TECHNISCHE UNIVERSITÄT
BERGAKADEMIE FREIBERG

Die Ressourcenuniversität. Seit 1765.

Fakultät für Mathematik und Informatik

Preprint 2015-02

Axel Klawonn, Martin Lanser, Oliver Rheinbach,
Holger Stengel and Gerhard Wellein

Hybrid MPI/OpenMP parallelization in FETI-DP methods

Axel Klawonn, Martin Lanser, Oliver Rheinbach, Holger Stengel and Gerhard Wellein

Hybrid MPI/OpenMP parallelization in FETI-DP methods

Hybrid MPI/OpenMP parallelization in FETI-DP methods

Axel Klawonn, Martin Lanser, Oliver Rheinbach, Holger Stengel, and Gerhard Wellein

Abstract We present an approach to hybrid MPI/OpenMP parallelization in FETI-DP methods using OpenMP with PETSc in the finite element assembly and using the shared memory parallel direct solver Pardiso in the FETI-DP solution phase. Our approach thus uses OpenMP parallelization on subdomains and MPI in between subdomains. We investigate the efficiency of this approach for a benchmark problem from two dimensional nonlinear hyperelasticity. We observe good scalability for up to 4 threads for each MPI rank on a state-of-the-art Ivy Bridge architecture and incremental improvements for up to 10 OpenMP threads for each MPI rank.

1 Introduction

The solution of large nonlinear and linear problems arising from the finite element discretization of partial differential equations requires fast and highly scalable parallel solvers. Domain decomposition methods [28, 29] are constructed to handle such large, discretized problems in parallel and are well known for their robustness and scalability, especially in structural mechanics and for material science problems.

In domain decomposition methods a geometrical decomposition of the computational domain is performed, and parallelization is based on a divide and conquer strategy. For static and quasi-static problems, an additional small coarse problem is essential for numerical and parallel scalability. In the family of domain decomposition methods, nonoverlapping approaches as FETI-DP (Finite Element Tearing and

Axel Klawonn and Martin Lanser
Mathematisches Institut, Universität zu Köln, Weyertal 86-90, 50931 Köln, Germany

Oliver Rheinbach
Fakultät für Mathematik und Informatik, Institut für Numerische Mathematik und Optimierung, Technische Universität Bergakademie Freiberg, 09596 Freiberg, Germany.

Holger Stengel and Gerhard Wellein
Erlangen Regional Computing Center, University of Erlangen–Nuremberg, Erlangen, Germany

Interconnecting - Dual Primal), first introduced by Farhat et al. [14, 13], have reduced communication compared to overlapping Schwarz methods [28, 29]. A classical FETI-DP implementation was awarded a Gordon Bell prize in in 2002 [9] for the solution of a structural mechanics problem using 3 400 cores. Modified versions, i.e., inexact FETI-DP methods have scaled up to 65 536 BlueGene/P cores for elasticity problems already in 2009 [25, 22].

In pure MPI parallel implementations of FETI- or Neumann-Neumann-type iterative substructuring methods, one or more subdomains are assigned to one MPI rank and the local calculations are performed sequentially. We will introduce a second level of parallelism by using shared memory parallelism on the subdomains. This approach was also taken, e.g., for a FETI-1 method in [16, 15] based on a different shared memory direct solver [17, 15].

The robustness of FETI-DP methods partially stems from the use of sparse direct solvers as building blocks. In the past, we have used, e.g., the sequential sparse solver package UMFPACK [12], known for its robustness, as well as the solver package MUMPS [3]. In this work, we apply the more recent shared memory parallel sparse solver PARDISO [27]. Additionally, we introduce an OpenMP parallel assembly of the local finite element problems. In this paper, we present numerical results and a detailed performance analysis of our hybrid FETI-DP implementation.

In this work, we will not consider large core counts as in [19], where weak scalability results for our current FETI-DP implementation for up to 262K cores have recently been presented. Instead, a fixed number of subdomains and MPI ranks, e.g., 4 or 64 subdomains and ranks, is considered. We then spawn different numbers of threads for this configuration. The goal of this paper is to investigate the efficiency of this approach on a state-of-the-art architecture. In the future, we will then apply our hybrid MPI/OpenMP approach for a large number of subdomains and MPI ranks on supercomputers where pure MPI implementations do not scale well enough.

2 Description of the FETI-DP Method

The FETI-DP (Finite Element Tearing and Interconnecting - Dual Primal) method is a domain decomposition method based on the geometric decomposition of a computational domain $\Omega \subset \mathbb{R}^d, d = 2, 3$, into N nonoverlapping subdomains Ω_i. Instead of solving a single and largeproblem $Ku = f$ arising from a finite element discretization on the domain Ω, several smaller problems $K^{(i)}u^{(i)} = f^{(i)}$, associated to the subdomains Ω_i, are solved in parallel. Here, the locally assembled stiffness matrices are denoted by $K^{(i)}$ and the local load vectors by $f^{(i)}$. A combination of a global subassembly in a few variables and dual conditions ensures the continuity of the solution on the interface $\Gamma := \bigcup_{i=0}^{N} \partial \Omega_i \setminus \partial \Omega$: We partition the degrees of freedom on Γ into sets of primal and dual variables, denoted by Π and Δ and also define the set of inner and dual variables $B := [I, \Delta]$. The assembly in some primal variables Π is performed by means of the standard FETI-DP partial assembly operator R_Π^T; see,

e.g., [22, 29]. The assembled system in the primal variables is also named FETI-DP coarse problem and can include more than just vertex constraints; c.f. [23]. To obtain continuity in the remaining interface variables we introduce the jump operator B_B, c.f. the definition in [22, 29], and Lagrange multipliers λ to enforce the continuity condition $B_B u_B = 0$ on the variables u_B. This leads to the FETI-DP master system

$$
\begin{bmatrix}
K_{BB} & \widetilde{K}_{\Pi B}^T & B_B^T \\
\widetilde{K}_{\Pi B} & \widetilde{K}_{\Pi\Pi} & 0 \\
B_B & 0 & 0
\end{bmatrix}
\begin{bmatrix}
u_B \\
\tilde{u}_\Pi \\
\lambda
\end{bmatrix}
=
\begin{bmatrix}
f_B \\
\tilde{f}_\Pi \\
0
\end{bmatrix}.
\tag{1}
$$

Here, K_{BB} is a block diagonal matrix and the blocks $K_{BB}^{(i)}$ are local to the subdomains Ω_i, while the matrices including primal variables Π are global but small. Elimination of u_B and $\tilde{u}_\Pi$ leads to

$$
F\lambda = d \tag{2}
$$

$$
\text{where} \qquad F = B_B K_{BB}^{-1} B_B^T + B_B K_{BB}^{-1} \widetilde{K}_{\Pi B}^T \widetilde{S}_{\Pi\Pi}^{-1} \widetilde{K}_{\Pi B} K_{BB}^{-1} B_B^T, \tag{3}
$$

$$
d = B_B K_{BB}^{-1} f_B + B_B K_{BB}^{-1} \widetilde{K}_{\Pi B}^T \widetilde{S}_{\Pi\Pi}^{-1} (\tilde{f}_\Pi - \widetilde{K}_{\Pi B} K_{BB}^{-1} f_B). \tag{4}
$$

Here, $\widetilde{S}_{\Pi\Pi} := \widetilde{K}_{\Pi\Pi} - \widetilde{K}_{\Pi B} K_{BB}^{-1} \widetilde{K}_{\Pi B}^T$ is the Schur complement on the primal variables. Finally, the FETI-DP method is the iterative solution of the preconditioned system

$$
M^{-1} F\lambda = M^{-1} d \tag{5}
$$

with a Krylov subspace method as CG or GMRES. In this paper, we always use the standard Dirichlet preconditioner $M_{FETI_D}^{-1} =: M^{-1}$ which is a weighted sum of local Schur complements $S^{(i)} := K_{\Delta\Delta}^{(i)} - K_{\Delta I}^{(i)} K_{II}^{(i)-1} K_{\Delta I}^{(i)T}$; see, e.g., [22, 29] for complete notation.

3 Classical MPI Parallelism in FETI-DP Methods

MPI-parallel FETI-DP implementations usually handle one or more subdomains per MPI rank. For simplicity, let us assume that we always assign exactly one subdomain to each MPI rank. In this case, the assembly of local problems and the LU factorizations are naturally parallelizable and communication is only required during the assembly process of the coarse problem $\widetilde{S}_{\Pi\Pi}$ and in each application of the jump operator B_B or its transpose.

In general, we can split the computations in three phases: *Local Assembly*, *FETI-DP Setup*, and *Krylov Iteration*. The first phase contains the local assembly of the matrices $K^{(i)}$ and the local right-hand sides $f^{(i)}$. This process is naturally parallelizable, since no communication between the subdomains is needed.

The second phase *FETI-DP Setup* is dominated by LU factorizations of subdomain matrices $K_{BB}^{(i)}$ and $K_{II}^{(i)}$. Here, sparse direct solver packages such as UMF-PACK [10, 11], MUMPS [1, 2] or PARDISO [24, 26, 27] are used. Once factorized,
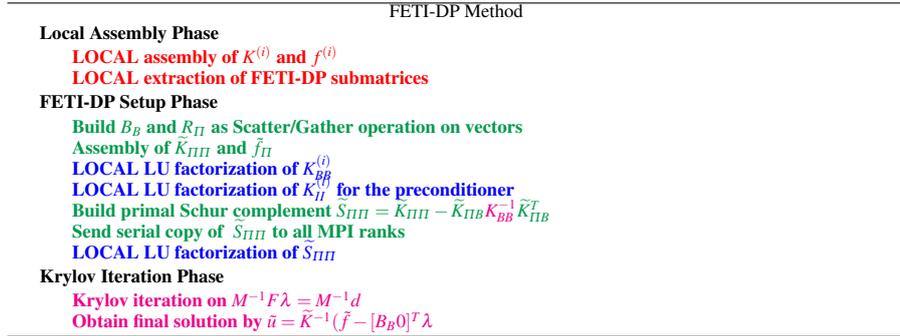

FETI-DP Method

Local Assembly Phase
    LOCAL assembly of $K^{(i)}$ and $f^{(i)}$
    LOCAL extraction of FETI-DP submatrices
FETI-DP Setup Phase
    Build B_B and R_Π as Scatter/Gather operation on vectors
    Assembly of $\widetilde{K}_{\Pi\Pi}$ and $\tilde{f}_\Pi$
    LOCAL LU factorization of $K_{BB}^{(i)}$
    LOCAL LU factorization of $K_{II}^{(i)}$ for the preconditioner
    Build primal Schur complement $\widetilde{S}_{\Pi\Pi} = \widetilde{K}_{\Pi\Pi} - \widetilde{K}_{\Pi B} K_{BB}^{-1} \widetilde{K}_{\Pi B}^T$
    Send serial copy of $\widetilde{S}_{\Pi\Pi}$ to all MPI ranks
    LOCAL LU factorization of $\widetilde{S}_{\Pi\Pi}$
Krylov Iteration Phase
    Krylov iteration on $M^{-1}F\lambda = M^{-1}d$
    Obtain final solution by $\tilde{u} = \widetilde{K}^{-1}(\tilde{f} - [B_B 0]^T \lambda$

Fig. 1 Structure of a FETI-DP method. The assembly phase (red), LU factorizations (blue), LU forward/backward substitutions (magenta). Parts with significant MPI communication are marked in green. Some MPI communication can be hidden behind local factorizations.

each application of K_{BB}^{-1} and $K_{II}^{(i)-1}$ requires only two forward-backward substitutions. The factorizations are completely parallel since all involved matrices are local to an MPI rank. The coarse operator $\widetilde{S}_{\Pi\Pi}$ is also assembled in the *FETI-DP Setup* phase. The assembly can be performed in parallel using a MPI parallel matrix structure. Here, of course, MPI communication is needed. Many implementations are then possible for the solution of the coarse problem. Since the FETI-DP coarse problem is not in the focus of this paper we here only discuss the simplest choice: $\widetilde{S}_{\Pi\Pi}$ is copied to all MPI ranks, and a redundant factorization of $\widetilde{S}_{\Pi\Pi}$ is performed on all ranks; refer, e.g., to [25, 22] for other possibilities.

The third phase is the Krylov subspace iteration on the preconditioned linear system $M^{-1}F\lambda = M^{-1}d$. Each iteration includes several matrix vector multiplications but the dominating parts are the forward backward substitutions caused by the applications of K_{BB}^{-1}, $K_{II}^{(i)-1}$ and $\widetilde{S}_{\Pi\Pi}^{-1}$; c.f., (2), (3), and (4).

The three FETI-DP phases are illustrated in Figure 1.

4 Shared Memory Parallelism in FETI-DP Methods

Modern supercomputer architectures have more and more cores per node while the amount of available memory per core even tends to decrease. To utilize all cores using a pure MPI implementation, we assign at least one FETI-DP subdomain (and one MPI rank) to each core. Sometimes even an overcommit can be beneficial, i.e., using multiple MPI ranks per core, e.g., on BlueGene/Q; see [19]. As a result, available memory per rank can be scarce. This can enforce a decomposition into small subdomains, which may not be optimal for scalability and performance.

A shared memory parallelization of the calculations on a single subdomain allows the decomposition in fewer subdomains than available cores. Larger subdomains can then be chosen and the FETI-DP coarse problem can be kept small. In an

extreme scenario only one subdomain and one MPI rank could be assigned to one complete compute node, although judging from the results in this paper, this may not be the most efficient choice.

To determine the hotspots of the FETI-DP algorithm and to decide which parts should be thread-parallel, we present a preliminary runtime profile of a current MPI parallel FETI-DP implementation. We have performed certain simplifications and optimizations of the assembly process, which are described in more detail in Section 7.1. These modifications are possible in the setting which we discuss here.

After these improvements, the *FETI-DP Setup* phase takes 62%, the assembly process 16%, the assembly of the right-hand side 6%, and the Krylov phase the remaining 16% of the average runtime of the Newton steps, c.f. Figure 2 (left). The *FETI-DP Setup* phase is strongly dominated by the local LU factorizations, which make more than 80% of the *FETI-DP Setup* runtime, c.f. Figure 2 (right). All in all, the assembly and the direct solver package add up to almost 80% of the FETI-DP runtime. Therefore a shared memory parallelization of these two phases has priority.

In the remainder of this paper, we discuss the optimization and OpenMP parallelization of the assembly process, which is based on a parallelization of the loop over the finite elements and also investigate the use of the shared memory parallel direct solvers PARDISO.

Let us summarize our approach: We use MPI parallelism in between the subdomains, which is well-established, and we provide an additional level of shared memory parallelism on the subdomains. In an extreme case, FETI-DP with one subdomain per node is possible.
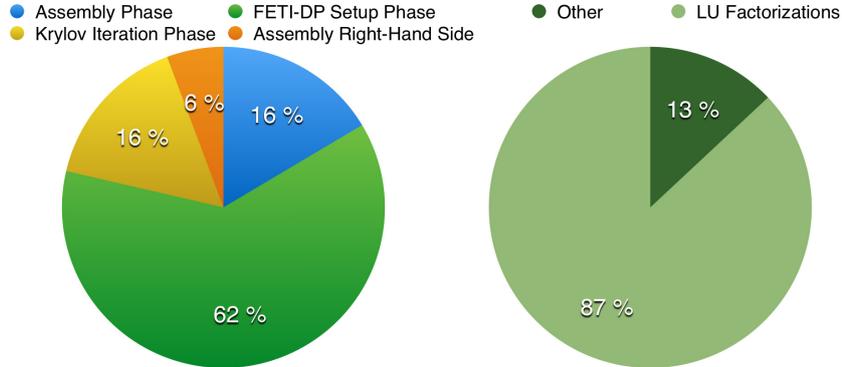

Fig. 2 Left: Runtime profile of the FETI-DP method for a Neo-Hookean hyperelasticity problem in 2D with $m_x = m_y = 200$, using the optimized assembly; see Section 7.1. Right: Distribution of the runtime in the *FETI-DP Setup* phase.

5 Hybrid Parallelization in Inexact and Nonlinear FETI-DP Approaches

In this paper, we consider nonlinear hyperelasticity as a model problem and follow a standard Newton-Krylov-FETI-DP approach for its solution. Moreover, since the number of subdomains is moderate, we always use exact FETI-DP methods [29, 14, 13], i.e., we can afford to solve the FETI-DP coarse problem exactly by a sparse direct solver. The findings presented in this paper on hybrid MPI/OpenMP parallelization are, nevertheless, also valid for inexact FETI-DP methods [21, 22, 20], which remain efficient also for a large number of subdomains. Since the families of the recently introduced nonlinear FETI-DP methods [21, 22, 20] use the same algorithmic building blocks as standard exact or inexact FETI-DP methods our results also apply to these approaches, and we will make use of the same hybrid parallelization in these methods, in the near future.

6 Experimental Setting

6.1 Model Problem

In all experiments in this paper, we consider nonlinear two-dimensional hyperelasticity problems. We consider a Neo-Hookean material with a soft matrix material and stiff circular inclusions; see Figure 3 for the geometry.

The strain energy density function W is given by [31]

$$W(u) = \frac{1}{2}(K - \frac{2}{3}G)\ln^2(\det F) + \frac{1}{2}G[\operatorname{tr}(F^T F) - 3 - 2\ln(\det F)]$$

$$\text{with} \qquad K = \frac{E}{3(1-2\nu)}, \; G = \frac{E}{2(1+\nu)}$$

and the deformation gradient $F(x) := \nabla \varphi(x)$; here, $\varphi(x) = x + u(x)$ denotes the deformation and $u(x)$ the displacement of x. The energy functional of which stationary points are computed, is given by

$$J(u) = \int_\Omega W(u) - F(u)dx - \int_\Gamma G(u)ds,$$

where $F(u)$ and $G(u)$ are functionals related to the volume and traction forces. In our experiments in 2D, we choose the following material parameters E and ν, see Figure 3 for the geometry: In the circular inclusions we have $E = 210\,000$ and in the surrounding matrix material $E = 210$. We have chosen $\nu = 0.3$ in the complete domain Ω. We use a C++ implementation of a Neo-Hooke material and use PETSc [4, 5, 7].
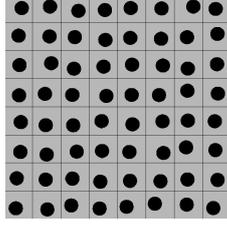
Fig. 3 Decomposition of the computational domain Ω into 64 subdomains; each subdomain has a (slightly randomly off-centered) circular inclusion of stiffer material.

The nonlinear elasticity problem is discretized with piecewise quadratic finite elements. Each square subdomain is discretized by $2 \times m_x \times m_y$ quadratic finite elements. This corresponds to $2(2m_x + 1)(2m_y + 1)$ degrees of freedom per subdomain. We have tested subdomain sizes of $m_x = m_y = 100$, $m_x = m_y = 200$, and $m_x = m_y = 300$ to cover a reasonable range relevant in real-life applications. But as performance results are qualitatively identical for those subdomain sizes, we consistently present performance values only for the subdomain size $m_x = m_y = 200$ to increase comparability.

We will first consider 4 subdomains and thus 4 MPI ranks, while the subdomain size is set to $m_x = m_y = 200$. If not noted otherwise, each of the 4 MPI processes is placed on a separate node. Appropriate measures were taken to pin processors and threads to the cores in a controlled way. The following parameters where used for the PARDISO solver:

```
-mat_pardiso_65 ${OMP_NUM_THREADS}    -mat_pardiso_68 0  -mat_pardiso_1 1
-mat_pardiso_24 1  -mat_pardiso_69 11  -mat_pardiso_2  3  -mat_pardiso_11 0
```

6.2 Hardware, Compiler, and Compiler Flags

All measurements for this report have been executed on a QDR InfiniBand cluster at RRZE. One node has two sockets, each equipped with an Intel Xeon 2660v2 "Ivy Bridge" chip (10 cores + SMT). To achieve consistent performance results, automatic frequency adjustment during socket scaling runs was prevented by fixing the processor frequency to the nominal core clock speed of 2.2 GHz. Each node has 64 GB of RAM (DDR3-1600), the available socket memory bandwidth is 42 GB/s.

We use the Intel Compiler 13.1.3.192 with optimization flags `-O3 -xAVX` and Intel MKL 11.0 Update 5. To use OpenMP within the application (additional flag `-openmp`), PETSc was built using the thread-aware version of the MPI library (Intel MPI flag `-mt_mpi`), which had no impact on the serial performance. For pinning of threads to cores and for performance analysis we use LIKWID [30] version 3.1.1.

7 Numerical Results

7.1 OpenMP Parallel Assembly

In a first quick runtime profile we observe that the assembly of the local stiffness matrix $K^{(i)}$ constitutes the major contribution to the overall program runtime (55%; 177s assembly, 320s overall). Therefore, the assembly subroutine is refactored as a preliminary step before parallelization.

Our setting allows certain simplifications in the assembly. In unstructured finite element meshes the number of nonzeros in each row can largely differ. In standard CSR-type matrix storage, as implemented in the PETSc-AIJ format, the performance penalty for insufficient preallocation of memory for the rows is very high. Thus, an upper bound for the nonzeros per row has to be estimated or the number of nonzeros has to be precomputed by looping over all finite elements. The latter approach will effectively double the assembly time since the loop over all finite elements has to be traversed twice. We often use a custom, flexible sparse matrix type that does not require memory preallocation and still allows efficient insertion of entries with a cost logarithmic in the number of nonzero entries. However, in the setting presented here, the performance of this matrix format is lower than expected. Moreover, here it is known that the number of nonzeros in each row is bounded by 38. A simple preallocation can therefore be used and the custom matrix class can be eliminated. This results in a reduction of the assembly time by 77% (177s vs. 40s). With this approach, the overall program runtime is halved (320s vs. 157s). As a result, the assembly is not the dominant runtime contribution any more (40 of 157s, about 25%). Nevertheless, we proceed with parallelizing the assembly phase.

The local stiffness matrices on the finite elements can be computed independently. Therefore, no data dependencies between iterations of the finite element loop exist, and the loop can be parallelized using `OMP_SCHEDULE=static`. After declaring all data structures to be local to each loop iteration, the insertion of the values into the target matrix with the PETSc function `MatSetValues` is protected by an OpenMP `critical` region. This is necessary, since in finite element matrices, matrix entries will generally be accessed multiple times. The serialization of the calls to `MatSetValues` exposed no drawback in terms of performance. An implementation of a matrix coloring scheme as in [18] is therefore not required in our case. This straightforward threading approach of the assembly routine shows good scalability. A speedup of 8 using 10 threads per process is reached; see the solid line in Figure 4.

In order to determine whether certain use of STL containers, as, e.g., calls to `vector.pushback()`, within the OpenMP parallel part of the assembly routine is performance critical, we also conduct a single-thread runtime profile (using `gprof` and Intel `loopprofileviewer`). The profiling reveals that the major part, indeed about 90%, of the assembly runtime is spent in the computation of the element stiffness matrix. This is to be expected. This function in turn is then dominated by the innermost of three nested loops. The outermost loop is a loop over the
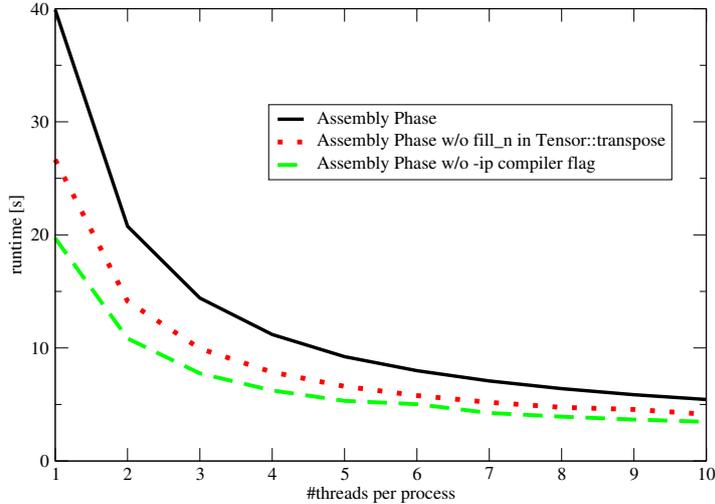
Fig. 4 Socket scaling of the OpenMP parallel assembly routine. Measurements are based on consecutive code modifications (Neo-Hookean hyperelasticity problem in 2D with $m_x = m_y = 200$).

integration points given by a numerical integration scheme. The two nested inner loops are loops over the degrees of freedom of the finite element.

For the computation of the element stiffness matrix, the transposes of matrices storing the derivatives of the finite element basis functions are required. The derivatives are implemented using a custom C++ class *Tensor*, which provides several tensor operations, as, e.g., a transpose method. Surprisingly, this transpose method from the Tensor class is the hotspot in the innermost for loop. In this method the standard constructor is used to construct a temporary Tensor object. About 50% (1.4 billion) of the calls to this constructor can be accounted to transpose. This constructor then initializes a Tensor element using `std::__fill_n`, which, surprisingly, turns out to be the largest runtime contribution in the whole application: According to the profile, this function makes up 25% of the overall runtime. As the dimension of the input and output Tensors in this case are identical, we have implemented a specialized transpose function that does not use the standard constructor to initialize the returned matrix but the implicit copy constructor, and therefore avoids excessive calls to `std::__fill_n`. A subsequent profiling run shows that this transpose now makes up only 0.05% of the overall runtime. Single-thread assembly runtime is reduced by about one third; cf. the solid and dotted lines in Figure 4. The reason for this, however, is not a reduction of computational complexity. The copy constructor in the specialized transpose as well as the `std::__fill_n` in the default constructor both loop over all Tensor elements. The reason is rather a reduction of code complexity which enables the compiler to inline and optimize the modified transpose method, which it was not able to do with the original version. As other functions still use `std::__fill_n`, including the standard constructor

and, e.g., the multiplication function, its share of overall runtime is now about 8% and subject to further optimizations.

In order to get a more detailed profile we have tried to reduce code inlining by compiling the application **without** the `-ip` (inter-procedural optimization) compiler flag. This did not influence profile detail, but had an unexpected influence on performance: The original single-thread assembly runtime is halved, regardless whether the original or specialized transpose function is used; cf. the solid and dashed lines in Figure 4. Apparently, the compiler itself is now able to perform the optimizations just described, i.e., it is able to use inlining. This is surprising, as inter-procedural optimization is intended to provide the compiler a better overview of the code. However, in the present case it caused a lack of appropriate inlining. With this, overall single-thread program runtime reduces to 134s, where the assembly phase requires 20s (15%). As building PETSc without the `-ip` compiler switch seems to have no impact on performance in our setting, we now generally omit it. Assembly of the load vector $f^{(i)}$ is currently of minor impact to overall performance. Nevertheless, the steps described for the stiffness matrix here should be considered there, too, in the future.

7.2 Shared Memory Parallel Direct Solver PARDISO

To provide shared memory parallelism in the *FETI-DP Setup* phase, which is dominated by the LU factorizations, we use the shared memory parallel direct solver package PARDISO [27] from the Intel MKL library. We interface PARDISO through PETSc using a third party interface [8]. The current PETSc version 3.5 ships with an included PARDISO interface [6]. We use PARDISO to perform the sparse LU factorizations (*FETI-DP Setup* phase in the graphs) and the forward/backward substitutions (*Krylov Iteration* phase in the graphs) in a threaded parallel fashion within each MPI process. Our PARDISO options are shown in Section 6.1.

The runtime of the threaded PARDISO-based phases *FETI-DP Setup* and *Krylov Iteration* is shown for a single socket in Figure 5. Both parts show a stepwise decrease in runtime and a speed-up of only 3.5 to 4 on ten cores. Substantial performance improvement can be seen when hitting "magic" thread counts, which are powers of two, while runtime stays constant otherwise. This behaviour is a first indication that the PARDISO phases are not limited by typical hardware bottlenecks such as the data transfer bandwidths. In Figure 6 we confirm this assumption by showing the bandwidths utilized over the different data paths as measured with our LIKWID tool. The *Krylov Iteration* has the highest main memory bandwidth utilization but still only exploits 50% of the maximum bandwidth of the processor. It is interesting to see that bandwidth over the on-chip data paths is almost identical to main memory bandwidth for the *Krylov Iteration*, clearly indicating that the current implementation is stream like and makes no reuse of data in L2 and L3 cache. A different characteristic shows up for the *FETI-DP Setup* phase where main memory bandwidth is of no relevance but cache reuse seems to be high due to strong increase
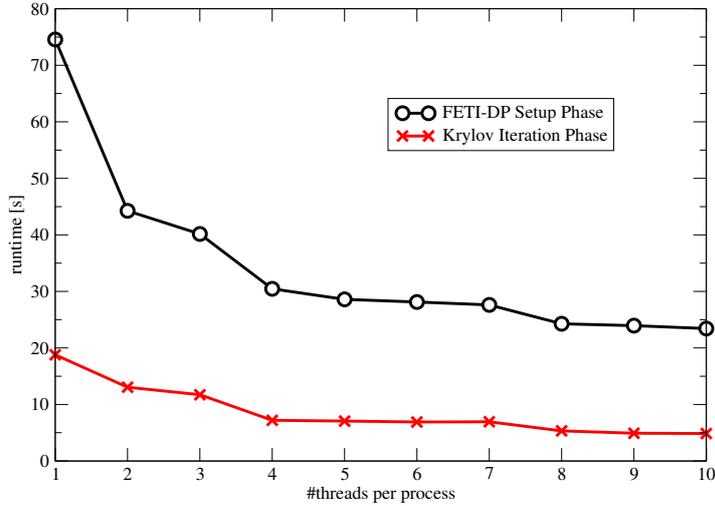
Fig. 5 Socket scaling of program parts parallelized with PARDISO (Neo-Hookean hyperelasticity problem in 2D with $m_x = m_y = 200$).

	Node	Socket 0										Socket 1									
		C0	C1	C2	C3	C4	C5	C6	C7	C8	C9	C0	C1	C2	C3	C4	C5	C6	C7	C8	C9
Setup 1	0	R0 / T0	T1	T2	T3	T4	T5	T6	T7	T8	T9										
	1	R1 / T0	T1	T2	T3	T4	T5	T6	T7	T8	T9										
	2	R2 / T0	T1	T2	T3	T4	T5	T6	T7	T8	T9										
	3	R3 / T0	T1	T2	T3	T4	T5	T6	T7	T8	T9										
Setup 2	0	R0 / T0	T1	T2	T3	T4	T5	T6	T7	T8	T9	R1 / T0	T1	T2	T3	T4	T5	T6	T7	T8	T9
	1	R2 / T0	T1	T2	T3	T4	T5	T6	T7	T8	T9	R3 / T0	T1	T2	T3	T4	T5	T6	T7	T8	T9
Setup 3	0	R0 / T0	T1	T2	T3	T4	R2 / T0	T1	T2	T3	T4	R1 / T0	T1	T2	T3	T4	R3 / T0	T1	T2	T3	T4
Setup 4	0	R0 / T0	T1	R1 / T0	T1	R2 / T0	T1	R3 / T0	T1												

Table 1 Distribution of threads (T) and processes (R, MPI ranks) on the machine for the four hybrid setups using 4 MPI ranks. Each node has two sockets, each socket holds an Ivy Bridge processor with 10 cores (C).

of utilized bandwidth for inner level caches. However the cache bandwidths shown in Figure 6 for *FETI-DP Setup* are still a factor of 5x-10x away from practical hardware limitations ruling them out as potential bottlenecks. Moreover, the LIKWID analysis has revealed a high AVX vectorization rate for *FETI-DP Setup* (which is the time critical phase) and pure scalar execution for the *Krylov Iteration*. This analysis gives a clear indication that improving the parallelization approach in the two phases is the most promising strategy to further boost performance for the direct solver phase. Work in this direction has been started with one of the developers of PARDISO (Olaf Schenk).
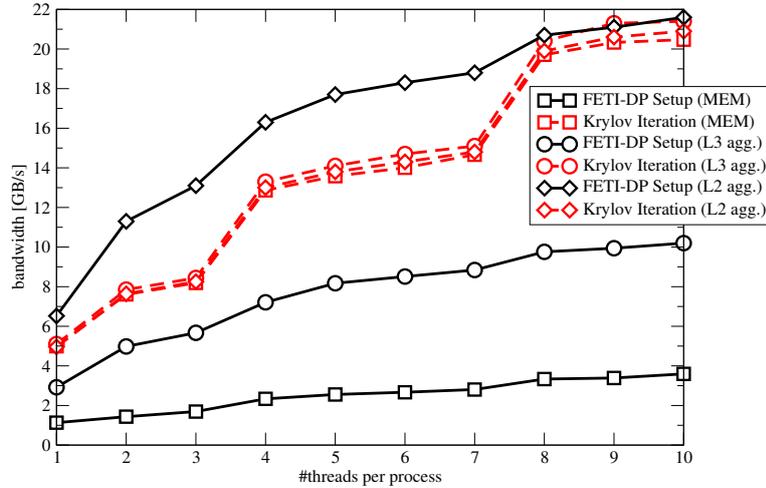
Fig. 6 Bandwidth utilization of the memory hierarchy levels of one socket for the two PARDISO parallel program parts. The different symbols mark bandwidth between memory and L3 cache (squares), L3 and L2 caches (circles), and L2 and L1 caches (diamonds). Cache bandwidth values are aggregated over all utilized cores. The available socket memory bandwidth of the IvyBridge system is 42 GB/s.

7.3 Runtime Study of the Full Threaded Software

We now consider the total time to solution for our Neo-Hooke material using the OpenMP threaded assembly and the shared memory solver PARDISO. We use 4 subdomains distributed to 4 MPI ranks. The 4 processes are placed on 4 nodes (1 process per node, Setup 1), on the 4 sockets of 2 nodes (Setup 2), on the 2 sockets of 1 node (Setup 3) or on a single socket (Setup 4). Then we vary the number of threads, i.e., each MPI process spawned up to 10 threads (Setups 1 and 2), up to 5 threads (Setup 3), and up to 2 threads (Setup 4). See Table 1 for an illustration of the hybrid setups. The results are presented in Figure 7. Interestingly, we observe that the quantitative scaling behavior (runtime variation) is virtually independent of the hybrid configuration, i.e., the distibution of processes and threads within nodes and sockets. The parallel efficiency falls below 50% when 5 or more threads are used.

Since the scaling behaviour is basically identical for the different hybrid setups (Setup 1,2,3, and 4), we report on detailed timings only for Setup 1; see Table 2. The scaling results are not optimal but encouraging. Satisfactory results could be achieved using up to 10 threads for the *Assembly* phase adopting our simple approach of protecting the calls to `MatSetValues` by an OpenMP `critical` region, and without parallelization of the load vector assembly. The scaling of the *FETI-DP Setup* and the *Krylov iteration* phase relies on the scaling behavior of PARDISO. It is known that perfect parallel scalability of direct solvers for sparse linear systems from PDEs is hard to achieve.
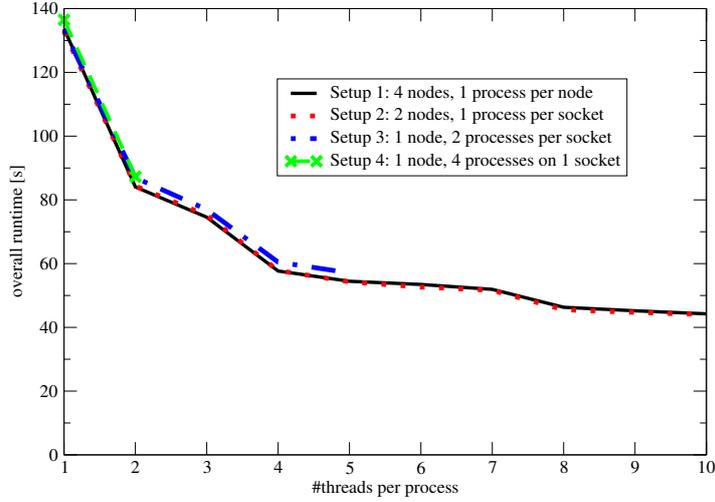
Fig. 7 Socket scaling of Neo-Hookean hyperelasticity problem in 2D with $m_x = m_y = 200$ using the four hybrid setups described in Section 7.3. Assembly is threaded using OpenMP, *FETI-DP Setup* and *Krylov Iteration* use the threaded PARDISO solver.
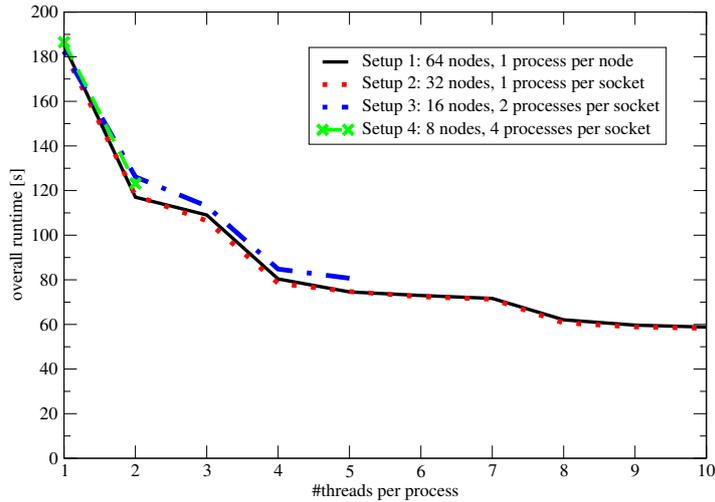

Fig. 8 Socket scaling of Neo-Hookean hyperelasticity problem in 2D with $m_x = m_y = 200$ as in Figure 7 but using 64 MPI ranks and a maximum of 640 threads.

From Table 2, we observe that for our hybrid MPI/OpenMP approach we obtain an overall efficiency of 58% using 4 OpenMP threads per MPI rank. Investing up to 10 threads per MPI rank still reduces the total execution time but at diminishing returns. Our results are thus similar to the findings in [16, Table II] for FETI-1.

In Figure 8 the experiments from Figure 7 are repeated but now using 64 MPI ranks and up to 640 threads. The results are qualitatively identical. Due to numerical

Threads	Assembly	Efficiency	FETI-DP Setup	Efficiency	Krylov Iteration	Efficiency	Overall	Efficiency
1	19.71s	100%	74.56s	100%	18.79s	100%	133.69s	100%
2	10.84s	91%	44.25s	84%	13.05s	72%	84.04s	80%
3	7.75s	85%	40.16s	62%	11.72s	53%	74.56s	60%
4	6.24s	79%	30.46s	61%	7.20s	65%	57.68s	58%
5	5.32s	74%	28.60s	52%	7.06s	53%	54.46s	49%
6	5.02s	65%	28.12s	44%	6.90s	45%	53.47s	42%
7	4.25s	66%	27.61s	39%	6.94s	39%	51.91s	37%
8	3.92s	63%	24.27s	38%	5.33s	44%	46.30s	36%
9	3.67s	60%	23.94s	35%	4.91s	43%	45.19s	33%
10	3.46s	57%	23.43s	32%	4.86s	39%	44.24s	30%

Table 2 Runtime in seconds of overall program and threaded phases for one up to ten threads per MPI process (Setup 1, i.e., one MPI process on each of the four nodes, Neo-Hookean hyperelasticity problem in 2D with $m_x = m_y = 200$).

effects the number of Krylov iterations is slightly larger. Instead of four Newton steps and between 61 and 67 Krylov iterations it now takes five Newton steps and between 209 and 211 Krylov iterations until convergence resulting in slightly higher runtimes.

Summary   We have presented a thorough profiling and code investigation in combination with a hardware bottleneck analysis for important steps of the FETI-DP method with special focus on node-level performance. It has been demonstrated that the impact of widely used C++ techniques on performance needs to be carefully investigated and rather basic code changes may result in large performance improvements. For two phases *FETI-DP Setup* and the *Krylov iteration* we have clearly shown that the code is "core-bound" and identified that the current parallelization strategy is the performance limiting factor on the socket level. Improving shared memory parallelization of the respective routines in the widely used PARDISO solver may not only provide additional performance for our application but will be of great interest for a large community.

References

1. P. R. Amestoy, I. S. Duff, and J.-Y. l'Excellent. Multifrontal parallel distributed symmetric and unsymmetric solvers. *Computer Methods in Applied Mechanics and Engineering*, 184(2-4):501–520, 2000.
2. P. R. Amestoy, I. S. Duff, J.-Y. l'Excellent, and J. Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM Journal of Matrix Analysis and Applications*, 23(1):15–41, 2001.
3. Patrick R. Amestoy, Iain S. Duff, Jean-Yves l'Excellent, and Jacko Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM Journal on Matrix Analysis and Applications*, 23(1):15–41, 2001.

4. Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Karl Rupp, Barry F. Smith, and Hong Zhang. PETSc Web page. `http://www.mcs.anl.gov/petsc`, 2014.

5. Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Karl Rupp, Barry F. Smith, and Hong Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.5, Argonne National Laboratory, 2014.

6. Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Karl Rupp, Barry F. Smith, and Hong Zhang. Changes in the petsc 3.5 version, 2014. http://www.mcs.anl.gov/petsc/documentation/changes/35.html.

7. Satish Balay, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press, 1997.

8. Jose D. Bermeo. Added support for mkl-pardiso solver, 11 2013. https://bitbucket.org/petsc/petsc/pull-request/105/added-support-for-mkl-pardiso-solver/commits.

9. Manoj Bhardwaj, Kendall H. Pierson, Garth Reese, Tim Walsh, David Day, Ken Alvin, James Peery, Charbel Farhat, and Michel Lesoinne. Salinas: A scalable software for high performance structural and mechanics simulation. In *ACM/IEEE Proceedings of SC02: High Performance Networking and Computing. Gordon Bell Award.*, pages 1–19, 2002.

10. T. A. Davis and I. S. Duff. An unsymmetric-pattern multifrontal method for sparse lu factorization. *SIAM J. Matrix Anal. Applic.*, 18(1):140–158, 1997.

11. T. A. Davis and I. S. Duff. A combined unifrontal/multifrontal method for unsymmetric sparse matrices. *ACM Trans. Math. Softw.*, 25(1):1–19, 1999.

12. Timothy A. Davis. A column pre-ordering strategy for the unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software*, 30(2):165–195, June 2004.

13. Charbel Farhat, Michel Lesoinne, Patrick LeTallec, Kendall Pierson, and Daniel Rixen. FETI-DP: A dual-primal unified FETI method - part i: A faster alternative to the two-level FETI method. *Internat. J. Numer. Methods Engrg.*, 50:1523–1544, 2001.

14. Charbel Farhat, Michel Lesoinne, and Kendall Pierson. A scalable dual-primal domain decomposition method. *Numer. Lin. Alg. Appl.*, 7:687–714, 2000.

15. Ibrahima Guèye. *Solving large linear systems arising in finite element approximations on massively parallel computers*. Theses, École Nationale Supérieure des Mines de Paris, December 2009.

16. Ibrahima Guèye, S. El Arem, Fédéric Feyel, François-Xavier Roux, and Georges Cailletaud. A new parallel sparse direct solver: Presentation and numerical experiments in large-scale structural mechanics parallel computing. *International Journal for Numerical Methods in Engineering*, 88(4):370–384, 2011.

17. Ibrahima Guèye, Xavier Juvigny, Fédéric Feyel, François-Xavier Roux, and Georges Cailletaud. A parallel algorithm for direct solution of large sparse linear systems, well suitable to domain decomposition methods. *European Journal of Computational Mechanics/Revue Europenne de Mcanique Numrique*, 18(7-8):589–605, 2009.

18. Xiaohu Guo, Gerard Gorman, Michael Lange, Andrew Sunderland, and Mike Ashworth. Developing hybrid openmp/mpi parallelism for fluidity-icom - next generation geophysical fluid modelling technology. http://www.hector.ac.uk/cse/distributedcse/reports/fluidity-icom02/fluidity-icom02.pdf, Final Report for DCSE ICOM.

19. Axel Klawonn, Martin Lanser, and Oliver Rheinbach. Towards extremely scalable nonlinear domain decomposition methods for elliptic partial differential equation. Technical Report 2014-13, Preprint Reihe, Fakultät für Mathematik, TU Bergakademie Freiberg, ISSN 1433-9407, 2014. `http://tu-freiberg.de/fakult1/forschung/preprints`. Submitted to SISC.

20. Axel Klawonn, Martin Lanser, and Oliver Rheinbach. A nonlinear FETI-DP method with an inexact coarse problem. In Thomas Dickopf, Martin J. Gander, Rolf Krause, and Luca F. Pavarino, editors, *Domain Decomposition Methods in Science and Engineering XXII*. Lecture Notes in Computational Science and Engineering, Springer-Verlag, 2015. Accepted for publication October 2014. Proceedings of the 22nd Conference on Domain Decomposition Methods in Science and Engineering, Lugano, Switzerland, September 16-20, 2013. Also `http://tu-freiberg.de/fakult1/forschung/preprints`.
21. Axel Klawonn and Oliver Rheinbach. Inexact FETI-DP methods. *Internat. J. Numer. Methods Engrg.*, 69(2):284–307, 2007.
22. Axel Klawonn and Oliver Rheinbach. Highly scalable parallel domain decomposition methods with an application to biomechanics. *ZAMM Z. Angew. Math. Mech.*, 90(1):5–32, 2010.
23. Axel Klawonn and Olof B. Widlund. Dual-Primal FETI Methods for Linear Elasticity. *Comm. Pure Appl. Math.*, 59(11):1523–1572, 2006.
24. A. Kuzmin, M. Luisier, and O. Schenk. Fast methods for computing selected elements of the greens function in massively parallel nanoelectronic device simulations. In F. Wolf, B. Mohr, and D. Mey, editors, *Euro-Par 2013 Parallel Processing*, volume 8097 of *Lecture Notes in Computer Science*, pages 533–544. Springer Berlin Heidelberg, 2013.
25. Oliver Rheinbach. Parallel iterative substructuring in structural mechanics. *Arch. Comput. Methods Eng.*, 16(4):425–463, 2009.
26. Olaf Schenk, Matthias Bollhöfer, and Rudolf A. Römer. On large-scale diagonalization techniques for the anderson model of localization. *SIAM Rev.*, 50(1):91–112, February 2008.
27. Olaf Schenk, Andreas Wächter, and Michael Hagemann. Matching-based preprocessing algorithms to the solution of saddle-point problems in large-scale nonconvex interior-point optimization. *Computational Optimization and Applications*, 36(2-3):321–341, 2007.
28. Barry F. Smith, Petter E. Bjørstad, and William Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*. Cambridge University Press, 1996.
29. Andrea Toselli and Olof Widlund. *Domain Decomposition Methods - Algorithms and Theory*, volume 34 of *Springer Series in Computational Mathematics*. Springer, 2004.
30. Jan Treibig, Georg Hager, and Gerhard Wellein. LIKWID: A lightweight performance-oriented tool suite for x86 multicore environments. In *PSTI2010, the First International Workshop on Parallel Software Tools and Tool Infrastructures*, pages 207–216, Los Alamitos, CA, USA, 2010. IEEE Computer Society.
31. O.C. Zienkiewicz and R.L. Taylor. *The Finite Element Method for Solid and Structural Mechanics*. Elsevier, Oxford, 2005.