

HYBRID-PARALLEL SPARSE MATRIX-VECTOR MULTIPLICATION WITH EXPLICIT COMMUNICATION OVERLAP ON CURRENT MULTICORE-BASED SYSTEMS

GERALD SCHUBERT, HOLGER FEHSKE

*Institute of Physics, University of Greifswald
Felix-Hausdorff-Str. 6, 17487 Greifswald, Germany*

GEORG HAGER, GERHARD WELLEIN

*Erlangen Regional Computing Center, University of Erlangen-Nuremberg
Martensstr. 1, 91058 Erlangen, Germany*

Received June, 2011

Revised July, 2011

Communicated by Guest Editors

ABSTRACT

We evaluate optimized parallel sparse matrix-vector operations for several representative application areas on widespread multicore-based cluster configurations. First the single-socket baseline performance is analyzed and modeled with respect to basic architectural properties of standard multicore chips. Beyond the single node, the performance of parallel sparse matrix-vector operations is often limited by communication overhead. Starting from the observation that nonblocking MPI is not able to hide communication cost using standard MPI implementations, we demonstrate that explicit overlap of communication and computation can be achieved by using a dedicated communication thread, which may run on a virtual core. Moreover we identify performance benefits of hybrid MPI/OpenMP programming due to improved load balancing even without explicit communication overlap. We compare performance results for pure MPI, the widely used “vector-like” hybrid programming strategies, and explicit overlap on a modern multicore-based cluster and a Cray XE6 system.

Keywords: sparse matrix-vector multiply, hybrid parallelization, multicore

1. Introduction

Many problems in science and engineering involve the solution of large eigenvalue problems or extremely sparse systems of linear equations arising from, e.g., the discretization of partial differential equations. Sparse matrix-vector multiplication (spMVM) is the dominant operation in many of those solvers and may easily consume most of the total run time. A highly efficient scalable spMVM implementation is thus fundamental, and complements advancements and new development in the high-level algorithms.

For more than a decade there has been an intense debate about whether the hi-

Listing 1: CRS sparse matrix-vector multiplication kernel

```

1 do i = 1, Nr
2   do j = row_ptr(i), row_ptr(i+1) - 1
3     C(i) = C(i) + val(j) * B(col_idx(j))
4   enddo
5 enddo

```

erarchical structure of current HPC systems needs to be considered in parallel programming, or if pure MPI is sufficient. Hybrid approaches based on MPI+OpenMP have been implemented in codes and kernels for various applications areas and compared with traditional MPI implementations. Most results are hardware-specific, and sometimes contradictory. In this paper we analyze hybrid MPI+OpenMP variants of a general parallel spMVM operation. Beyond the naive approach of using OpenMP for parallelization of kernel loops (“vector mode”) we also employ a hybrid “task mode” to overcome or mitigate a weakness of standard MPI implementations: the lack of truly asynchronous communication in nonblocking MPI calls. We test our implementation against pure MPI approaches for two application scenarios on an InfiniBand cluster as well as a Cray XE6 system.

1.1. *Related work*

In recent years the performance of various spMVM algorithms has been evaluated by several groups [1, 2, 3]. Covering different matrix storage formats and implementations on various types of hardware, they have reviewed a more or less large number of publicly available matrices and reported on the obtained performance. Scalable parallel spMVM implementations have also been proposed [4, 5], mostly based on an MPI-only strategy. Hybrid parallel spMVM approaches have already been devised before the emergence of multicore processors [6, 7]. Recently a “vector mode” approach could not compete with a scalable MPI implementation for a specific problem on a Cray system [4]. There is no up-to-date literature that systematically investigates novel features like multicore, ccNUMA node structure, and simultaneous multithreading (SMT) for hybrid parallel spMVM.

1.2. *Sparse matrix-vector multiplication and node-level performance model*

A possible definition of a “sparse” matrix is that the number of its nonzero entries grows only linearly with the matrix dimension; however, not all problems are easily scaled, so in general a sparse matrix may be defined to contain “mainly” zero entries. Since keeping such a matrix in computer memory with all zeros included is usually out of the question, an efficient format to store the nonzeros only is required. The most widely used variant is “Compressed Row Storage” (CRS) [8]. It does not exploit specific features that may emerge from the underlying physical

problem like, e.g., block structures, symmetries, etc., but is broadly recognized as the most efficient format for general sparse matrices on cache-based microprocessors. All nonzeros are stored in one contiguous array `val(:)`, row by row, and the starting offsets of all rows are contained in a separate array `row_ptr(:)`. Array `col_idx(:)` contains the original column index of each matrix entry. A matrix-vector multiplication with a right-hand-side (RHS) vector `B(:)` and a result vector `C(:)` can then be written as shown in Listing 1. Here N_r is the number of matrix rows. While arrays `C(:)` and `val(:)` are traversed contiguously, access to `B(:)` is indexed and may potentially cause very low spatial and temporal locality in this data stream.

The performance of spMVM operations on a single compute node is often limited by main memory bandwidth. Code balance [9] is thus a good metric to establish a simple performance model. We assume the average length of the inner (j) loop to be $N_{nzt} = N_{nz}/N_r$, where N_{nz} is the total number of nonzero matrix entries. Then the contiguous data accesses in the CRS code generate $(8+4+16/N_{nzt})$ bytes of memory traffic for a single inner loop iteration, where the first two contributions come from the matrix `val(:)` (8 bytes) and the index array `col_idx(:)` (4 bytes), while the last term reflects the update of `C(i)` (write allocate + evict). The indirect access pattern to `B(:)` is determined by the sparsity structure of the matrix and can not be modeled in general. However, `B(:)` needs to be loaded at least once from main memory, which adds another $8/N_{nzt}$ bytes per inner iteration. Limited cache size and nondiagonal access typically require loading at least parts of `B(:)` multiple times in a single MVM. This is quantified by a machine- and problem-specific parameter κ : For each additional time that `B(:)` is loaded from main memory, κ increases by $8/N_{nzt}$. Together with the computational intensity of 2 flops per j iteration the code balance is

$$B_{\text{CRS}} = \left(\frac{12 + 24/N_{nzt} + \kappa}{2} \right) \frac{\text{bytes}}{\text{flop}} = \left(6 + \frac{12}{N_{nzt}} + \frac{\kappa}{2} \right) \frac{\text{bytes}}{\text{flop}} . \quad (1)$$

On the node level B_{CRS} can be used to determine an upper performance limit by measuring the node memory bandwidth (e.g., using the STREAM benchmark) and assuming $\kappa = 0$. Moreover, κ can be determined experimentally from the spMVM floating point performance and the memory bandwidth drawn by the CRS code (see Sect. 2). Since the “slimmest” matrices used here have $N_{nzt} \approx 7 \dots 15$, each additional access to `B(:)` incurs a nonnegligible contribution to the data transfer in those cases.

Note that this simple model neglects performance-limiting aspects beyond bandwidth bottlenecks like in-cache transfer time, load imbalance, communication and/or synchronization overhead, and the adverse effects of nonlocal memory access across ccNUMA locality domains (LDs).

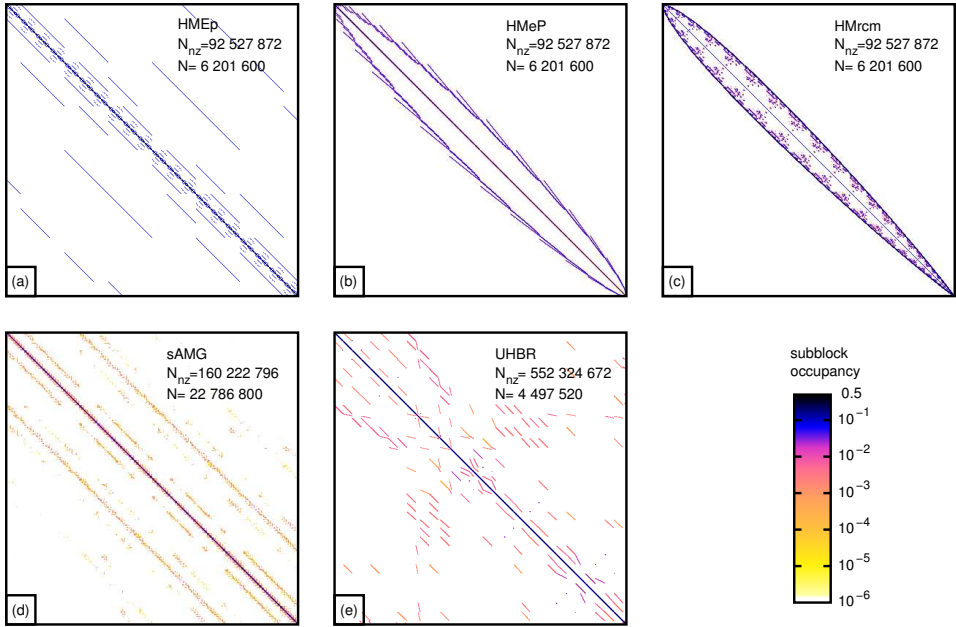


Fig. 1: Sparsity patterns of the matrices described in Sect. 1.3.1. (a)–(c) describe the same physical system, but use a different numbering of the basis elements. See text for details. Square subblocks have been aggregated and color-coded according to occupancy to improve visibility.

1.3. Experimental setting

1.3.1. Test matrices

Since the sparsity pattern may have substantial impact on the single node performance and parallel scalability, we have chosen three application areas known to generate extremely sparse matrices.

As a first test case we use a matrix from exact diagonalization of strongly correlated electron-phonon systems in solid state physics. Here generic microscopic models are used to treat both charge (electrons) and lattice (phonons) degrees of freedom in second quantization. Choosing a finite-dimensional basis set, which is the direct product of basis sets for both subsystems (electrons \otimes phonons), the generic model can be represented by a sparse Hamiltonian matrix. Iterative algorithms such as Lanczos or Jacobi-Davidson are used to compute low-lying eigenstates of the Hamilton matrices, and more recent methods based on polynomial expansion allow for computation of spectral properties [10] or time evolution of quantum states [11]. In all those algorithms, spMVM is the most time-consuming step.

In this paper we consider the Holstein-Hubbard model (cf. [12] and references therein) and choose six electrons (subspace dimension 400) on a six-site lattice coupled to 15 phonons (subspace dimension 1.55×10^4). The resulting matrix of

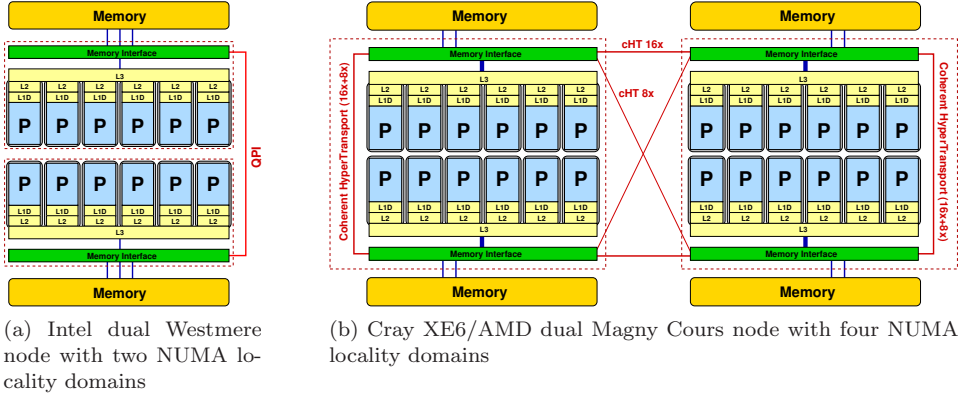


Fig. 2: Node topology of the benchmark systems. Dashed boxes indicate sockets.

dimension 6.2×10^6 is very sparse ($N_{\text{nzt}} \approx 15$) and can have two different sparsity patterns, depending on whether the phononic or the electronic basis elements are numbered contiguously (see Figs. 1 (a) and (b), respectively). We also applied the well-known “Reverse Cuthill-McKee (RCM)” algorithm [13] to the Hamilton matrix in order to improve spatial locality in the access to the right hand side vector, and to optimize interprocess communication patterns towards near-neighbor exchange. Since the RCM-optimized structure (Fig. 1 (c)) showed no performance advantage over the HMeP variant (Fig. 1 (b)) neither on the node nor on the highly parallel level, we will not consider RCM any further in the following.

The second matrix was generated by the adaptive multigrid code sAMG (see [14, 15], and references therein) for the irregular discretization of a Poisson problem on a car geometry. Its matrix dimension is 2.2×10^7 with an average of $N_{\text{nzt}} \approx 7$ entries per row (see Fig. 1 (d)).

The UHBR matrix (see Fig. 1 (e)) originates from aeroelastic stability investigations of an ultra-high bypass ratio (UHBR) turbine fan of the German Aerospace Center (DLR) with a linearized Navier-Stokes solver [16]. This solver is part of the parallel simulation system TRACE (Turbo-machinery Research Aerodynamic Computational Environment) which was developed by DLR’s Institute for Propulsion Technology. Its matrix dimension is 4.5×10^6 with an average of $N_{\text{nzt}} \approx 123$ entries per row, making it a rather “densely populated” sparse matrix in comparison to the other test cases.

For symmetric matrices as considered here it would be sufficient to store the upper triangular matrix elements and perform, e.g., a parallel symmetric CRS sp-MVM [4]. The data transfer volume is then reduced by almost a factor of two, allowing for a corresponding performance improvement. We do not use this optimization here for two major reasons. First, the discussion of the hybrid parallel vs. MPI-only implementation should not be restricted to the special case of explic-

itly symmetric matrices. Second, an efficient shared-memory implementation of a symmetric CRS spMVM base routine has been presented only very recently [17].

1.3.2. *Test machines*

Intel Nehalem EP / Westmere EP The two Intel platforms represent a “tick” step within Intel’s “tick-tock” product strategy. Both processors only differ in a few microarchitectural details; the most important difference is that Westmere, due to the 32 nm production process, accommodates six cores per socket instead of four while keeping the same L3 cache size per core (2 MB) as Nehalem. The processor chips (Xeon X5550 and X5650) used for the benchmarks run at 2.66 GHz base frequency with “Turbo Mode” and Simultaneous Multithreading (SMT) enabled. A single socket forms its own ccNUMA LD via three DDR3-1333 memory channels (see Fig. 2 (a)), allowing for a peak bandwidth of 32 GB/s. We use standard dual-socket nodes that are connected via fully nonblocking QDR InfiniBand (IB) networks. The Intel compiler in version 11.1 and the Intel MPI library in version 4.0.1 were used throughout. Thread-core affinity was controlled with the LIKWID [18] toolkit.

Cray XE6 / AMD Magny Cours The Cray XE6 system is based on dual-socket nodes with AMD Magny Cours 12-core processors (2.1 GHz Opteron 6172) and the latest Cray “Gemini” interconnect. The internode bandwidth of the 2D torus network is beyond the capability of QDR InfiniBand. The single node architecture depicted in Fig. 2(b) reveals a unique feature of the AMD Magny Cours chip series: The 12-core package comprises two 6-core chips with separate L3 caches and memory controllers, tightly bound by “1.5” HyperTransport (HT) 16x links. Each 6-core unit forms its own NUMA LD via two DDR3-1333 channels, i.e., a two-socket node comprises four NUMA locality domains. In total the AMD design uses eight memory channels, allowing for a theoretical main memory bandwidth advantage of 8/6 over a Westmere node. The Cray compiler in version 7.2.8 was used for the Cray/AMD measurements.

In Sect. 3.1 we also show MPI performance results for an older Cray XT4 system based on AMD Opteron “Barcelona” processors.

2. Node-level performance analysis

The basis for each parallel program must be an efficient single core/node implementation. Assuming general sparse matrix structures, the CRS format presented above is very suitable for modern cache-based multicore processors [19]. Even advanced machine-specific optimizations such as nontemporal prefetch instructions for Opteron processors provide only minor benefits [4] and are thus not considered here. A simple OpenMP parallelization of the outermost loop, together with an appropriate NUMA-aware data placement strategy has proven to provide best node-level performance. We choose the HMeP, HMEp, and UHBR matrices as reference cases for our performance model.

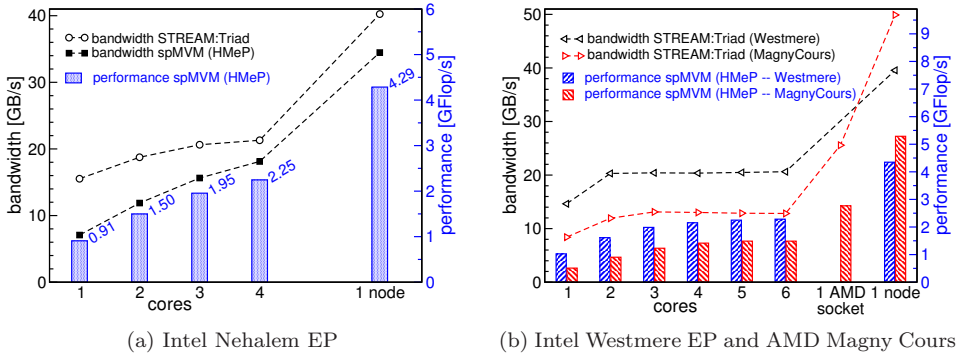


Fig. 3: Node-level performance for the test systems. Effective STREAM triads bandwidth^a, and performance for spMVM using the HMeP matrix (bars) is shown. In (a) we also report the measured memory bandwidth for the spMVM operation.

Intrsocket and intranode spMVM scalability should always be discussed together with effective STREAM triad numbers, which form a practical upper bandwidth limit.^a Figure 3(a) shows the memory bandwidth on the Nehalem EP platform drawn by the STREAM triad and the spMVM as measured with LIKWID [18]. While the STREAM bandwidth soon saturates within a socket, the spMVM bandwidth and the corresponding GFlop/s numbers still benefit from the use of all cores. This is a typical behavior for codes with (partially) irregular data access patterns. However, the fact that more than 85% of the STREAM bandwidth can be reached with spMVM indicates that our CRS implementation makes good use of the resources. The maximum spMVM performance can be estimated by dividing the memory bandwidth by the code balance (1), using $N_{nzr} = 15$ and $\kappa = 0$. For a single socket the spMVM draws 18.1 GB/s (STREAM triads: 21.2 GB/s), allowing for a maximum performance of 2.66 GFlop/s (3.12 GFlop/s). Combining the measured performance (2.25 GFlop/s) and bandwidth of the spMVM operation with $B_{CRS}(\kappa)$ we find $\kappa = 2.5$, i.e., 2.5 additional bytes of memory traffic on $B(\cdot)$ per inner loop iteration (37.3 bytes per row) are required due to limited cache capacity. Thus the complete vector $B(\cdot)$ is loaded six times from main memory to cache, but each element is used $N_{nzr} = 15$ times on average. This ratio gets worse if the matrix bandwidth increases. For the HMeP matrix we found $\kappa = 3.79$, which translates to a 50% increase in the additional data transfers for $B(\cdot)$. The code balance implies a performance drop of about 10%, which is consistent with our measurements.

The UHBR matrix represents an interesting case, since the average number of nonzeros per row is $N_{nzr} \approx 123$. At a measured spMVM bandwidth of 18.9 GB/s and a performance of 2.99 GFlop/s per socket we arrive at $\kappa = 0.43$, which means

^aNontemporal stores have been suppressed in the STREAM measurements and the bandwidth numbers reported have been scaled appropriately ($\times 4/3$) to account for the write-allocate transfer.

that each element of the RHS is loaded 8 times; however, it is used 123 times, which leads to the conclusion that the contribution of the RHS to the memory traffic is minor. We have included this example here because it shows that the data transfer for the RHS may be negligible even if it is loaded many times — N_{nzt} plays a decisive role. Nevertheless, since this matrix shows perfect scaling in the highly parallel case we will not discuss it any further in this work.

In Fig. 3 (b) we summarize the performance characteristics for Intel Westmere and AMD Magny Cours, which both comprise six cores per locality domain. While the AMD system is slower on a single LD, its node-level performance is about 25% higher than on Westmere due to its four LDs per node. Within the domains spMVM saturates at four cores on both architectures, leaving ample room to use the remaining cores for other tasks, like communication (see Sect. 3.5). In the following we will report results for the Westmere and Magny Cours platforms only.

3. Distributed-memory parallelization

3.1. *Nonblocking point-to-point communication in MPI*

Strong scaling of MPI-parallel spMVM is inevitably limited by communication overhead. Hence, it is vital to find ways to hide communication costs as far as possible. A widely used approach is to employ nonblocking point-to-point MPI calls for overlapping communication with useful work. However, it has been known for a long time that most MPI implementations support progress, i.e., actual data transfer, only when MPI library code is executed by the user process, although the hardware even on standard InfiniBand-based clusters does not hinder truly asynchronous point-to-point communication.^b

Very simple benchmark tests can be used to find out whether the nonblocking point-to-point communication calls in an MPI library do actually support truly asynchronous transfers. Listing 2 (adapted from [9]) shows an example where an `MPI_Irecv()` operation is set off before a function (`do_work()`) performs register-only operations for a configurable amount of time. If the nonblocking message transfer overlaps with computations, the overall runtime of the code will be constant as long as the time for computation is smaller than the time for message transfer. We have used a constant large message length of 80 MB to get accurate measurements. Note that the results of such tests may depend crucially on tunable parameters like, e.g., the message size for the cross-over from an “eager” to a “rendezvous” protocol, especially for small messages. For the application scenarios described later, most messages are still beyond such limits. Figure 4 shows overall runtime versus time for computation on the Intel Westmere cluster and the Cray XT4/XE6 systems, respectively. We only report internode results, since no current MPI implementation on any system supports asynchronous nonblocking intranode communication.

^bIn fact, dedicated “offload” communication hardware was not unusual in historic supercomputer architectures, the Intel Paragon of the early 1990s being a typical example.

Listing 2: A simple benchmark to determine the capability of the MPI library to perform asynchronous nonblocking point-to-point communication for large messages (receive variant).

```

1  if(rank==0) {
2      stime = MPI_Wtime();
3      MPI_Irecv(rbuf,mcount,MPI_DOUBLE,1,0,MPI_COMM_WORLD,&req);
4      do_work(calctime);
5      MPI_Wait(req, &status);
6      etime = MPI_Wtime();
7      cout << calctime << " " << etime-stime << endl;
8  } else MPI_Send(sbuf,mcount,MPI_DOUBLE,0,0,MPI_COMM_WORLD);

```

On the Intel cluster we compared three different MPI implementations: Intel MPI, OpenMPI, and MVAPICH2. The latter was compiled with the `--enable-async-progress` flag. OpenMPI 1.5 supports a similar setting, but it is documented to be still under development in the current version (1.5.3), and we were not able to produce a stable configuration with progress threads activated. The results show that only OpenMPI (even without progress threads explicitly enabled) was capable of asynchronous nonblocking communication, albeit only when sending data via a nonblocking send. The nonblocking receive is not asynchronous, however.

Comparing the Cray XT4 and XE6 systems, it is striking that only the older XT4 has an MPI implementation that supports asynchronous nonblocking transfers for large messages.

In summary, one must conclude that the naive assumption that “nonblocking” and “asynchronous” are the same thing cannot be upheld for most current MPI implementations; as a consequence, overlapping computation with communication is often a matter of explicit programming.

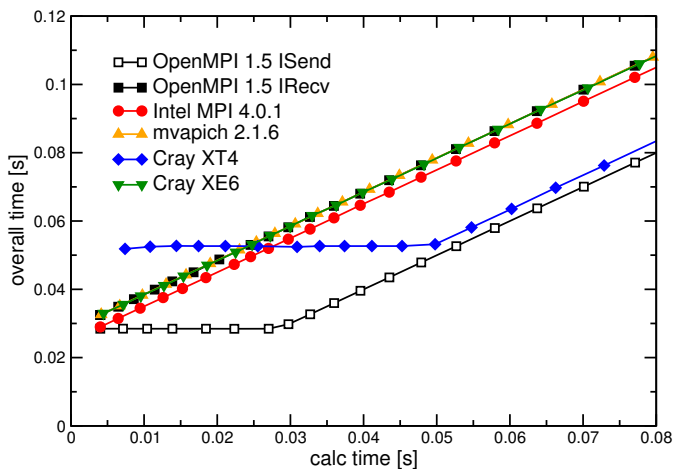


Fig. 4: Internode results for the nonblocking MPI benchmark on the Westmere-based test cluster and on Cray XT4 and XE6 systems. Unless indicated otherwise, results for nonblocking send and receive are almost identical.

3.2. *MPI-parallel sparse MVM*

In the following sections we will contrast the “naive” overlap via nonblocking MPI with an approach that uses a dedicated OpenMP thread for explicitly asynchronous transfers. We adopt the nomenclature from [6] and [9] and distinguish between “vector mode” and “task mode.”

MPI parallelization of spMVM is generally done by distributing the nonzeros (or, alternatively, the matrix rows), the right hand side vector $\mathbf{B}(\cdot)$, and the result vector $\mathbf{C}(\cdot)$ evenly across MPI processes. Due to off-diagonal nonzeros, every process requires some parts of the RHS vector from other processes to complete its own chunk of the result, and must send parts of its own RHS chunk to others. Note that it is generally difficult to establish good load balancing for computation and communication at the same time. Unless indicated otherwise we use a balanced distribution of nonzeros across the MPI processes here. At a given number of processes, the resulting communication pattern depends only on the sparsity structure, so the necessary bookkeeping needs to be done only once. The actual spMVM computations can be performed either by a single thread or, if threading is available, by multiple threads inside the MPI process.

3.3. *Vector-like parallelization: Vector mode without overlap*

Gathering the data to be sent by a process into a contiguous send buffer may be done after the receive has been initiated, potentially hiding the cost of copying (see Fig. 5 (a)). We call this naive approach “hybrid vector mode,” since it strongly resembles the programming model for vector-parallel computers [6]: The time-consuming (although probably parallel) computation step does not overlap with communication overhead. This is actually how “MPI+OpenMP hybrid programming” is still defined in most publications. The question whether and why using multiple threads per MPI process may improve performance compared to a pure MPI version on the same hardware is not easy to answer. Depending on the problem, different aspects come into play, and there is no general rule [20].

3.4. *Vector-like parallelization: Vector mode with naive overlap*

As an alternative one may consider hybrid vector mode with nonblocking MPI (see Fig. 5 (b)) to potentially overlap communication with the part of spMVM that can be completed using local RHS elements only. After the nonlocal elements have been received, the remaining spMVM operations can be performed. A disadvantage of splitting the spMVM in two parts is that the local result vector must be written twice, incurring additional memory traffic. The performance model (1) can be modified to account for an additional data transfer of $16/N_{\text{nzr}}$ bytes per inner loop iteration, leading to a modified code balance of

$$B_{\text{CRS}}^{\text{split}} = \left(6 + \frac{20}{N_{\text{nzr}}} + \frac{\kappa}{2} \right) \frac{\text{bytes}}{\text{flop}} . \quad (2)$$

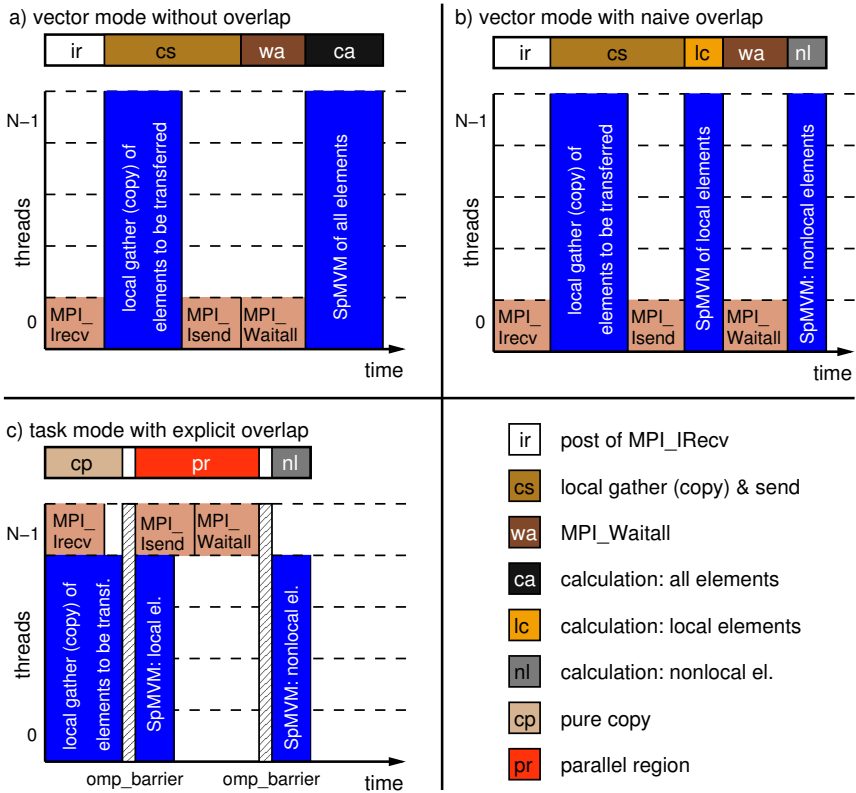


Fig. 5: Schematic timeline view of the implemented hybrid kernel versions: (a) no communication/calculation overlap, (b) naive overlap using nonblocking MPI, and (c) explicit overlap by a dedicated communication thread. The abbreviations in the top bars indicate the individual contributions that will be discussed in the following sections.

For $N_{\text{nzr}} \approx 7 \dots 15$ and assuming $\kappa = 0$, one may expect a node-level performance penalty between 15% and 8%, and even less if $\kappa > 0$.

For simplicity we will also use the term “vector mode” for pure MPI versions with single-threaded computation.

3.5. Task mode with explicit overlap

A safe way to ensure overlap of communication with computation is to use a separate communication thread and leave the computational loops to the remaining threads. We call this “hybrid task mode,” because it establishes a functional decomposition of tasks (communication vs. computation) across the resources (see Fig. 5 (c)): One thread executes MPI calls only, while all others are used to copy data into send buffers, perform the spMVM with the local RHS elements, and finally (after all

communication has finished) do the remaining spMVM parts. Since spMVM saturates at about 3–5 threads per locality domain (as shown in Fig. 3 (b)), at least one core per LD is available for communication without adversely affecting node-level performance. On architectures with SMT, like the Intel Westmere, there is also the option of using one compute thread per physical core and bind the communication thread to a logical core. Note that, even with perfect overlap, one may expect the speedup compared to any vector mode to be always less than a factor of two.

Apart from the additional memory traffic due to writing the result vector twice (see Sect. 3.4), another drawback of hybrid task mode is that the standard OpenMP loop worksharing directive cannot be used, since there is no concept of “subteams” in the current OpenMP standard. Work distribution is thus implemented explicitly, using one contiguous chunk of nonzeros per compute thread.

4. Internode performance results and discussion

In this section we present strong scaling results for the Holstein-Hubbard (both basis numberings) and sAMG matrices. Besides a discussion of the benefits of hybrid task mode we also provide evidence that hybrid vector mode, even without overlap, may improve performance due to better load balancing.

4.1. Basis ordering for the Holstein-Hubbard matrix

Despite the different sparsity patterns of HMeP and HMEp their node-level performance differs only by roughly 10% (HMEp: 3.89 GFlop/s, HMeP: 4.34 GFlop/s on a Westmere EP node). The question arises whether it is possible to choose an appropriate partitioning of the matrix (or, equivalently, a certain number of processes) so that communication overhead is greatly reduced, and whether the basis ordering plays a relevant role. An analysis of the individual contributions of the parallel spMVM shows the paramount role assumed by the sparsity pattern as soon as communication becomes an issue. In Fig. 6 we show for different numbers of nodes and one MPI process per core (vector modes) or one MPI process per node (task mode) the cost for computing and communicating (time \times number of nodes); each box with whiskers denotes the variation across all processes in the parallel run (10th/90th and 25th/75th percentiles). The broadening of the boxes and whiskers with increasing node count is a consequence of load imbalance; see Sect. 4.2.1 for details about this issue.

4.1.1. Vector modes

The purely computational cost (ca, lc, nl) is roughly on par in both matrices, and scales almost linearly with the number of nodes (approximately horizontal trend for the slowest processes in the cost plot), no matter which variant of vector mode is chosen. In contrast, the cost for communication spent in `MPI_Waitall` grows with the number of processors, which is to be expected since the local matrix blocks

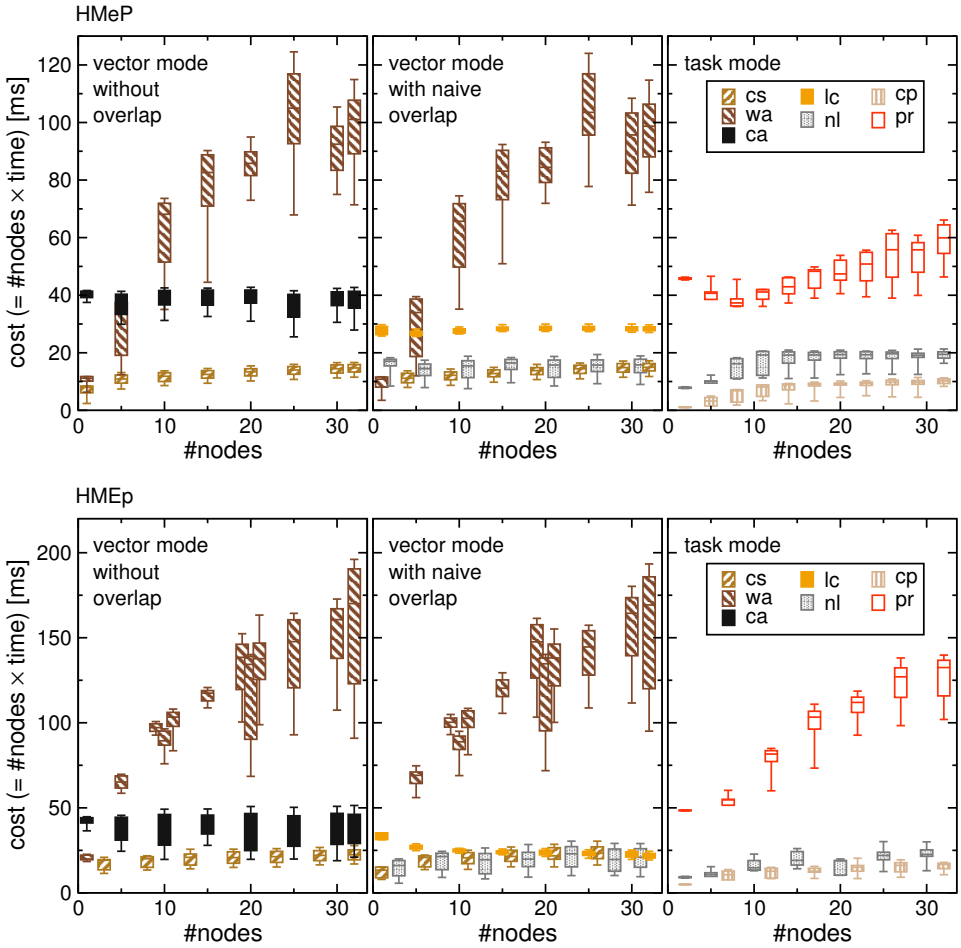


Fig. 6: Cost of the contributions to parallel spMVM vs. number of nodes for HMeP (top row, constant number of nonzeros per process) and HMEp (bottom row, constant number of rows per process) on the Westmere cluster. All physical cores of a node were used, either by MPI processes (vector modes) or additional OpenMP threads (task mode). Results are given in terms of a ‘cost’ function, which is the product of time required and the number of nodes. See Fig. 5 for a description of the abbreviations. Costs for posting the MPI_IRecv are marginal on this scale. For clarity reasons not all node counts are shown.

become smaller. Furthermore, the inhomogeneous matrix structures result in increasing variations of the transferred data volume (and thus communication time). While both matrices show this trend, there are particular node counts for which the communication pattern of HMEp is obviously more favorable (10, 20). At these points the number of cores is commensurable with the diagonal block structure of

the matrix, the majority of the communication happens inside the nodes, and only few large messages are passed between nodes. But even then the communication cost for HMEp is still larger than for HMeP.

4.1.2. *Task mode*

In task mode, which was used here with one MPI process per node, the overall cost is dominated by the parallel region (pr). The reduced number of communicating MPI processes alleviates the load balancing problem in the communication scheme, which will be discussed in the following section. More importantly, up to around 15 nodes the cost for the parallel region is roughly constant in the HMeP case, implying that communication is hidden completely behind computation as discussed in Sect. 3.5. Beyond this point, communication time starts to become dominant at least for some processes, as indicated by the slow rise of the top whisker (90th percentile) for the parallel region. However, we still expect decent performance scaling for this setup. The HMEp matrix, due to its unfavorable communication pattern, does not allow for sufficient overlap.

4.2. *Testcase HMeP*

4.2.1. *Analysis of runtime contributions in the MPI-parallel case*

In order to pinpoint the relevant performance-limiting aspects in the MPI-parallel case we show in Fig. 7 the different contributions to the runtime of a single spMVM for each MPI process when using one of the two vector mode variants, with one process per core and at 12, 96, and 384 processes, respectively. In these graphs the overall runtime is always given by the highest bar; variations in runtime across processes are a sign of load imbalance. Owing to the separation of local from nonlocal spMVM parts, vector mode with naive overlap (right column) is always slower than vector mode without overlap (left column).

The histograms reflect roughly the shape of the sparsity pattern for HMeP as shown in Fig. 1: As the number of processes is increased, “speeders,” i.e., processes that are faster than the rest, start to show up mainly at the top and bottom ends of the matrix (low and high ranks). This imbalance is chiefly caused by a smaller amount of communication (`MPI_Waitall`), whereas the contribution of computation to the runtime is much more balanced across all processes. One may thus expect that a lower number of MPI processes on the same number of cores (i.e., using multiple threads per process) improves performance due to better load balancing even without explicit overlap. At larger process counts, execution time starts to be dominated by communication. Hence, even with multiple threads per MPI process and therefore improved load balancing, explicit overlap of communication with the local part of the spMVM (labeled “lc” in the right group of diagrams in Fig. 7) is expected to show significant speedups.

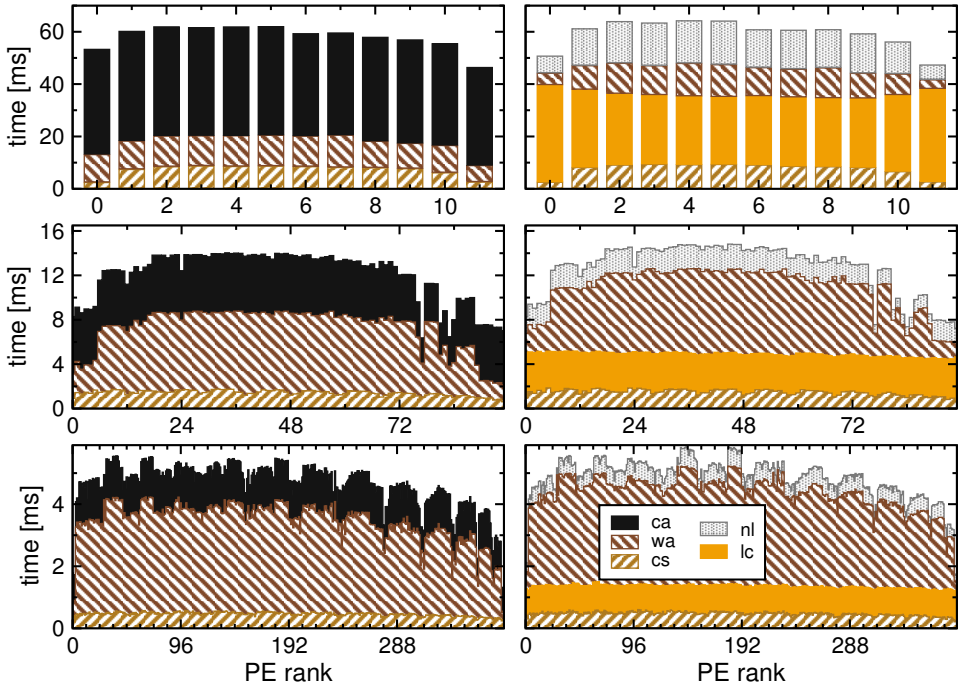


Fig. 7: Contributions to the runtime of one spMVM (HMeP) for each MPI process at different numbers of processes (12, 96, and 384 from top to bottom; strong scaling) on the Westmere cluster. The first column corresponds to vector mode without overlap (see Fig. 5 (a)) whereas the second column presents the runtimes for vector mode with naive overlap (see Fig. 5 (b)). One core per process was used throughout. Abbreviations as in Fig. 5.

4.2.2. Performance results

At one MPI process per physical core (left panel in Fig. 8), vector mode with naive overlap is always slower than the variant without overlap because the additional data transfer on the result vector cannot be compensated by overlapping communication with computation. Task mode was implemented here with one communication thread per MPI process, running on the second virtual core. In this case, point-to-point transfers explicitly overlap with the local spMVM, leading to a noticeable performance boost. One may conclude that MPI libraries with support for progress threads could follow the same strategy and bind those threads to unused logical cores, thereby allowing overlap even with single-threaded user code.

With one MPI process per NUMA locality domain (middle panel) the advantage of task mode is even more pronounced. Also the plain vector mode without overlap shows some notable speedup compared to the MPI-only version, which was expected from the discussion of load balancing in the previous section. Since the memory bus of an LD is already saturated with four threads, it does not make a difference




Fig. 8: Strong scaling performance data for spMVM with the HMeP matrix (constant number of nonzeros per process) on the Intel Westmere cluster for different pure MPI and hybrid variants (kernel version (a) – (c)) as in Fig. 5). Symbols on each data set indicate the 50% parallel efficiency point with respect to the best single-node version. The best variant on the Cray XE6 system is shown for reference (see text for details).

whether six worker threads are used with one communication thread on a virtual core, or whether a complete physical core is devoted to communication. The same is true with only one MPI process (12 threads) per node (right panel).

The symbols in Fig. 8 indicate the 50% parallel efficiency point (with respect to the best single-node performance as reported in Fig. 3 (b)) on each data set. In practice one would not go beyond this number of nodes because of bad resource utilization. For the matrix and the system under investigation it is clear that task mode allows strong scaling to much higher levels of parallelism with acceptable parallel efficiency than any variant of vector mode. However, even the vector mode variants show a significant performance advantage at multiple threads per process due to improved load balancing.

Contrary to expectations based on the single-node performance numbers (Fig. 3 (b)), the Cray XE6 can generally not match the performance of the Westmere cluster at larger node counts, with the exception of pure MPI where both are roughly on par (left panel, Cray results for vector mode with naive overlap). When using threaded MPI processes (middle and right panel), task mode performs best on the Cray system. The advantage over the other kernel variants is by far not as pronounced as on Westmere, however. We have observed a strong influence of job topology and machine load on the communication performance over the 2D torus network. Since spMVM requires significant non-nearest-neighbor communication with growing process counts, the nonblocking fat tree network on the Westmere cluster seems to be better suited for this kind of problem. The presented results are

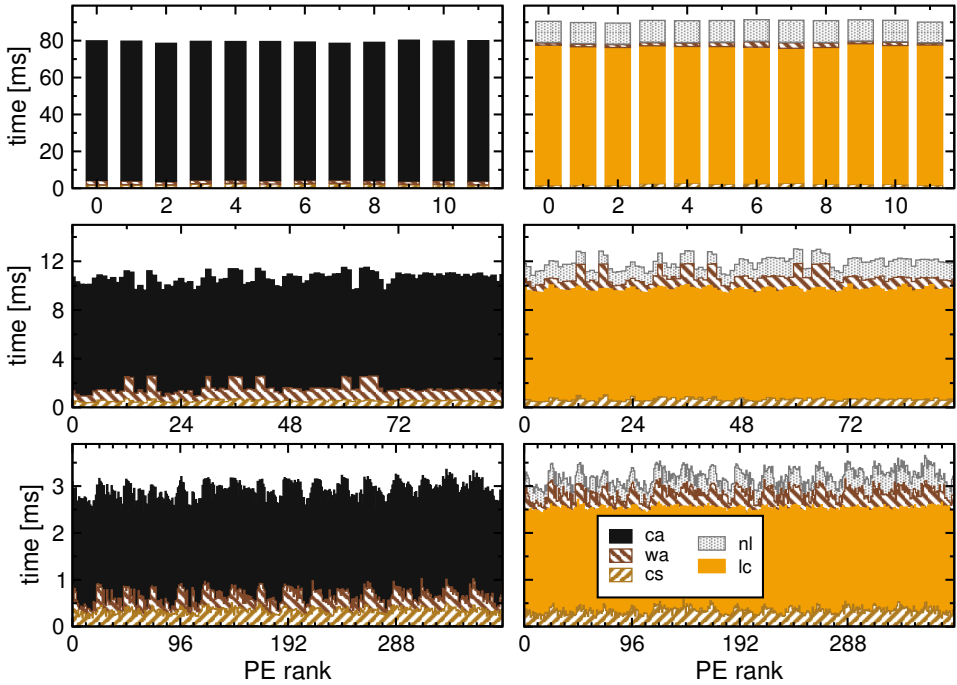


Fig. 9: Contributions to the runtime of one spMVM (sAMG) for each MPI process. Parameters and abbreviations as in Fig. 7.

best values obtained on a dedicated XE6 machine.

There is also a universal drop in scalability beyond about six nodes, which is largely independent of the particular hybrid mode. This can be ascribed to a strong decrease in overall internode communication volume when the number of nodes is small. The effect is somewhat less pronounced for pure MPI, since the overhead of intranode message passing cannot be neglected.

4.3. Testcase sAMG

4.3.1. Analysis of runtime contributions in the MPI-parallel case

As shown in Fig. 9, there is only a slight load imbalance for the sAMG matrix even at 384 processes (lower diagrams). We thus only expect a marginal performance benefit from using threaded MPI processes. Also the overall fraction of communication overhead is quite small; task mode with explicit overlap of communication will hence not lead to a significant speedup.

4.3.2. Performance results

As expected from the analysis in the previous section, all variants and hybrid modes (pure MPI, one process per LD, and one process per node) show similar scaling




Fig. 10: Strong scaling performance data for spMVM with the sAMG matrix (same variants as in Fig. 8). Parallel efficiency is above 50% for all versions up to 32 nodes. The Cray system performed best in vector mode without overlap for all cases.

behavior on the Westmere cluster, and there is no advantage of task mode over vector mode without overlap or over pure MPI (see Fig. 10). This observation supports the general rule that it makes no sense to consider MPI+OpenMP hybrid programming if the pure MPI code already scales well and behaves in accordance with a single-node performance model.

On the Cray XE6, vector mode without overlap performs best across all hybrid modes, with a significant advantage when running one MPI process with six threads per LD. This aspect is still to be investigated.

5. Summary and outlook

We have investigated the performance properties of pure MPI and hybrid MPI+OpenMP hybrid variants of sparse matrix-vector multiplication on two multicore-based parallel systems, using matrices with different sparsity patterns. The single-node performance analysis on Intel Westmere and AMD Magny Cours processors showed that memory-bound sparse MVM saturates the memory bus of a locality domain already at about four threads, leaving free cores for explicit computation/communication overlap. As most current MPI libraries do not support truly asynchronous point-to-point transfers, explicit overlap enabled substantial performance gains for strong scaling of communication-bound problems. Since the communication thread can run on a virtual core, MPI implementations could use the same strategy for internal “progress threads” and so enable asynchronous progress without changes in MPI-only user code.

We have also identified the clear advantage of using threaded MPI processes, even without explicit communication overlap, in cases where computation is well balanced and load imbalance is caused by the communication pattern.

Acknowledgments

We thank J. Treibig, R. Keller and T. Schönemeyer for valuable discussions, A. Basermann for providing and supporting the RCM transformation and the UHBR test case as well as K. Stüben and H. J. Plum for providing and supporting the sAMG test case. We also thank CSCS Manno for granting access to their Cray X6E system. Financial support from KONWIHR II (project HQS@HPC II) and from Deutsche Forschungsgemeinschaft through SFB TR 24 (project A5) is gratefully acknowledged.

References

- [1] G. Goumas, K. Kourtis, N. Anastopoulos, V. Karakasis, N. Koziris: *Performance evaluation of the sparse matrix-vector multiplication on modern architectures*. J. Supercomputing **50**(1), 36–77 (2008).
- [2] S. Williams, L. Oliker, R. Vuduc, J. Shalf, K. Yelick, J. Demmel: *Optimization of sparse matrix-vector multiplications on emerging multicore platforms*. Parallel Computing **35**, 178 (2009)
- [3] N. Bell and M. Garland: *Implementing sparse matrix-vector multiplication on throughput-oriented processors*. Proceedings of SC09.
- [4] M. Krotkiewski and M. Dabrowski: *Parallel symmetric sparse matrix-vector product on scalar multi-core CPUs*. Parallel Computing **36** (4), 181–198 (2010).
- [5] R. Geuss and S. Röllin: *Towards a fast parallel sparse symmetric matrix-vector multiplication*. Parallel Computing **27** (1), 883–896 (2001).
- [6] R. Rabenseifner and G. Wellein: *Communication and Optimization Aspects of Parallel Programming Models on Hybrid Architectures*. International Journal of High Performance Computing Applications **17**, 49–62 (2003).
- [7] G. Wellein, G. Hager, A. Basermann, and H. Fehske: *Fast sparse matrix-vector multiplication for TFlop/s computers*. In: Proceedings of VECPAR2002, LNCS 2565, Springer Berlin (2003).
- [8] R. Barrett *et al.*: *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*. SIAM, ISBN 978-0898713282, (1993).
- [9] G. Hager and G. Wellein: *Introduction to High Performance Computing for Scientists and Engineers*. CRC Press, ISBN 978-1439811924, (2010).
- [10] A. Weiße, G. Wellein, A. Alvermann, and H. Fehske: *The kernel polynomial method*. Rev. Mod. Phys. **78**, 275 (2006).
- [11] A. Weiße and H. Fehske: *Chebyshev expansion techniques*. In: Computational Many Particle Physics, Lecture Notes in Physics 739, pp. 545–577, Springer Berlin Heidelberg (2008).
- [12] H. Fehske, G. Wellein, G. Hager, A. Weiße, and A. R. Bishop: *Quantum lattice dynamical effects on the single-particle excitations in 1D Mott and Peierls insulators*. Phys. Rev. B **69**, 165115 (2004).
- [13] E. Cuthill and J. McKee: *Reducing the bandwidth of sparse symmetric matrices*. Proceedings of 24th national conference (ACM '69), ACM, New York, NY, USA, 157–172.
- [14] K. Stüben: *An Introduction to Algebraic Multigrid*. In: Multigrid: Basics, Parallelism and Adaptivity, Eds. U. Trottenberg *et al.*, Academic Press (2000).
- [15] <http://www.scai.fraunhofer.de/en/business-research-areas/numerical-software/products/samg.html>
- [16] A. Basermann *et al.*: *HICFD - Highly Efficient Implementation of CFD Codes for HPC Many-Core Architectures*. In: Proceedings of CiHPC, Springer 2011 [in print]

- [17] A. Buluc, S.W. Williams, L. Oliker, and J. Demmel: *Reduced-Bandwidth Multi-threaded Algorithms for Sparse-Matrix Vector Multiplication*. Proc. IPDPS 2011 (to appear). <http://gauss.cs.ucsb.edu/~aydin/ipdps2011.pdf>
- [18] <http://code.google.com/p/likwid>
- [19] G. Schubert, G. Hager, and H. Fehske: *Performance limitations for sparse matrix-vector multiplications on current multicore environments*. In: High Performance Computing in Science and Engineering, Garching/Munich 2009, 13–26, Springer Berlin Heidelberg (2010).
- [20] R. Rabenseifner, G. Hager, and G. Jost: *Hybrid MPI/OpenMP Parallel Programming on Clusters of Multi-Core SMP Nodes*. In: Proceedings of PDP 2009.