Let  $V(\tau_h) \subset H^2(\tau_h)$  denote any approximate subset of the broken Sobolev space  $H^2(\tau_h) := \{q \in L^2(\Omega) \mid q|_t \in H^2(t) \forall t \in \tau_h\}$ . We call  $p_h(\boldsymbol{\mu}) \in V(\tau_h)$  an approximate solution of (1), if

$$b_h(p_h(\boldsymbol{\mu}), v; \boldsymbol{\mu}) = l_h(v; \boldsymbol{\mu}) \quad \text{for all } v \in V(\tau_h). \tag{3}$$

Here, the DG bilinear form  $b_h$  and the right hand side  $l_h$  are chosen according to the SWIPDG method [20], i.e.

$$b_h(v, w; \boldsymbol{\mu}) := \sum_{t \in \tau_h} \int_t A(\boldsymbol{\mu}) \nabla v \cdot \nabla w + \sum_{e \in \mathcal{F}(\tau_h)} b_h^e(v, w; \boldsymbol{\mu})$$

$$l_h(v; \boldsymbol{\mu}) := \sum_{t \in \tau_h} \int_t f v,$$

where the DG coupling bilinear forms  $b_h^e$  for a face  $e$  is given by

$$b_h^e(v, w; \boldsymbol{\mu}) := \int_e \langle A(\boldsymbol{\mu}) \nabla v \cdot \mathbf{n}_e \rangle [w] + \langle A(\boldsymbol{\mu}) \nabla w \cdot \mathbf{n}_e \rangle [v] + \frac{\sigma_e(\boldsymbol{\mu})}{|e|^\beta} [v][w].$$

The LRBMS method allows for a variety of discretizations, i.e. approximation spaces  $V(\tau_h)$ . As a particular choice of an underlying high dimensional approximation space we choose  $V(\tau_h) = Q_h^k := \bigoplus_{T \in \mathcal{T}_H} Q_h^{k,T}$ , where the discontinuous local spaces are defined as

$$Q_h^{k,T} := Q_h^{k,T}(\tau_h(T)) := \{q \in L^2(T) \mid q|_t \in \mathbb{P}_k(t) \forall t \in \tau_h(T)\}.$$

## 4.2 Model Reduction

For model order reduction in the LRBMS method we choose the reduced space  $Q_{\text{red}} := \bigoplus_{T \in \mathcal{T}_H} Q_{\text{red}}^T \subset Q_h^k$  with local reduced approximation spaces  $Q_{\text{red}}^T \subset Q_h^{k,T}$ . We denote  $p_{\text{red}}(\boldsymbol{\mu})$  to be the reduced solution of (3) in  $Q_{\text{red}}$ . This formulation naturally leads to solving a sparse blocked linear system similar to a DG approximation with high polynomial degree on the coarse subdomain grid.

The construction of subdomain reduced spaces  $Q_{\text{red}}^T$  is again very flexible. Initialization with shape functions on  $T$  up to order  $k$  ensures a minimum fidelity. Basis extensions can be driven by a discrete weak greedy approach which incorporates localized solutions of the global system. Depending on available computational resources, and given a suitable localizable a-posteriori error estimator  $\eta(p_{\text{red}}(\boldsymbol{\mu}), \boldsymbol{\mu})$ , we can forego computing global high-dimensional solutions altogether and only rely on online enrichment to extend  $Q_{\text{red}}^T$  ‘on the fly’. With online enrichment, given a reduced solution  $p_{\text{red}}(\boldsymbol{\mu})$  for some arbitrary  $\boldsymbol{\mu} \in \mathcal{P}$ , we first compute local error indicators  $\eta^T(p_{\text{red}}(\boldsymbol{\mu}), \boldsymbol{\mu})$  for all  $T \in \mathcal{T}_H$ . If  $\eta^T(p_{\text{red}}(\boldsymbol{\mu}), \boldsymbol{\mu})$  is greater than some prescribed bound  $\delta_{\text{tol}} > 0$ , we solve on a overlay region  $\mathcal{N}(T) \supset T$  and extend  $Q_{\text{red}}^T$  with  $p_{\mathcal{N}(T)}(\boldsymbol{\mu})|_T$ . Inspired by results in [17] we set the overlay region’s diameter  $\text{diam}(\mathcal{N}(T))$  of the order  $\mathcal{O}(\text{diam}(T)|\log(\text{diam}(T))|)$ . In practice we use the completely on-line/off-line decomposable error estimator developed in [49, Sec. 4] which in turn is based on the idea of producing a

conforming reconstruction of the diffusive flux  $\lambda(\boldsymbol{\mu})\kappa_\epsilon \nabla_h p_h(\boldsymbol{\mu})$  in some Raviart-Thomas-Nédélec space  $V_h^l(\tau_h) \subset H_{div}(\Omega)$  presented in [21, 56].

This process is then repeated until either a maximum number of enrichment steps occur or  $\eta^T(p_{\text{red}}(\boldsymbol{\mu}), \boldsymbol{\mu}) \leq \delta_{\text{tol}}$ .

---

#### Algorithm 4 Schematic representation of the LRBMS pipeline

---

**Require:**  $P_{\text{train}} \subset \mathcal{P}$

**Require:** Reconstruction operator  $R_h(p_{\text{red}}(\boldsymbol{\mu})) : Q_{\text{red}}(\mathcal{T}_H) \rightarrow Q_h^k(\tau_h)$

```

1: function GREEDYBASISGENERATION( $\delta_{\text{grdy}}, \eta(p_{\text{red}}(\boldsymbol{\mu}), \boldsymbol{\mu})=\text{None}$ )
2:   if  $\eta(p_{\text{red}}(\boldsymbol{\mu}), \boldsymbol{\mu})$  is not None then
3:      $E \leftarrow \{\eta(p_{\text{red}}(\boldsymbol{\mu}_i), \boldsymbol{\mu}_i) \mid \boldsymbol{\mu}_i \in P_{\text{train}}\}$ 
4:   else
5:      $E \leftarrow \{\|R_h(p_{\text{red}}(\boldsymbol{\mu}_i)) - p_h(\boldsymbol{\mu}_i)\| \mid \boldsymbol{\mu}_i \in P_{\text{train}}\}$ 
6:   while  $E \neq \emptyset$  AND  $\max(E) \geq \delta_{\text{grdy}}$  do
7:      $i \leftarrow \text{argmax}(E)$ 
8:     compute  $p_h(\boldsymbol{\mu}_i)$ 
9:     for all  $T \in \mathcal{T}_H$  do
10:      extend  $Q_{\text{red}}^T$  with  $p_h(\boldsymbol{\mu})|_T$ 
11:      $E \leftarrow E \setminus E_i$ 

12: Generate  $\mathcal{T}_H$  ▷ Offline Phase
13: for all  $T \in \mathcal{T}_H$  do
14:   create  $\tau_h(T)$ 
15:   init  $Q_{\text{red}}^T$  with DG shape functions of order  $k$ 

16: GREEDYBASISGENERATION( $\dots$ ) ▷ Optional
17: compute  $p_{\text{red}}(\boldsymbol{\mu})$  for arbitrary  $\boldsymbol{\mu}$  ▷ Online phase
18: for all  $T \in \mathcal{T}_H$  do ▷ Optional Adaptive Enrichment
19:    $\eta \leftarrow \eta^T(p_{\text{red}}(\boldsymbol{\mu}), \boldsymbol{\mu})$ 
20:   while  $\eta \geq \delta_{\text{tol}}$  do
21:     compute  $p_{N(T)}(\boldsymbol{\mu})$ 
22:      $Q_{\text{red}}^T \leftarrow p_{N(T)}(\boldsymbol{\mu})|_T$ 

```

---

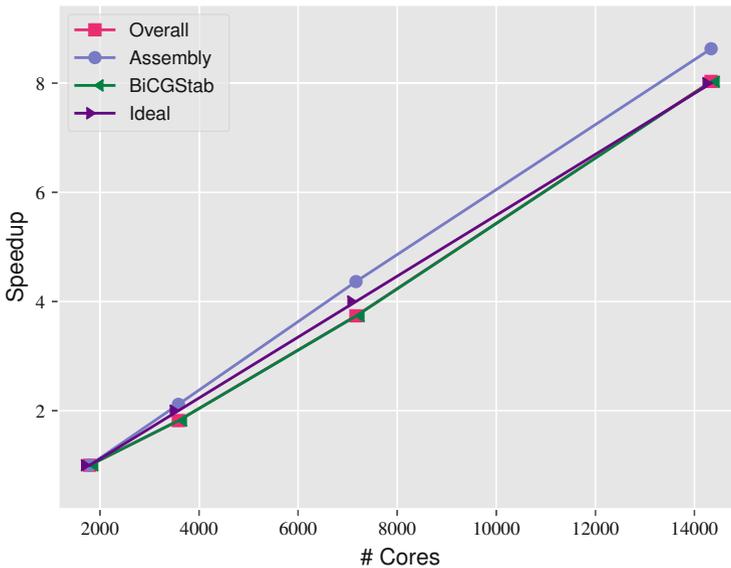
### 4.3 Implementation

We base our MPI-parallel implementation of LRBMS on the serial version developed previously. In this setup the high-dimensional quantities and all grid structures are implemented in DUNE. The model order reduction as such is implemented in Python using pyMOR [45]. The model reduction algorithms in pyMOR follow a solver agnostic design principle. Abstract interfaces allow for example projections, greedy algorithms or reduced data reconstruction to be written without knowing details of the PDE solver backend. The global macro grid  $\mathcal{T}_H$  can be any MPI-enabled DUNE grid manager with adjustable overlap size for the domain decomposition, we currently use DUNE-YASPGRID. The fine grids  $\tau_h(T)$  are constructed using the same grid manager as on the macro scale, with MPI subcom-

municators. These are currently limited to a size of one (rank-local), however the overall scalability could benefit from dynamically sizing these subcommunicators to balance communication overhead and computational intensity as demonstrated in [36, Sec. 2.2]. The assembly of the local (coupling) bilinear forms is done in DUNE-GDT [24], with pyMOR/Python bindings facilitated through DUNE-XT [46], where DUNE-GRID-GLUE [19] generates necessary grid views for the SWIPDG coupling between otherwise unrelated grids. Switching to DUNE-GRID-GLUE constitutes a major step forward in robustness of the overall algorithm, compared to our previous manually implemented approach to matching independent local grids for coupling matrices assembly.

We have identified three major challenges in parallelizing all the steps in LRBMS:

1. **Global solutions  $p_h(\mu)$  of the blocked system in Eq. (3) with an appropriate MPI-parallel iterative solver.** With the serial implementation already using DUNE-ISTL as the backend for matrix and vector data, we only had to generate an appropriate communication configuration for the blocked SWIPDG matrix structure to make the BiCGStab solver usable in our context. We tested this setup on the SuperMUC Petascale System in Garching. The results in Fig. 14 show very near ideal speedup from 64 nodes with 1792 MPI ranks up to a full island with 512 nodes and 14336 ranks.



**Fig. 14** Localized Reduced Basis Method: Block-SWIPDG speedup results; linear system solve (green), discretization and system assembly (blue), theoretic ideal speedup (violet) and actual achieved speedup for the overall run time (red). Simulation on  $\sim 7.9 \cdot 10^6$  cubical cells shows minimum 94% parallel efficiency, scaling from 64 to 512 nodes (SuperMUC Phase 2)

2. **(Global) Reduced systems also need a distributed solver.** By design all reduced quantities in pyMOR are, at the basic, unabstracted level, NumPy arrays [57]. Therefore we cannot easily re-use the DUNE-ISTL based solvers for the high-dimensional systems. Our current implementation gathers these reduced system matrices from all MPI-ranks to rank 0, recombines them, solves the system with a direct solver and scatters the solution. There is great potential in making this step more scalable by either using a distributed sparse direct solver like Mumps [3] or translating the data into the DUNE-ISTL backend.
3. **Adaptive online enrichment is inherently load imbalanced due to its localized error estimator guidance.** The load imbalance results from one rank idling while waiting to receive updates to a basis on a subdomain in its overlap region from another rank. This idle time can be minimized by encapsulating the update in a `MPIFuture` described in Sect. 2.1. This will allow the rank to continue in its own enrichment process until the updated basis is actually needed in a subsequent step.

## 5 Uncertainty Quantification

The solution of stochastic partial differential equations (SPDEs) is characterized by extremely high dimensions and poses great (computational) challenges. Multilevel Monte Carlo (MLMC) algorithms attract great interest due to their superiority over the standard Monte Carlo approach. Based on Monte Carlo (MC), MLMC retains the properties of independent sampling. To overcome the slow convergence of MC, where many computationally expensive PDEs have to be solved, MLMC combines in a proper way cheap MC estimators and expensive MC estimators, achieving (much) faster convergence. One of the critical components of the MLMC algorithms is the way in which the coarser levels are selected. The exact definition of the levels is an open question and different approaches exist. In the first funding phase, Multiscale FEM was used as a coarser level in MLMC. During the second phase, the developed parallel MLMC algorithms for uncertainty quantification were further enhanced. The main focus was on exploring the capabilities of the renormalization approach for defining the coarser levels in the MLMC algorithm, and on using MLMC as a coarse grained parallelization approach.

Here, we employ MLMC to exemplarily compute the mean flux through saturated porous media with prescribed pressure drop and known distribution of the random coefficients.

**Mathematical Problem** As a model problem in  $\mathbb{R}^2$  or  $\mathbb{R}^3$ , we consider steady state single phase flow in random porous media:

$$-\nabla \cdot [k(x, \omega)\nabla p(x, \omega)] = 0 \text{ for } x \in D = (0, 1)^d, \omega \in \Omega$$

subject to the boundary conditions  $p_{x=0} = 1$  and  $p_{x=1} = 0$  and zero flux on other boundaries. Here  $p$  is pressure,  $k$  is scalar permeability, and  $\omega$  is a random vector. The quantity of interest is the mean (expected value)  $E[Q]$  of the total flux  $Q$  through the inlet of the unit cube i.e.,  $Q(x, \omega) := \int_{x=0} k(x, \omega) \partial_n p(x, \omega) dx$ . Both the coefficient  $k(x, \omega)$  and the solution  $p(x, \omega)$  are subject to uncertainty, characterized by the random vector  $\omega$  in a properly defined random space  $\Omega$ . For generating permeability fields we consider the practical covariance  $C(x, y) = \sigma^2 \exp(-\|x - y\|_2 / \lambda)$ . An algorithm based on forward and inverse Fourier transform over the circulant covariance matrix is used to generate the permeability field. For solving the deterministic PDEs a Finite Volume method on a cell centered grid is used [32]. More details and further references can be found in a previous paper [47].

**Monte Carlo Simulations** To quantify the uncertainty, and compute the mean of the flux we use a MLMC algorithm. Let  $\omega_M$  be a random vector over a properly defined probability space, and  $Q_M$  be the corresponding flux. It is known that  $E[Q_M]$  can be made arbitrarily close to  $E[Q]$  by choosing  $M$  sufficiently large. The standard MC algorithm convergences very slowly, proportionally to the variance over the square root of the number of samples, which makes it often unfeasible. MLMC introduces  $L$  levels with the  $L$ -th level coinciding with the considered problem, and exploits the telescopic sum identity:

$$E[Q_M^L(\omega)] = E[Q_M^0(\omega)] + E[Q_M^1(\omega) - Q_M^0(\omega)] + \dots + E[Q_M^L(\omega) - Q_M^{L-1}(\omega)]$$

The notation  $Y^l = Q^l - Q^{l-1}$  is also used. The main idea of MLMC is to properly define levels, and combine a large number of cheap simulations, that are able to approximate the variance well, with a small number of expensive correction simulations providing the needed accuracy. For details on Monte Carlo and MLMC we refer to previous publications [32, 47] and the references therein. Here, the target is to estimate the mean flux on a fine grid, and we define the levels as discretizations on coarser grids. In order to define the permeability at the coarser levels we use the renormalization approach.

MLMC has previously run the computations at each level with the same tolerance. However, in order to evaluate the number of samples needed per level, one has to know the variance at each level. Because the exact variance is not known in advance, MLMC starts by performing simulations with a prescribed, moderate number of samples per level. The results are used to evaluate the variance at each level, and thus to evaluate the number of samples needed per level. This procedure can be repeated several times in an Estimate–Solve cycle. At each estimation step, information from all levels is needed, which leads to a synchronization point in the parallel realization of the algorithm. This may require dynamic redistribution of the resources after each new evaluation.

MLMC can provide a coarse graining in the parallelization. A well balanced algorithm has to account for several factors: (1) How many processes should be allocated per level; (2) how many processes should be allocated per deterministic problem including permeability generation; (3) how to parallelize the permeability

generation; (4) which of the parallelization algorithms for deterministic problems available in EXA-DUNE should be used; (5) should each level be parallelized separately and if not, how to group the levels for efficient parallelization. The last factor is the one giving coarse grain parallelization opportunities. For the generation of the permeability, we use the parallel MPI implementation of the FFTW library. As deterministic solver, we use a parallel implementation of the conjugate gradient scheme preconditioned with AMG, provided by DUNE-ISTL. Both of them have their own internal domain decomposition.

We shortly discuss one Estimate-Solve cycle of the MLMC algorithm. Without loss of generality we assume 3-level MLMC. Suppose that we have already computed the required number of samples per level (i.e., we are after Estimate and before Solve). Let us denote by  $N_i, i = \{0, 1, 2\}$  the number of required realizations per level for  $\widehat{Y}_l$ , by  $p_i$  the number of processes allocated per  $\widehat{Y}_i$ , by  $p_{l_i}^g$  the respective group size of processes working on a single realization, by  $n$  the number of realizations for each group of levels, with  $t_i$  the respective time for solving a single problem once, and finally with  $p^{\text{total}}$  the total number of available processes. Then we can compute the total CPU time for the current Estimate-Solve cycle as

$$T_{\text{CPU}}^{\text{total}} = N_0 t_0 + N_1 t_1 + N_2 t_2.$$

Ideally each process should take  $T_{\text{CPU}}^p = T_{\text{CPU}}^{\text{total}}/p^{\text{total}}$ . Dividing the CPU time needed for one  $\widehat{Y}_i$  by  $T_{\text{CPU}}^p$ , we get a continuous value for the number of processes on a given level  $p_i^{\text{ideal}} = N_i t_i / T_{\text{CPU}}^p$  for  $i = \{0, 1, 2\}$ . Then we can take  $p_i = \lfloor p_i^{\text{ideal}} \rfloor$ . To obtain an integer value for the number of processes allocated per level, first we construct a set of all possible splits of the original problem as a combination of subproblems (e.g., parallelize level 2 separately and the combination of levels 0 and 1, or parallelize all levels simultaneously, etc.). Each element of this set is evaluated independently, and all combinations of upper and lower bounds are calculated, such that  $p_i^{\text{ideal}}$  is divisible by  $p_{l_i}^g$ ,  $\sum_{i=0}^2 p_i < p^{\text{total}}$  and  $p_i \leq N_i p_{l_i}^g$ . Traversing, computing and summing the computational time needed for each element gives us a time estimation. Then we select the element (grouping of levels) with minimal computational time. To tackle the distribution of the work on a single level, a similar approach can be employed. Due to the large dimension of the search tree a heuristic technique can be employed. Here we consider a simple predefined group size for each deterministic problem, having in mind that when using AMG the work for a single realization at the different levels is proportional to the unknowns at this level.

**Numerical Experiments** Results for a typical example are shown in Fig. 15. The parameters are  $\sigma = 2.75, \lambda = 0.3$ . The tests are done on SuperMUC-NG, LRZ Munich on a dual Intel Skylake Xeon Platinum 8174 node. Due to the stochasticity of the problem, we plot the speedup multiplied with the proportion of the tolerance. The renormalization has shown to be a very good approach for defining the coarser levels in MLMC. The proposed parallelization algorithm gives promising scalability results. It is weakly coupled to the number of samples that MLMC estimates.

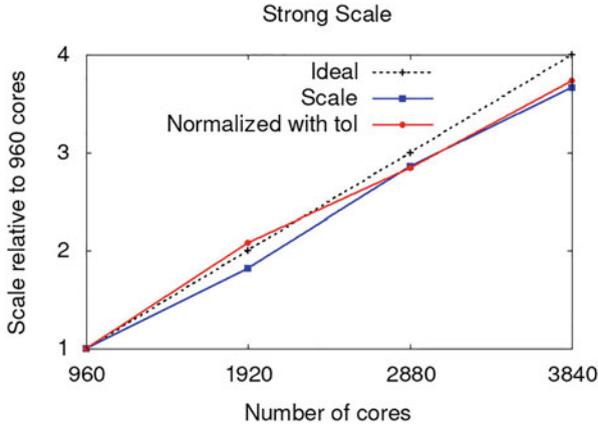


Fig. 15 Scalability of the MLMC approach

Although the search for an optimal solution is an NP-hard problem, the small number of levels enables a full traversing of the search tree. It can be further improved by automatically selecting the number of processes per group that solves a single problem. One also may consider introducing interrupts between the MPI communicators on a level to further improve the performance.

## 6 Land-Surface Flow Application

To test some of the approaches developed in the EXA-DUNE project, especially the usage of sum-factorized operator evaluation with more complex problems, we developed an application to simulate coupled surface-subsurface flow for larger geographical areas. This is a topic with high relevance for a number of environmental questions from soil protection and groundwater quality up to weather and climate prediction.

One of the central aims of the new approach developed in the project is to be able to relate a physical meaning to the parameter functions used in each grid cell. This is not possible with the traditional structured grid approaches as the necessary resolution would be prohibitive. To avoid the excessive memory requirements of completely unstructured grids we build on previous results for block-structured meshes and use a two-dimensional unstructured grid on the surface which is extruded in a structured way in vertical direction. However, more flexible discretization schemes are needed for such grids, compared to the usual cell-centered Finite-Volume approaches.

## 6.1 Modelling and Numerical Approach

To describe subsurface flow we use the Richards equation [52]

$$\frac{\partial \theta(\psi)}{\partial t} - \nabla \cdot [k(\psi) (\nabla \psi + \mathbf{e}_g)] + q_w = 0$$

where  $\theta$  is the volumetric water content,  $\psi$  the soil water potential,  $k$  the hydraulic conductivity,  $\mathbf{e}_g$  the unit vector pointing in the direction of gravity and  $q_w$  a volumetric source or sink term.

In nearly all existing models for coupled surface-subsurface flow, the kinematic-wave approximation is used for surface flow, which only considers surface slope as driving force and does not even provide a correct approximation of the steady-state solution. The physically more realistic shallow-water-equations are used rarely, as they are computationally expensive. We use the diffusive-wave approximation, which still retains the effects of water height on run-off, yields a realistic steady-state solution and is a realistic approximation for flow on vegetation covered ground [2]:

$$\frac{\partial h}{\partial t} - \nabla \cdot [D(h, \nabla h) \nabla (h + z)] = f_c, \quad (4)$$

where  $h$  is the height of water over the surface level  $z$ ,  $f_c$  is a source-sink term (which is used for the coupling) and the diffusion coefficient  $D$  is given by

$$D(h, \nabla h) = \frac{h^\alpha}{C \cdot \|\nabla (h + z)\|^{1-\gamma}}$$

with  $\|\cdot\|$  referring to the Euclidean norm and  $\alpha$ ,  $\gamma$  and  $C$  are empirical constants. In the following we use  $\alpha = \frac{5}{3}$  and  $\gamma = \frac{1}{2}$  to obtain Manning's formula and a friction coefficient of  $C = 1$ .

Both equations are discretised with a cell-centered Finite-Volume scheme and alternatively with a SWIPDG scheme in space (see Sect. 3) and an appropriate diagonally implicit Runge–Kutta scheme in time for the subsurface and an explicit Runge–Kutta scheme for the surface flow. Upwinding is used for the calculation of conductivity in subsurface flow [5] and for the water height in the diffusion term in surface flow.

First tests have shown that the formulation of the diffusive-wave approximation from the literature as given by Eq. (4) does not result in a numerically stable solution if the gradient becomes very small, as then a gradient approaching zero is multiplied by a diffusion coefficient going to infinity. A much better behaviour is achieved by rewriting the equation as

$$\frac{\partial h}{\partial t} - \nabla \cdot \left[ \frac{h^\alpha}{C} \cdot \frac{\nabla (h + z)}{\|\nabla (h + z)\|^{1-\gamma}} \right] = f_c,$$

where the rescaled gradient  $\frac{\nabla(h+z)}{\|\nabla(h+z)\|^{1-\gamma}}$  is always going to zero when  $\nabla(h+z)$  is going to zero as long as  $\gamma < 1$  and the new diffusion coefficient  $h^\alpha/C$  is bounded.

Due to the very different time-scales for surface and subsurface flow, an operator-splitting approach is used for the coupled system. A new coupling condition has been implemented, which is a kind of Dirichlet-Neumann coupling, but guarantees a mass-conservative solution. With a given height of water on the surface (from the initial condition or the last time step modified by precipitation and evaporation), subsurface flow is calculated with a kind of Signorini boundary condition, where all surface water is infiltrated in one time step as long as the necessary gradient is not larger than the pressure resulting from the water ponding on the surface (in infiltration) and potential evaporation rates are maintained as long as the pressure at the surface is not below a given minimal pressure (during evaporation). The advantage of the new approach is that it does not require a tracking of the sometimes complicated boundary between wet and dry surface elements, that it yields no unphysical results and that the solution is mass-conservative even if not iterated until convergence.

Parallelisation is obtained by an overlapping or non-overlapping domain-decomposition (depending on the grid). However, only the two-dimensional surface grid is partitioned whereas the vertical direction is kept on one process due to the strong coupling. Thus there is also no need for communication of surface water height for the coupling, as the relevant data is always stored in the same process. The linear equation systems are solved with the BiCGstab-solver from DUNE-ISTL with Block-ILU0 as preconditioner. The much larger mesh size in horizontal direction compared to the vertical direction results in strong coupling of the unknowns in the vertical direction. The Block-ILU0 scheme provides an almost exact solver of the strongly coupled blocks in the vertical direction and is thus a very effective scheme. Furthermore, one generally has a limited number of cells in the vertical direction and extends the mesh in horizontal direction to simulate larger regions. Thus the good properties of the solver are retained when scaling up the size of the system.

## 6.2 Performance Optimisations

As the time steps in the explicit scheme for surface flow can get very small due to the stability limit, a significant speedup can be achieved by using a semi-implicit scheme, where the non-linear coefficients are calculated with the solution from the previous time step or iteration. However, if the surface is nearly completely dry, this could lead to numerical problems, thus an explicit scheme is still used under nearly dry conditions with an automatic switching between both.

While matrix-free DG solvers with sum-factorization can yield excellent per node performance (Sect. 3.3), it is a rather tedious task to implement them for new partial differential equations. Therefore, a code-generation framework is currently being developed in a project related to EXA-DUNE [33]. The framework is used to

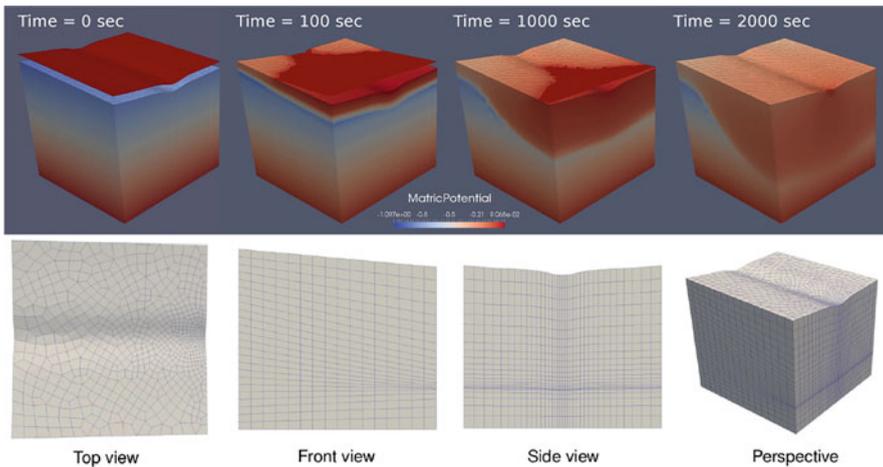
implement an alternative optimized version of the solver for the Richards equation as this is the computationally most expensive part of the computations. Expensive material functions like the van Genuchten model including several power functions are replaced by cubic spline approximations, which allow a fast vectorized incorporation of flexible material functions to simulate strongly heterogeneous systems. Sum-factorisation is used in the operator evaluations for the DG-discretization with a selectable polynomial degree.

A special pillar grid has been developed as a first realisation of a 2.5D grid [33]. It adds a vertical dimension to a two-dimensional grid (which is either structured or unstructured). However, as the current version still produces a full three-dimensional mesh at the moment, future developments are necessary to exploit the full possibilities of the approach.

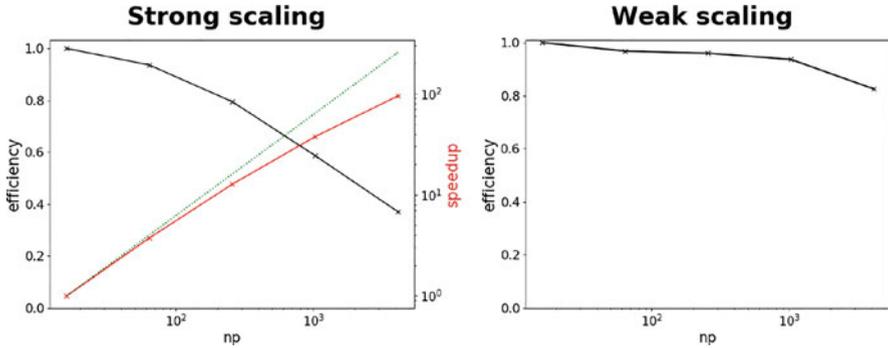
### 6.3 Scalability and Performance Tests

Extensive tests covering infiltration as well as exfiltration have been performed (e.g. Fig. 16) to test the implementation and the new coupling condition. Good scalability is achieved in strong as well as in weak scaling experiments on up to 256 nodes and 4096 cores of the bwForCluster in Heidelberg (Fig. 17). Simulations for a large region with topographical data taken from a digital elevation model (Fig. 18) have been conducted as well.

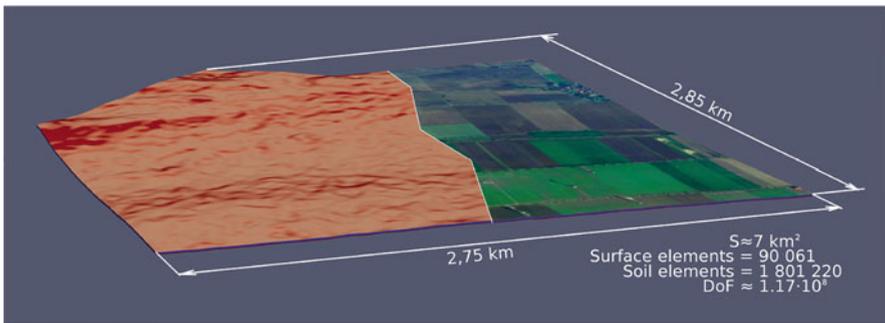
With the generated code-based solver for the Richards equation a substantial fraction of the system’s peak performance (up to 60% on a Haswell-CPU) can be



**Fig. 16** Surface runoff and infiltration of 5 cm water into a dry coarse sand (top) and the unstructured 2.5D mesh used for the simulations (bottom)

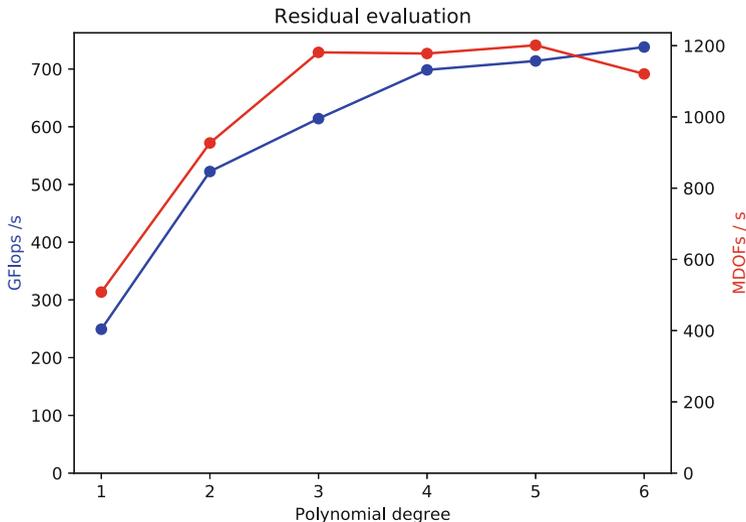


**Fig. 17** Results of a strong (left) and weak (right) scalability test with a coupled run-off and infiltration experiment on 1 to 256 nodes (16 to 4096 Intel Xeon E5-2630 v3 2.4 GHz CPU cores) of the bwForCluster at IWR in Heidelberg



**Fig. 18** Pressure distribution with an overlay of the landscape taken from Google Earth calculated in a simulation of water transport in a real landscape south of Brunswick simulated on 30 nodes with 1200 cores of HLRN-IV in Göttingen ( $2 \times$  Intel Skylake Gold 6148 2.4 GHz CPUs)

utilized due to the combination of sum factorization and vectorisation (Fig. 19). For the Richards equation (as for other PDEs before) the number of millions of degrees of freedom per second is independent of the polynomial degree with this approach. We measure a total speedup of 3 compared to the naive implementation in test simulations on a Intel Haswell Core i7-5600U 2.6 GHz CPU with first order DG base functions on a structured  $32 \times 32 \times 32$  mesh for subsurface and  $32 \times 32$  grid for surface flow. Even higher speedups are expected with higher-order base functions and matrix-free iterative solvers. The fully-coupled combination of the Richards solver obtained with the code generation framework and surface-runoff is tested with DG up to fourth order on structured as well as on unstructured grids. Parallel simulations are possible as well.



**Fig. 19** Performance of the Richards solver implemented with the code generation framework for EXA-DUNE

## 7 Conclusion

In EXA-DUNE we extended the DUNE software framework in several directions to make it ready for the exascale architectures of the future which will exhibit a significant increase in node level performance through massive parallelism in form of cores and vector instructions. Software abstractions are now available that enable asynchronicity as well as parallel exception handling and several use cases for these abstractions have been demonstrated in this paper: resilience in multigrid solvers and several variants of asynchronous Krylov methods. Significant progress has been achieved in hardware-aware iterative linear solvers: we developed preconditioners for the GPU based on approximate sparse inverses, developed matrix-free operator application and preconditioners for higher-order DG methods and our solvers are now able to vectorize over multiple right hand sides. These building blocks have then been used to implement adaptive localized reduced basis methods, multilevel Monte-Carlo methods and a coupled surface-subsurface flow solver on up to 14k cores. The EXA-DUNE project has spawned a multitude of research projects, running and planned, as well as further collaborations in each of the participating groups. We conclude that the DUNE framework has made a major leap forward due to the EXA-DUNE project and work on the methods investigated here will continue in future research projects.

## References

1. Abadi, M.: TensorFlow: large-scale machine learning on heterogeneous systems (2015). <http://tensorflow.org/>
2. Alonso, R., Santillana, M., Dawson, C.: On the diffusive wave approximation of the shallow water equations. *Eur. J. Appl. Math.* **19**(5), 575–606 (2008)
3. Amestoy, P.R., Duff, I.S., Koster, J., J.-Y., L.: A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Anal. Appl.* **23**(1), 15–41 (2001)
4. Ascher, D., Dubois, P., Hinsén, K., Huginin, J., Oliphant, T.: Numerical Python. Lawrence Livermore National Laboratory, Livermore (1999)
5. Bastian, P.: A fully-coupled discontinuous Galerkin method for two-phase flow in porous media with discontinuous capillary pressure. *Computat. Geosci.* **18**(5), 779–796 (2014)
6. Bastian, P., Blatt, M., Dedner, A., Engwer, C., Klöforn, R., Kornhuber, R., Ohlberger, M., Sander, O.: A generic grid interface for parallel and adaptive scientific computing. Part II: Implementation and tests in DUNE. *Computing* **82**(2–3), 121–138 (2008). <https://doi.org/10.1007/s00607-008-0004-9>
7. Bastian, P., Blatt, M., Dedner, A., Engwer, C., Klöforn, R., Ohlberger, M., Sander, O.: A generic grid interface for parallel and adaptive scientific computing. Part I: abstract framework. *Computing* **82**(2–3), 103–119 (2008). <https://doi.org/10.1007/s00607-008-0003-x>
8. Bastian, P., Blatt, M., Scheichl, R.: Algebraic multigrid for discontinuous Galerkin discretizations of heterogeneous elliptic problems. *Numer. Linear Algebra Appl.* **19**(2), 367–388 (2012). <https://doi.org/10.1002/nla.1816>
9. Bastian, P., Engwer, C., Fahlke, J., Geveler, M., Göddeke, D., Iliev, O., Ippisch, O., Milk, R., Möring, J., Müthing, S., Ohlberger, M., Ribbrock, D., Turek, S.: Advances concerning multiscale methods and uncertainty quantification in EXA-DUNE. In: *Software for Exascale Computing—SPPEXA 2013–2015. Lecture Notes in Computational Science and Engineering*. Springer Berlin (2016)
10. Bastian, P., Müller, E., Müthing, S., Piatkowski, M.: Matrix-free multigrid block-preconditioners for higher order discontinuous Galerkin discretisations. *J. Comput. Phys.* **394**, 417–439 (2019). <https://doi.org/10.1016/j.jcp.2019.06.001>
11. Bland, W., Bosilca, G., Bouteiller, A., Hérault, T., Dongarra, J.: A proposal for user-level failure mitigation in the MPI-3 standard. Department of Electrical Engineering and Computer Science, University of Tennessee (2012)
12. Buis, P., Dyksen, W.: Efficient vector and parallel manipulation of tensor products. *ACM Trans. Math. Softw.* **22**(1), 18–23 (1996). <https://doi.org/10.1145/225545.225548>
13. Cantwell, C., Nielsen, A.: A minimally intrusive low-memory approach to resilience for existing transient solvers. *J. Sci. Comput.* **78**(1), 565–581 (2019). <https://doi.org/10.1007/s10915-018-0778-7>
14. Christie, M., Blunt, M.: Tenth SPE comparative solution project: a comparison of upscaling techniques. In: *SPE Reservoir Simulation Symposium* (2001). <https://doi.org/10.2118/72469-PA>
15. Chronopoulos, A., Gear, C.W.: s-step iterative methods for symmetric linear systems. *J. Comput. Appl. Math.* **25**(2), 153–168 (1989)
16. Dongarra, J., Hittinger, J., Bell, J., Chacón, L., Falgout, R., Heroux, M., Hovland, P., Ng, E., Webster, C., Wild, S., Pao, K.: Applied mathematics research for exascale computing. Tech. rep., U. S. Department of Energy, Office of Science, Advanced Scientific Computing Research Program (2014). <http://science.energy.gov/~media/ascr/pdf/research/am/docs/EMWGreport.pdf>
17. Elfverson, D., Ginting, V., Henning, P.: On multiscale methods in Petrov–Galerkin formulation. *Numer. Math.* **131**(4), 643–682 (2015). <https://doi.org/10.1007/s00211-015-0703-z>
18. Engwer, C., Dreier, N.A., Altenbernd, M., Göddeke, D.: A high-level C++ approach to manage local errors, asynchrony and faults in an MPI application. In: *Proceedings of 26th EuroMicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP 2018)* (2018)

19. Engwer, C., Müthing, S.: Concepts for flexible parallel multi-domain simulations. In: Dickopf, T., Gander, M.J., Halpern, L., Krause, R., Pavarino, L.F. (eds.) *Domain Decomposition Methods in Science and Engineering XXII*, pp. 187–195. Springer, Berlin (2016)
20. Ern, A., Stephansen, A., Vohralík, M.: Guaranteed and robust discontinuous Galerkin a posteriori error estimates for convection-diffusion-reaction problems. *J. Comput. Appl. Math.* **234**(1), 114–130 (2010). <https://doi.org/10.1016/j.cam.2009.12.009>
21. Ern, A., Stephansen, A., Zunino, P.: A discontinuous Galerkin method with weighted averages for advection-diffusion equations with locally small and anisotropic diffusivity. *IMA J. Numer. Anal.* **29**(2), 235–256 (2009). <https://doi.org/10.1093/imanum/drm050>
22. Fehn, N., Wall, W., Kronbichler, M.: A matrix-free high-order discontinuous Galerkin compressible Navier–Stokes solver: a performance comparison of compressible and incompressible formulations for turbulent incompressible flows. *Numer. Methods Fluids* **89**(3), 71–102 (2019)
23. Fog, A.: VCL C++ vector class library (2017). <http://www.agner.org/optimize/vectorclass.pdf>
24. Fritze, R., Leibner, T., Schindler, F.: dune-gdt (2016). <https://doi.org/10.5281/zenodo.593009>
25. Gagliardi, F., Moreto, M., Olivieri, M., Valero, M.: The international race towards exascale in Europe. *CCF Trans. High Perform. Comput.* **1**(1), 3–13 (2019). <https://doi.org/10.1007/s42514-019-00002-y>
26. Geveler, M., Ribbrock, D., Göddeke, D., Zajac, P., Turek, S.: Towards a complete FEM-based simulation toolkit on GPUs: unstructured grid finite element geometric multigrid solvers with strong smoothers based on sparse approximate inverses. *Comput. Fluids* **80**, 327–332 (2013). <https://doi.org/10.1016/j.compfluid.2012.01.025>
27. Ghysels, P., Vanroose, W.: Hiding global synchronization latency in the preconditioned conjugate gradient algorithm. *Parallel Comput.* **40**(7), 224–238 (2014)
28. Gmeiner, B., Huber, M., John, L., Rüde, U., Wohlmuth, B.: A quantitative performance study for Stokes solvers at the extreme scale. *J. Comput. Sci.* **17**, 509–521 (2016). <https://doi.org/10.1016/j.jocs.2016.06.006>
29. Göddeke, D., Altenbernd, M., Ribbrock, D.: Fault-tolerant finite-element multigrid algorithms with hierarchically compressed asynchronous checkpointing. *Parallel Comput.* **49**, 117–135 (2015). <https://doi.org/10.1016/j.parco.2015.07.003>
30. Gupta, S., Patel, T., Tiwari, D., Engelmann, C.: Failures in large scale systems: long-term measurement, analysis, and implications. In: *Proceedings of the 30th IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC)* (2017)
31. Huber, M., Gmeiner, B., Rüde, U., Wohlmuth, B.: Resilience for massively parallel multigrid solvers. *SIAM J. Sci. Comput.* **38**(5), S217–S239
32. Iliev, O., Mohring, J., Shegunov, N.: Renormalization based MLMC method for scalar elliptic SPDE. In: *International Conference on Large-Scale Scientific Computing*, pp. 295–303 (2017)
33. Kempf, D.: Code generation for high performance PDE solvers on modern architectures. Ph.D. thesis, Heidelberg University (2019)
34. Kempf, D., Heß, R., Müthing, S., Bastian, P.: Automatic code generation for high-performance discontinuous Galerkin methods on modern architectures (2018). arXiv:1812.08075
35. Keyes, D.: Exaflop/s: the why and the how. *Comptes Rendus Mécanique* **339**(2–3), 70–77 (2011). <https://doi.org/10.1016/j.crme.2010.11.002>
36. Klawonn, A., Köhler, S., Lanser, M., Rheinbach, O.: Computational homogenization with million-way parallelism using domain decomposition methods. *Comput. Mech.* **65**(1), 1–22 (2020). <https://doi.org/10.1007/s00466-019-01749-5>
37. Kohl, N., Thönnies, D., Drzisga, D., Bartuschat, D., Rüde, U.: The HyTeG finite-element software framework for scalable multigrid solvers. *Int. J. Parallel Emergent Distrib. Syst.* **34**(5), 477–496 (2019). <https://doi.org/10.1080/17445760.2018.1506453>
38. Kretz, M., Lindenstruth, V.: Vc: A C++ library for explicit vectorization. *Softw. Practice Exp.* **42**(11), 1409–1430 (2012). <https://doi.org/10.1002/spe.1149>
39. Kronbichler, M., Kormann, K.: A generic interface for parallel cell-based finite element operator application. *Comput. Fluids* **63**, 135–147 (2012). <https://doi.org/10.1016/j.compfluid.2012.04.012>

40. Lengauer, C., Apel, S., Bolten, M., Chiba, S., Rüde, U., Teich, J., Größlinger, A., Hannig, F., Köstler, H., Claus, L., Grebhahn, A., Groth, S., Kronawitter, S., Kuckuk, S., Rittich, H., Schmitt, C., Schmitt, J.: ExaStencils: Advanced multigrid solver generation. In: *Software for Exascale Computing—SPPEXA 2nd phase*. Springer, Berlin (2020)
41. Liang, X., Di, S., Tao, D., Li, S., Li, S., Guo, H., Chen, Z., Cappello, F.: Error-controlled lossy compression optimized for high compression ratios of scientific datasets. In: *IEEE Big Data*, pp. 438–447 (2018). <https://doi.org/10.1109/BigData.2018.8622520>
42. Lu, Y.: Paving the way for China exascale computing. *CCF Trans. High Perform. Comput.* **1**(2), 63–72. (2019). <https://doi.org/10.1007/s42514-019-00010-y>
43. Luporini, F., Ham, D.A., Kelly, P.H.J.: An algorithm for the optimization of finite element integration loops. *ACM Trans. Math. Softw.* **44**(1), 3:1–3:26 (2017). <https://doi.org/10.1145/3054944>
44. Message Passing Interface Forum: MPI: a message-passing interface standard version 3.1 (2015). <https://www.mpi-forum.org/docs/>
45. Milk, R., Rave, S., Schindler, F.: pyMOR—generic algorithms and interfaces for model order reduction. *SIAM J. Sci. Comput.* **38**(5), S194–S216 (2016). <https://doi.org/10.1137/15M1026614>
46. Milk, R., Schindler, F., Leibner, T.: Extending DUNE: the dune-xt modules. *Arch. Numer. Softw.* **5**(1), 193–216 (2017). <https://doi.org/10.11588/ans.2017.1.27720>
47. Mohring, J., Milk, R., Ngo, A., Klein, O., Iliev, O., Ohlberger, M., Bastian: uncertainty quantification for porous media flow using multilevel Monte Carlo. In: *International Conference on Large-Scale Scientific Computing*, pp. 145–152 (2015)
48. Müthing, S., Piatkowski, M., Bastian, P.: High-performance implementation of matrix-free high-order discontinuous Galerkin methods (2017). arXiv:1711.10885
49. Ohlberger, M., Schindler, F.: Error control for the localized reduced basis multiscale method with adaptive on-line enrichment. *SIAM J. Sci. Comput.* **37**(6), A2865–A2895 (2015). <https://doi.org/10.1137/151003660>
50. Orszag, S.: Spectral methods for problems in complex geometries. *J. Comput. Phys.* **37**(1), 70–92 (1980). [https://doi.org/10.1016/0021-9991\(80\)90005-4](https://doi.org/10.1016/0021-9991(80)90005-4)
51. Piatkowski, M., Müthing, S., Bastian, P.: A stable and high-order accurate discontinuous Galerkin based splitting method for the incompressible Navier–Stokes equations. *J. Comput. Phys.* **356**, 220–239 (2018). <https://doi.org/10.1016/j.jcp.2017.11.035>
52. Richards, L.A.: Capillary conduction of liquids through porous mediums. *Physics* **1**, 318–333 (1931)
53. Ruelmann, H., Geveler, M., Turek, S.: On the prospects of using machine learning for the numerical simulation of PDEs: training neural networks to assemble approximate inverses. *Eccomas Newsletter* 27–32 (2018)
54. Teranishi, K., Heroux, M.: Toward local failure local recovery resilience model using MPI-ULFM. In: *Proceedings of the 21st European MPI Users’ Group Meeting, EuroMPI/ASIA ’14*, pp. 51:51–51:56 (2014). <https://doi.org/10.1145/2642769.2642774>
55. Turek, S., Göddeke, D., Becker, C., Buijssen, S., Wobker, S.: FEAST—realisation of hardware-oriented numerics for HPC simulations with finite elements. *Concurrency Comput. Pract. Exp.* **22**(6), 2247–2265 (2010)
56. Vohralík, M.: A posteriori error estimates for lowest-order mixed finite element discretizations of convection-diffusion-reaction equations. *SIAM J. Numer. Anal.* **45**(4), 1570–1599 (2007). <https://doi.org/10.1137/060653184>
57. van der Walt, S., Colbert, S.C., Varoquaux, G.: The NumPy array: a structure for efficient numerical computation. *Comput. Sci. Eng.* **13**(2), 22–30 (2011). <https://doi.org/10.1109/MCSE.2011.37>
58. Xu, J.: Iterative methods by space decomposition and subspace correction. *SIAM Rev.* **34**(4), 581–613 (1992). <https://doi.org/10.1137/1034116>

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

