



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Work-stealing prefix scan: Addressing load imbalance in large-scale image registration

Citation for published version:

Copik, M, Grosser, T, Hoefler, T, Bientinesi, P & Berkels, B 2022, 'Work-stealing prefix scan: Addressing load imbalance in large-scale image registration', *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 3. <https://doi.org/10.1109/TPDS.2021.3095230>

Digital Object Identifier (DOI):

[10.1109/TPDS.2021.3095230](https://doi.org/10.1109/TPDS.2021.3095230)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

IEEE Transactions on Parallel and Distributed Systems

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Work-stealing prefix scan: Addressing load imbalance in large-scale image registration

Marcin Copik¹, Tobias Grosser², *Member, IEEE*, Torsten Hoefler¹, *Senior Member, IEEE*,
Paolo Bientinesi³, Benjamin Berkels^{4,5}

¹Department of Computer Science, ETH Zurich; ²School of Informatics, University of Edinburgh;

³Department of Computing Science, Umeå University; ⁴AICES Graduate School, RWTH Aachen University; ⁵Institute for Geometry and Practical Mathematics, RWTH Aachen University

Abstract—Parallelism patterns (e.g., map or reduce) have proven to be effective tools for parallelizing high-performance applications. In this paper, we study the recursive registration of a series of electron microscopy images – a time consuming and imbalanced computation necessary for nano-scale microscopy analysis. We show that by translating the image registration into a specific instance of the prefix scan, we can convert this seemingly sequential problem into a parallel computation that scales to over thousand of cores. We analyze a variety of scan algorithms that behave similarly for common low-compute operators and propose a novel work-stealing procedure for a hierarchical prefix scan. Our evaluation shows that by identifying a suitable and well-optimized prefix scan algorithm, we reduce time-to-solution on a series of 4,096 images spanning ten seconds of microscopy acquisition from over 10 hours to less than 3 minutes (using 1024 Intel Haswell cores), enabling derivation of material properties at nanoscale for long microscopy image series.

Index Terms—prefix sum, parallel algorithms, work stealing, load balancing, image registration

1 INTRODUCTION

Many seemingly sequential algorithms in which the computation of an element y_{i+1} depends on the preceding element y_i can be parallelized with a *prefix scan* operation. Such an operation takes a binary and associative operator \odot and an input array x_0, x_1, \dots, x_n and produces the output array y_0, y_1, \dots, y_n . Every element y_i in the output array is the result of the binary and inclusive combination of all “previous” elements and the current one in the input array: $y_i = x_0 \odot x_1 \odot \dots \odot x_i$. An *exclusive prefix scan* computes an output array where the combination does not include the current element. One can show that, if a sequentially-dependent transformation consuming y_i and producing y_{i+1} can be expressed as such associative and binary operator, then the problem can be parallelized using a prefix scan.

This powerful construct has numerous uses in parallel computing. It enables parallelization of multiple non-trivial problems that might seem to be inherently sequential, including finite state machines, solving linear tridiagonal systems, parallelization of many sequential loops with dependencies, or sequential chains of computations that can be modeled as a function composition [1], [2], [3], [4]. However, the sequentiality in the original problem causes some overheads—the resulting parallel algorithm is either highly parallel or work efficient but not both at the same time. The more processes can be used effectively, the more additional work (i.e., applications of \odot) has to be performed in parallel. If the workload is balanced and all processes run at the same speed, then the additional work does not delay the processing. In this scenario, a prefix scan is as fast as a simple reduction albeit with a higher energy consumption due to additional computation.

In this work, we parallelize an application in the area

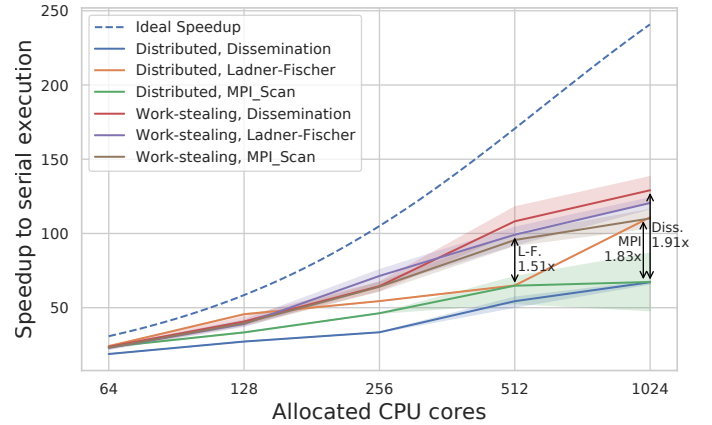


Fig. 1: The strong scaling of distributed prefix scan (Section 4) on the *scan* image registration for three variants of prefix scan algorithms (Section 2.1). Experimental results (solid lines) were obtained on the Piz Daint [5] system for 4,096 images. Theoretical bound (3) is discussed in Section 5.2.

of large-scale image registration in electron microscopy. Given the importance of microscopy data for analysis of material properties and temporal changes at nanoscale, a critical objective is to enable the processing of very long sequences of microscopy images by a domain specialist without running days long computations. We look at the grand scheme of registration and represent this seemingly sequential procedure as a composition of two steps, a massively parallel preprocessing phase and a prefix scan. In contrast to most other applications of prefix scans, our appli-

cation requires *expensive and highly load-imbalanced operators with nearly trivial communication*. We are first to show (1) how the communication pattern in near-optimal circuits by Ladner and Fischer can be tuned for MPI execution on large-scale compute clusters and (2) we develop a novel node-local work-stealing algorithm for general prefix scans, found in a large variety of recursive and seemingly sequential computations. The load balancing scan enables the parallelization of problems that would otherwise be considered inefficient given imbalanced computation, sparse iteration space, and a tightly constrained form of prefix scan.

We apply the load balancing prefix scan to the registration procedure and we show that the performance of distributed prefix scans can be significantly improved even for a highly imbalanced application. In the strong scaling experiment (Section 5.2), our hierarchical dynamic approach achieves speedups of up to 1.51x, 1.83x, and 1.91x for different scan algorithms, as presented in Figure 1, while decreasing the CPU allocation time and overall energy consumption up to 1.87x and 2.23x times, respectively (Section 5.4).

Our paper makes the following contributions:

- A novel node-local, work-stealing prefix scan that exploits the hierarchy of parallel workers and memories to (1) decrease performance and energy costs of a distributed prefix scan and (2) exploit the additional levels of a shared-memory parallelization to construct an efficient load balancing step. To the best of our knowledge, this is the first scan algorithm designed for problems with unbalanced workloads, and we offer an open-source implementation¹.
- A scalable and efficient parallelization strategy for recursive image registration that enables analysis of temporal changes in long microscopy acquisitions. With our dynamic prefix scan, the performance of image registration is improved up to two times while decreasing energy costs by over two times.
- A novel example of a parallel scan problem, which performance challenges have not been addressed by research on parallel algorithms and MPI collectives, and a generic solution for expensive and unbalanced scan operators to fill this important gap in the quality of MPI collectives.

2 BACKGROUND AND MOTIVATION

Recent advances in transmission electron microscopy have allowed for a more precise visualization of materials and physical processes, such as metal oxidation, at nanometer resolution. Yet, many environmental factors negatively affect the quality of microscopy images. A novel registration method [6] has been proposed to mitigate these limitations by acquiring a series of low dose microscopy frames and aligning each frame to the first frame with an image registration procedure (Section 2.3). With this strategy, the increased amount of reliable information extracted from noisy microscopy data is paid for by a computationally intensive and sequential process that becomes a bottleneck of the analysis. By phrasing the task of registering an image series as a special instance of the prefix scan (Section 2.1), we can

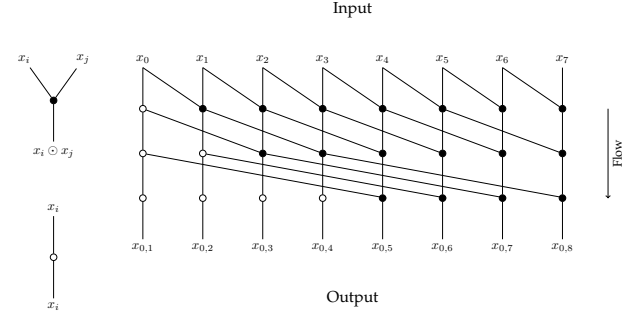


Fig. 2: The dissemination prefix scan. Black dots represent an application of the operator while a white dot indicates a communication that does not involve any computation on the receiver. An optimal logarithmic depth is achieved by performing $N - 2^i$ operations in i -th iteration. For input data of size 8, 17 operator applications are necessary to obtain results in 3 iterations.

use the universal parallel pattern to propose parallelization strategies for this recursive computation. We show that there are no known prefix scans that can handle very well problems incorporating a high computation to communication ratio and an unpredictable and variable execution time (2.2).

2.1 Prefix Scan

The importance and complexity of prefix scans make it one of the most studied basic patterns in parallel computing. Numerous algorithms exist that trade-off additional work and lower parallel depth. The work-depth relation of prefix scans was an open research problem for many decades [7], [8]. The sketch of the dissemination prefix sum in Figure 2 depicts the main idea applied to parallelize the prefix scan: a decrease in depth is obtained by performing multiple computations on a single data element. Depth-optimal algorithms cannot be zero-deficient [8], i.e., an increase in work must be larger than the decrease in depth. Although depth minimization is the primary goal when designing scalable algorithms, a huge work intensity usually implies an excessive communication. Work-inefficient algorithms are more sensitive to deviations in execution time since they require more applications of the binary operator. Imbalanced operators will affect differently various scan algorithms due to differences in propagation of dependencies.

A tree-based prefix scan is one of the classical parallel prefix scan strategies, as presented by Blelloch [9] and Brent et al. [10]. For both algorithms, the depth is bounded by a double traversal of a binary tree. The dissemination prefix scan, also known as the recursive doubling [11], was presented by Kogge et al. [12] and Hillis et al. [13]. The recursive family of prefix circuits presented by Ladner et al. [14] achieve an asymptotically smaller work overhead at optimal time but are rarely used in practice due to a less favorable communication pattern. For most of the well-known parallel prefix circuits, the depth is given as $C_1 \log_2 N + C_2$, where C_1 and C_2 are integer constants. The constant C_2 is non-zero for algorithms such as a tree-based inclusive scan presented by Brent et al., which has one layer less than the exclusive Blelloch scan. The Ladner-Fischer scan is designed with a constant C_2 that controls the depth-work balance.

1. <https://github.com/berkels/match-series>

Name	Type	Depth	Work
Sequential	I	$N - 1$	$N - 1$
Blelloch	E	$2 \log_2 N$	$2(N - 1)$
Dissemination	I	$\log_2 N$	$N \log_2 N - N + 1$
Ladner–Fischer	I	$\log_2 N$	$< 4N - 5$

TABLE 1: Major *I*-nclusive and *E*-xclusive parallel prefix scan algorithms. The exact work for the Ladner–Fischer scan is given by a recursive equation in N .

Table 1 presents a comparison of the discussed parallel prefix scan algorithms. Exclusive and inclusive variants are specified in the Message Passing Interface standard as the collective operations `MPI_Exscan` and `MPI_Scan` [15], respectively. They are implemented using either simple algorithms that achieve the optimal runtime of $\log_2 P$ rounds on P processes or tree-based algorithms that optimize the communication latency [16]. Most scan implementations are optimized for the common case that *communication time dominates computation time and that computation is balanced*.

2.2 Related Work

Standard strategies for a prefix scan when data size significantly exceeds the number of parallel workers have been frequently presented by other authors. Kruskal et. al [17] presented such algorithm on an EREW model [17]. It was later applied on a binary tree network of processors by Meijer [18] and to solving a tridiagonal linear system on a hypercube architecture by Eğecioğlu et al. [11]. There, the authors define the algorithm with a fixed choice of the dissemination as a global scan algorithm, whereas we present in Section 4 a generic distributed prefix sum and consider various scan algorithms in the global phase. Eğecioğlu et. al [19] introduced a recursive algorithm for a distributed prefix scan, where it was found to have better efficiency than the previous approaches when the discrepancy between computation and communication cost is significant. This work was the first one to introduce more complex data distribution for prefix scan and to design a prefix scan strategy for computationally intensive operators. Data segments are distributed according to the number of arithmetical steps performed by each processor, an information that cannot be either estimated or predicted in problems such as image registration. Thus, designing a prefix scan for applications with an unknown load balance is an open problem. Chatterjee et. al. [1] defined for the vectorization of prefix scan on CRAY-MP a strategy known as the *reduce-then-scan*. Maleki et al. [20] present a hierarchical prefix sum implementation for GPU architectures. Although we use the same strategies as a basis for hierarchical and work-stealing scan, we extend it extensively with a dynamic accumulation of partial results and load balancing between neighboring threads.

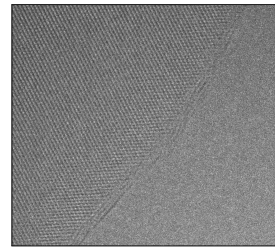
Research on tuning `MPI_Scan` and `MPI_Exscan` collectives is focused on reducing the communication cost and improving bandwidth on memory-bound operators with computation being far cheaper than communication. Sanders et al. [16] used pipelined binary trees and improved later the performance of the prefix scan in message-passing systems by exploiting a bidirectional communication [21], [22]. The improvements are limited to prefix operators

bounded by network latency, which is not the case for the image registration.

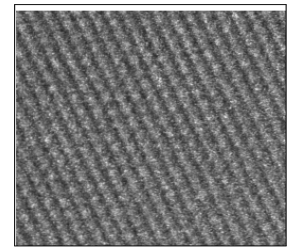
2.2.1 Specific prefix scan operators

Although the prefix scan research has been dominated by optimizations dedicated to trivial operators, there has been few examples of prefix scans with computationally intensive operators. Maleki et al. [23] consider prefix scan solution to the linear-tropical dynamic programming problem where the operation is a matrix-matrix multiplication. Gradl et al. [24] presented a parallel prefix algorithm for accumulation of matrix multiplications in quantum control. Waldherr et al. [25] and Auckenthaler [26] showed later that prefix scan parallelization of this operation is outperformed by a sequential prefix scan with parallel matrix multiplication operator. These applications of prefix scan resulted in neither tuning nor designing a scan algorithm for operators where computation time is significantly larger than communication.

In addition to performance improvements shown on the image registration problem, our work-stealing scan can be applied to improve the efficiency of other imbalanced scans as well, and excellent examples of imbalanced operators are sparse linear algebra operations, found in the scan parallelization of neural network backpropagation with sparse matrix operations [27]. Prefix scans are essential for the automatic parallelization of loop-carried dependencies [28]. While polyhedral techniques allow for approximating a balanced distribution of non-uniform loop nests, a dynamic work-stealing would improve the performance when static scheduling is not possible due to dynamic and data-dependent control-flow.



(a) Frame 26



(b) Deformed Frame 26, top left

Fig. 3: Frame 26 of the acquisition (left) and magnified after alignment to Frame 25 (right). The movement of the frame along vertical axis is visible as white stripe on the top. We observe a low variability between images acquired in a short timespan. TEM data courtesy of Sarah Haigh, University of Manchester.

2.3 Image Registration

We consider a series of two-dimensional, noisy atomic-scale electron microscopy images f_0, f_1, \dots, f_N that are used instead of a single, high quality frame acquired with a high-dose electron beam. Short-exposure image series are used since they allow to obtain a higher precision than a single image in the electron microscopy setting [29]. This replacement requires an aggregation of the information contained in the entire image series, usually done by averaging the images. However, the images cannot be averaged directly,

since they are affected by environmental noise of the observed sample during acquisition. Considering that electron microscopy allows for a magnification by more than 10 million, even movements of the sample by just half of the width of an atom result in shifts of the observed images by several pixels. To mitigate the effects of sample drift, each frame is registered to the first image. Since the images are showing atomic grids, they have a high degree of self similarity in the form of (nearly) periodic structures, cf. Fig. 3. This periodicity makes the registration much more difficult: Given a pair of (nearly) periodic images without any prior information on their relative shift, registration can only determine the shift up to a multiple of the period of the images, which is not sufficient for the reconstruction. If the estimated shifts are off by a multiple of the period, unrelated positions will be averaged, which will blur and duplicate deviations, but this deviations are what is actually interesting for the applications, since they can significantly influence the material properties. The shift between non-consecutive images can be large, which prevents directly registering non-consecutive images in this setting.

The need for HPC arises when the temporal behavior of the observed sample needs to be studied. In this setting, a series can consist of hundred thousand or more high resolution images that are still subject to the problem of periodicity. As of now, such series are simply not analyzed as a whole but only manually selected subsets. Being able to analyze such series in their entirety has a large potential to lead to new insights in materials science that are otherwise inaccessible.

3 PREFIX SCAN IMAGE PROCESSING

To enable parallelization and large-scale processing of long electron microscopy sequences, we represent this problem as a specific and unique case of the prefix scan. To this end, we explore the recursive nature of image series registration and define the prefix scan operator that inherits all properties of the computationally-intensive and iterative registration method. The resulting scan operator has three distinct features that set it apart from most of the other problems analyzed in the context of prefix scan parallelization: (1) non-trivial characteristics of the operator's associativity, (2) an unusually large ratio of computation to communication cost, and (3) unpredictable execution time causing load imbalance issues.

3.1 Image Registration

We define the problem of image registration for two-dimensional images $\mathcal{R}, \mathcal{T} \in \mathcal{I}$, known as reference and template images, respectively.

Definition 3.1. Image registration problem Given a distance measure $\mathcal{D} : \mathcal{I} \times \mathcal{I} \rightarrow \mathbb{R}$ and two images $\mathcal{R}, \mathcal{T} \in \mathcal{I}$, find a transformation $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that

$$\mathcal{D}(\mathcal{R}, \mathcal{T} \circ \phi)$$

is minimized. $\phi(x) = R(\alpha) \cdot x + t$ is a rigid transformation with angle α , rotation matrix $R(\alpha) \in \mathbb{R}^{2 \times 2}$ and translation $t \in \mathbb{R}^2$.

Intuitively, we want to find ϕ such that the deformed template image is aligned to the reference: $\mathcal{T} \circ \phi \approx \mathcal{R}$. We use the image registration procedure proposed by Berkels et. al. [6]. The approach defines a normalized cross-correlation functional as the distance measure and proposes a combination of a multilevel scheme with a gradient flow minimization process to solve the registration problem. The objective functional is characterized by the presence of multiple local minima. The computed deformation may vary not only between different starting points for the minimization but also among various implementations of the same algorithm, resulting in unpredictable computation time. We refer to an implementation of this technique as the function **A**. It

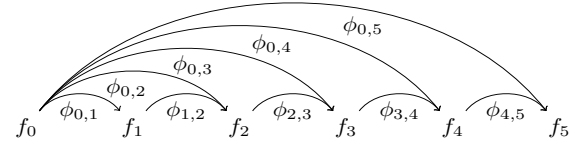


Fig. 4: The image registration process for a series of 6 frames. For an image f_i , the result from its predecessor $\phi_{0,i-1}$ is combined with a neighboring deformation $\phi_{i-1,i}$.

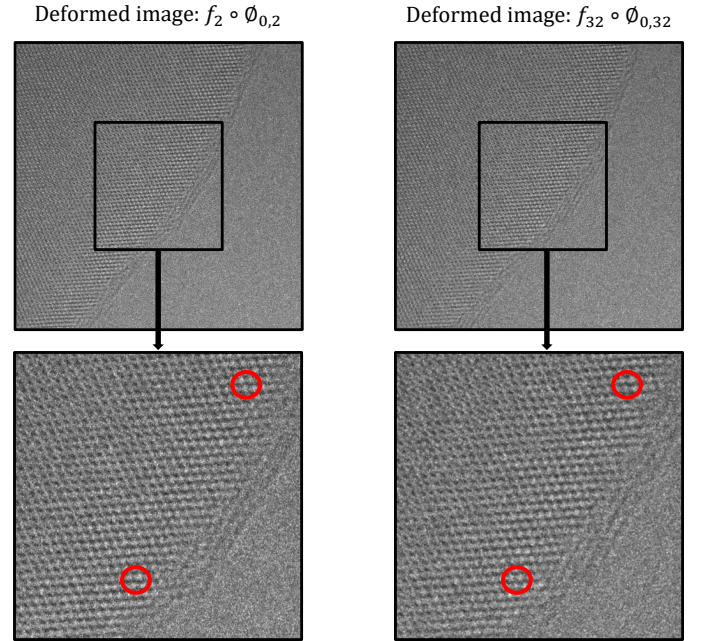


Fig. 5: The comparison of frames f_2 and f_{32} after applying the registration to f_0 . Red circles indicate points of interest used for verification of correct alignment.

accepts two images \mathcal{R} and \mathcal{T} , and an initial guess for the deformation ϕ' , and estimates a deformation $\phi_{\mathcal{R}, \mathcal{T}}$ with the proposed algorithm:

$$\phi_{\mathcal{R}, \mathcal{T}} = \mathbf{A}(\mathcal{R}, \mathcal{T}, \phi'),$$

3.2 Series Registration

The alignment problem requires a dedicated approach when the images are (nearly) periodic and that is the case for

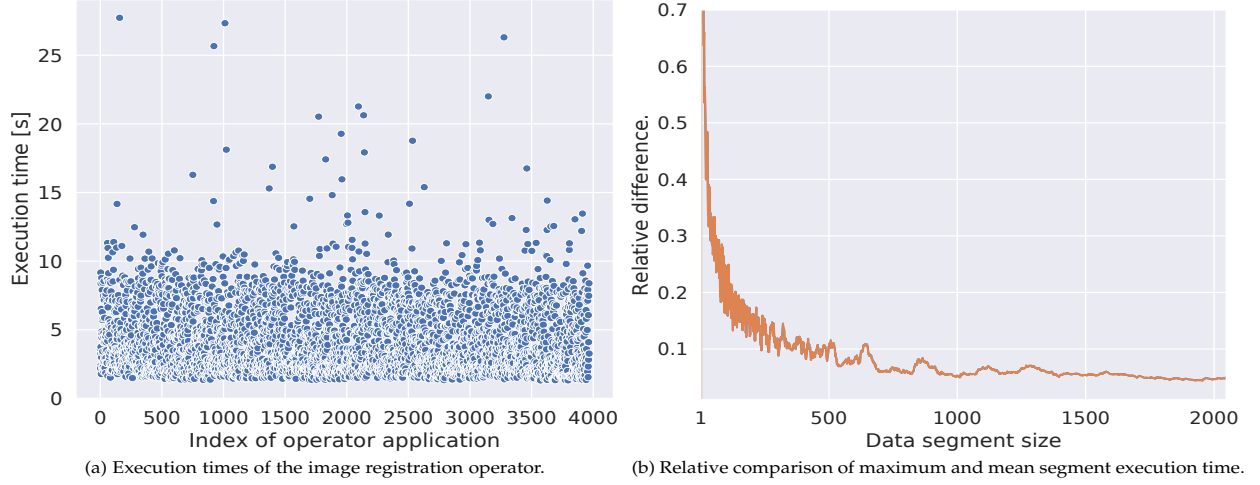


Fig. 6: The computationally intensive image registration operator: (a) execution times t_1, \dots, t_N of the operator in the first local phase of the prefix scan and (b) load imbalance for a static distribution with segment size $S = \frac{N}{P}$, where we estimate a relative difference between the mean ($\mu = \frac{S}{N} \sum T_s$) and maximum execution time ($\max_s T_s$) across all segments. The measurements were obtained on an Intel E5-2690 v3 CPU with 2.60 GHz base frequency.

electron micrographs. A correct registration of two frames is possible using the identity mapping as initial guess if the shift between them is smaller than half of the period. The validity of this assumption can be guaranteed only for neighboring frames f_i and f_{i+1} . For the generic registration of f_0 and any frame f_i , this limitation can be bypassed by taking into account all neighboring frames in between where the procedure is deemed to be accurate.

Given deformations $\phi_{0,1}$ and $\phi_{1,2}$ estimating $f_1 \circ \phi_{0,1} \approx f_0$ and $f_2 \circ \phi_{1,2} \approx f_1$, respectively, we can safely assume that the composition of deformations $\phi_{1,2} \circ \phi_{0,1}$ is a decent initial guess to register f_0 and f_2 , since

$$f_2 \circ (\phi_{1,2} \circ \phi_{0,1}) = (f_2 \circ \phi_{1,2}) \circ \phi_{0,1} \approx f_1 \circ \phi_{0,1} \approx f_0.$$

First, we define the registration procedure for two consecutive frames using the function **A**. We use the identity transformation I_ϕ since it is a suitable initial guess for consecutive images:

$$\forall i \in \mathbb{N} \phi_{i,i+1} = \mathbf{A}(f_i, f_{i+1}, I_\phi),$$

Then, we approximate the deformation for two non-consecutive frames by using the composition of two deformations as an initial guess. We reuse the previous function to define a new function **B** to handle the non-consecutive indices. It accepts two deformations $\phi_{j,k}$ and $\phi_{i,j}$ with the images f_i and f_k , and it estimates the deformation $\phi_{i,k}$:

$$\begin{aligned} \forall i, k \in \mathbb{N}, k > i + 1 \\ \forall j \in \mathbb{N}, i < j < k \\ \phi_{i,k} = \mathbf{B}(\phi_{i,j}, \phi_{j,k}) = \mathbf{A}(f_i, f_k, \phi_{j,k} \circ \phi_{i,j}). \end{aligned}$$

In particular, if we iterate consecutively from the first to i -th image (Figure 4):

$$\forall i \in \mathbb{N}, i > 1 \phi_{0,i} = \mathbf{B}(\phi_{0,i-1}, \phi_{i-1,i}).$$

In the formulation of the series registration problem above, the recursive nature is immediately seen: for any image f_i , the task of aligning to f_0 requires solving the

registration problem for f_0 and f_{i-1} first. Each final deformation $\phi_{0,i}$ can be obtained by consecutively applying the registration algorithm to neighboring deformations $\phi_{0,1}, \phi_{1,2}, \dots, \phi_{i-1,i}$. The accumulation of partial solutions can be represented as a prefix scan with the operator \odot_B defined as follows:

$$\begin{aligned} \phi_{i,j} \odot_B \phi_{j,k} &= \mathbf{B}(\phi_{i,j}, \phi_{j,k}) \\ \phi_{0,j} &= \phi_{0,1} \odot_B \phi_{1,2} \odot_B \dots \odot_B \phi_{i-2,i-1} \end{aligned}$$

Clearly, this operator is neither commutative nor has an inverse.

3.3 Associativity

As we have already seen above, image registration is a non-convex optimization problem with multiple local minima. Thus, it may seem that the corresponding prefix scan operator is not associative. The special precautions we had to take for our specific setting with periodic structures are the key to get associativity in practice. This is due to our assumption that the shift between two consecutive images is smaller than half of the period and the way we construct initial guesses for the deformation for non-consecutive images, which should ensure that we start the minimization sufficiently close to the global minimum. The registration process converges to correct results as long as deformations accumulate between adjacent images, ensuring that the shift between images is sufficiently small. Prefix scan preserves the guarantee, and thanks to the iterative optimization process, each operator application will converge to the best local solution even if changes in computation produce slightly different partial results.

The integrity of the data was verified with a manual inspection on small scale experiments, which included examples where various deformations provide equally suitable matches. To compare two frames, we locate points of interest such as groups of atoms, and verify that their absolute position inside frames is the same. An example of the manual verification process is presented in Figure 5.

A numerical comparison of cost function scores between sequential and parallel runs is not possible because the optimization process deals with a high level of noise in the input data and there, a different score does not necessarily indicate a worse or better match.

3.4 Computation Cost

The simplest case of a prefix scan operator found in the literature, which happens to be the one most frequently evaluated, is integer addition. More complex examples still involve relatively cheap operations, such as polynomial evaluation and addition of summed area tables with multiple integer and floating-point multiplications. As a result, parallel prefix scan algorithms tend to be optimized for memory-bound operators with a low execution time. Image registration does not fit into this category, as it can be seen in Figure 6a. A single operator application usually takes up to 10 seconds, with noticeable outliers going for up to 30 seconds. The resulting deformation stores only three floating-point values and the cost of sharing such data is dominated by the latency. The computation time is much larger than latency introduced by network communication and this discrepancy will not change even with significant serial optimizations of the operator.

3.5 Load Imbalance

Another particular feature of prefix scan operators that is commonly seen in the literature, is a deterministic execution time that does not change between applications. In contrast to operations with a predictable and constant runtime, here the actual computation cost is not only unpredictable but highly variant. Due to the iterative nature of the registration algorithm, we can not foresee for a given input data how many iterations are necessary to reach a stopping criterion. The time measurements presented in Figure 6a show that significant outliers do not form any regular distribution and estimation of an efficient distribution is not possible. For the same dataset, we studied the load imbalance of a static data distribution to learn how the distributed run might be affected when the increase in computing resources leads to smaller data chunks available to each rank. We look at the difference between mean and maximum execution time across data segments. Intuitively, if the imbalance of computational effort between segments is large, then the larger is the difference between mean completion time and the slowest worker. The results in Figure 6b show how the increase in execution time raises from roughly 5% for large data segments to over 20% when each segment contains less than 100 deformations. These results indicate how speedups of our parallel image registration are going to change when we scale the problem to the point where only a few dozens of images are available per MPI rank. The performance is going to decrease not only because of the raising cost of a global scan but also due to increasing influence of load imbalance. Since we want our configuration to be located on the part of the plot with a low imbalance factor, we have to choose a sufficiently large segment size. To that end, we present a hierarchical decomposition of prefix scan in Section 4.2 to group parallel workers and increasing the segment size on certain levels of hierarchy.

4 DISTRIBUTED SCAN STRATEGIES

Prefix scan has been successfully applied in distributed and accelerated computations and many of these attempts rediscovered the same standard strategies, *scan-then-map* and *reduce-then-scan*, when data size significantly exceeds the number of workers (Section 4.1). We show how these strategies can be extended to a hierarchy of parallel workers (Section 4.2), such as the one used in a hybrid computation with multiple threads per each distributed worker, with neither a loss of generality nor an increase of algorithm depth. We exploit the additional level of a shared-memory parallelization (Section 4.3) to provide a load balancing step, improving performance of scan operators with an unbalanced and unpredictable execution time, such as the one introduced in the previous section.

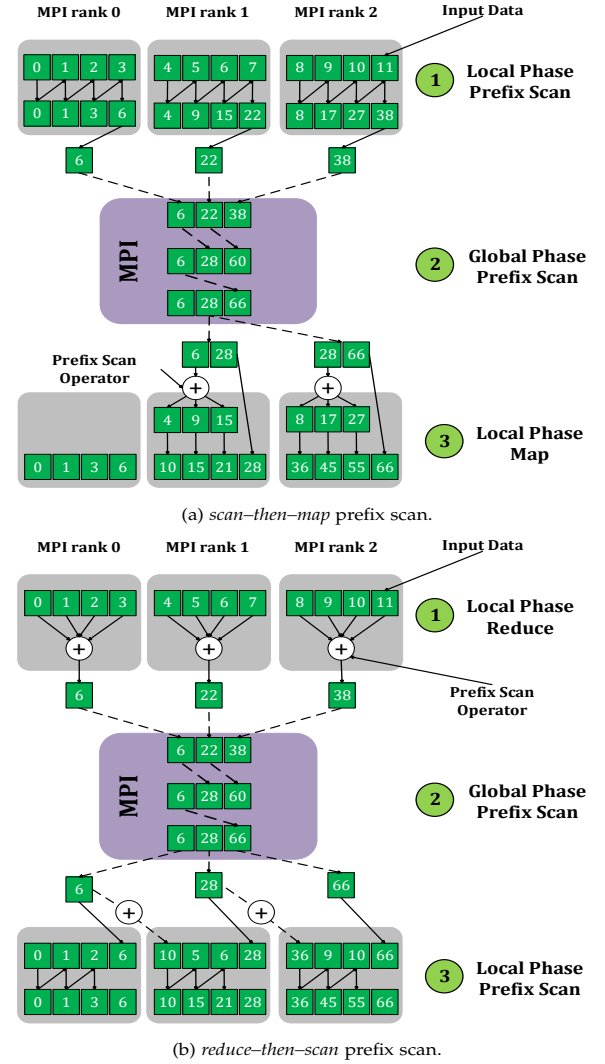


Fig. 7: Examples of distributed prefix scan strategies for integer addition on three MPI ranks. Operations local to MPI rank are shown with \rightarrow whereas \dashrightarrow indicates global communication.

4.1 Distributed Prefix Scan

The classical parallel prefix scan algorithms were designed to minimize the depth when the number of processing

elements is equal to the number of data elements. Although such case is common in circuit design, it is not well-suited as a general solution since the length of input sequence x_0, x_1, \dots, x_{N-1} is usually larger than the number of workers P . For simplicity, we limit the analysis to the case of even data distribution, and each worker is assigned $K = \frac{N}{P}$ input elements with boundary indices l_I and r_I .

The main strategy for a distributed scan follows a principle of splitting the work to local-global-local sequence of computations, as presented in Figure 7. First, each rank is assigned a data segment to process independently in the local phase ①, computing a sum of all elements in the local segment x_{l_I, r_I} . The result is passed to a global prefix scan of size P , computing an accumulated result x_{0, r_I} on each rank ②. The computation is finalized with an update of local data segments with an accumulated sum of x_{0, r_I-1} ③. As global phase ②, one can use any distributed scan implementation, such as `MPI_Scan`. The local phases ① and ③ can be defined in two ways, either as a scan that updates local data and requires only adding global result in the end or as a reduction that leaves the segment intact and finishes the computation with a prefix scan. The *scan-then-map* procedure is usually preferred over the *reduce-then-scan* approach since the former exhibits a slightly lower depth and decreased workload due to the first parallel worker inactivity in last phase. We use the former approach for evaluation of a standard, distributed prefix scan. While in the *scan-then-map* algorithm the work distribution has to be defined before execution, a dynamic determination of workload per thread is possible when the first phase is a reduction. Such property is desirable for imbalanced computations to allow work-stealing and decrease disproportions in workload. In next two paragraphs we discuss advantages and disadvantages of each approach in detail.

4.1.1 Scan-then-map

In this approach, presented in Figure 7a, a *scan* is computed initially in the first phase LP_1 , creating a new sequence of partial results that requires only an *application* of global scan result in the last phase. Depth of the first phase is straightforward: $\frac{N}{P} - 1$.

The last element computed by the local scan x_{l_I, r_I} is the sum needed for a global scan. In the second local phase, each local result $x_{l_I, j}$ is combined with the exclusive value x_{0, l_I-1} . An exception is the data segment assigned to worker 0 which is already finished. This requires exactly K applications of the operator. However, since the prefix scan is inclusive, the last value x_{0, r_I} is already computed in the global phase which saves one application of the operator and requires $\frac{N}{P} - 1$ steps.

The depth of the algorithm is given as follows:

$$D_{DS}(N, P) = 2\frac{N}{P} - 2 + D_{GS}(N, P) \quad (1)$$

The analysis of critical path is possible for an even distribution of data. Otherwise, critical paths of local phases might be provided by different workers and simple summation would yield an incorrect result. The last phase can be parallel and balanced since each element is updated independently. Yet, there's no possibility to decrease load imbalance before global phase.

4.1.2 Reduce-then-scan

As depicted in Figure 7b, each worker computes sequentially a *reduction* in the first phase, leaving local data elements untouched until the *scan* in the second local phase. There, the global result x_{0, l_I-1} is added to the first local element x_{l_I} and the scan updates each value with x_{0, l_I-1} . For the first phase, workload and depth does not change since the very first element x_{l_I} can be used as an initial value of the sum. There's a difference in last phase, however. Although we can still use the trick with inclusive result, one needs first to apply the global result to the first element and the first worker is no longer idle in that phase, requiring $\frac{N}{P}$ operator applications. The depth of the algorithm is given as follows:

$$D_{DS}(N, P) = 2\frac{N}{P} - 1 + D_{GS}(N, P) \quad (2)$$

Contrary to the other approach, here the first phase allows for further parallelization due to less strict nature of reduction. The strictly sequential last phase is a minor disadvantage.

4.2 Hierarchical Prefix Scan

We now present the strategy for a distributed scan that includes a hierarchical distribution of work and data, similar to the decomposition techniques applied for prefix scans on GPU devices [20]. We show that such redistribution can be performed with a constant increase in the algorithmic depth in the worst case. Even though we do not achieve reduction in depth, the hierarchization decreases the number of ranks participating in the global scan. This change reduces negative performance effects of an unbalanced global scan on many ranks and decreases the pressure and dependence on network communication by performing more computation intra-node. Moreover, applying the hierarchical scan to distributed computation introduces a lower hierarchy layer with shared-memory environment that allows for an efficient implementation of work stealing, as discussed in Section 4.3. Although we consider here the most common case of a hybrid MPI application with local threads assigned to each rank, the general principle extends to an arbitrary number of levels.

For a multithreaded implementation, we assume that an allocation of P MPI ranks is replaced with P' ranks and T threads such that $P' \cdot T = P$.

- 1) Local Phase on $P' \cdot T$ workers.
For both scan and reduce, there is no change in either depth or work performed since each segment of size $\frac{N}{P}$ is replaced with a new one of length $\frac{N}{P' \cdot T}$.
- 2) Local Scan on T local segments.
We assume that internally each rank uses the same parallel prefix scan algorithm as in the global scan, with $D' = C_1 \log_2 T + C_2$.
- 3) Global Scan on P' ranks.
In each scan iteration, the result received from other rank is applied by T threads to T scan results corresponding to inclusive prefix scan over all segments. Only the last result is used for communication.

4) Second Local Phase on $P' \cdot T$ workers.

The entire computation proceeds without major changes since each thread owns the scan result continuously updated in the global phase.

We observe that the composition of local and global parallel scans does not change the asymptotic performance since $C_1 \log_2 T + C_1 \log_2 P' = C_1 \log_2 P$ and only change is visible in constants C_2 . Nevertheless, this increase in depth does not apply to depth-optimal scans, where $C_2 = 0$, which are of special interest.

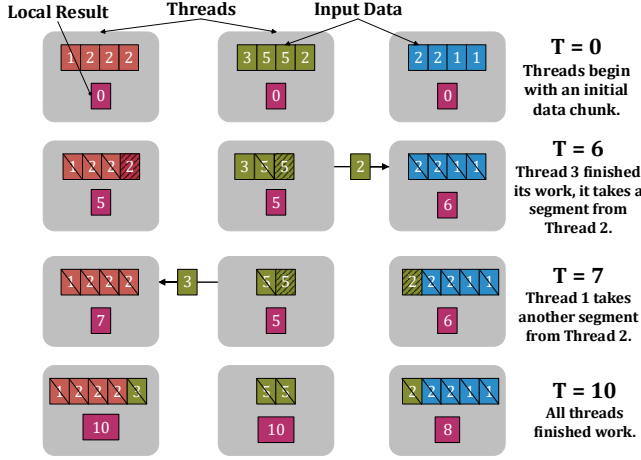


Fig. 8: An example of work-stealing in the computing reduction over three data segments. Data values indicate the computation time. Compared to the static data distribution, where the result of middle segment arrives after $t = 15$ units of time, the computation is more balanced due to additional redistribution and the global phase is started earlier.

4.3 Dynamic Hierarchical Prefix Sum

We now move away from one of the core assumptions that has been always made for a prefix scan operator - the computational cost is constant and easily predictable. Although this assumption is valid in many applications, it does not always hold as it is the case of image registration. If a static data distribution does not provide a balanced workload, the performance of the entire application is affected: not only it will take more time to process a single segment but the disparity will be later propagated due to a synchronous nature of the prefix scan. The ideally balanced data distribution can be estimated if each operator cost is known a priori which is not feasible for iterative computations with data-dependent stopping conditions. A fine-grain work distribution is inefficient due to the increased depth of the global scan phase.

Therefore, we focus on techniques that can react to ongoing changes in workload balance. Well-known load balancing solutions do not apply directly to prefix scan due to the sequential nature of the scan computation and its limited ability to redistribute work chunks. We use the hierarchical prefix scan representation and introduce the shared-memory parallelism with threads. We attempt to detect when a certain thread is processing its workload

faster and let it steal work from its neighbors to balance the computation effort. Thus, we improve the performance of global scan, a main bottleneck in large scale computations, by (1) decreasing the work imbalance in first phase of computation and (2) restricting the global phase only to parallel workers on the highest hierarchy levels, in our case MPI ranks. The *scan-then-map* approach imposes a strict evaluation order from left to right since otherwise the new sequence would not contain a correct prefix scan. Fortunately, this requirement does not exist in the *reduce-then-scan* strategy where the first phase computes only a sum of the entire segment x_{l_I, r_I} . Given the associativity of the operator, there shall be no change in result if elements are processed from left to the right as in a prefix scan, from right to the left or from the middle of data segment in both directions. This observations allows us to consider flexible segment boundaries.

An example of such problem is presented in Figure 8. The static data distribution leads to an unbalanced workload and effectively slows down the prefix scan to the slowest thread. By changing the order of evaluation to left-to-right in the lowest-numbered segment, right-to-left in the highest-numbered segment, and to middle-outward for other segments, we leave an option for each thread to acquire more work in case its neighbor is processing slower. We note that there is no cost associated with changing segment boundaries since for the entire hierarchy it is only relevant that segment boundaries are aligned with each other. Any load balancing procedure will be restricted to exchanges between neighboring threads, due to the requirement that a sum must be computed across consecutive data elements. We focus on intra-node work-stealing due to diminishing returns of inter-node synchronization between logically adjacent threads. The main point of an efficient heuristic is to decide in which direction to accumulate data after starting on the middle element. Since the imbalance between neighbors cannot be predicted, the most sensible option is a greedy approach where threads always move in the direction of whichever adjacent thread is slower. Let pl_I and pr_I be the boundaries of processed elements for thread I . For each neighbor, we define the processing rate $t_{I\pm 1}$ as the ratio of computation time to the number of operator applications. Let s_I correspond to the number of data elements left unprocessed between threads I and $I + 1$. The Algorithm 1 presents the heuristic. For simplicity, we omit the initial step where threads always move to the right.

5 EVALUATION

For evaluation we use two supercomputing system: the Piz Daint supercomputer and a local cluster with Ivy Bridge CPUs, summarized in Table 2. We use 12 and 20 threads per rank on Piz Daint and IvyBridge, respectively, without hyper-threading and with each thread pinned to a physical core. All prefix scan algorithms were implemented in C++ as a part of the quocmesh library [30]. The work-stealing implementation splits the work across OpenMP threads and performs a local scan over partial results with the dissemination pattern since its implementation is simpler than a Ladner-Fischer scan and the difference in work performed is negligible when only a dozen or so threads participate in

Algorithm 1 Load balancing on threads $1, \dots, T$

```

1: while  $s_{I-1} > 0 \vee s_{I+1} > 0$  do
2:   if  $s_{I-1} > 0 \wedge s_{I+1} > 0$  then
3:     if  $t_{I-1} > t_{I+1}$  then
4:        $d \leftarrow LEFT$ 
5:     else
6:        $d \leftarrow RIGHT$ 
7:     end if
8:   else
9:     if  $s_{I-1} > 0$  then  $d \leftarrow LEFT$  else  $d \leftarrow RIGHT$ 
10:  end if
11:  if  $d == LEFT$  then
12:     $pl_I \leftarrow pl_I - 1, res_I \leftarrow x_{pl_I} \odot res_I$ 
13:  else
14:     $pr_I \leftarrow pr_I + 1, res_I \leftarrow res_I \odot x_{pr_I}$ 
15:  end if
16: end while

```

the scan. Images are available to all ranks through the high-performance filesystem and the communication is limited to 20 bytes of deformation data and indices. Algorithms use point-to-point communication with the exception of the Ladner-Fischer that uses `MPI_Broadcast` in certain iterations. For each experiment, measurements were repeated five times and we show on plots mean value with 95% confidence interval.

As image series, we consider data from an experiment where ultrahigh vacuum high-resolution TEM (UVH HRTEM) has been applied to capture the process of aluminum oxidation [31]. The images have been acquired at a resolution of $1,920 \times 1,856$ and a rate of 400 frames per second. An example of a single frame is presented in the Figure 3a.

	Piz Daint	Ivy Bridge
CPU	Intel Xeon E5-2690 CPU 2.60GHz	Intel Xeon E5-2680 v2 2.80GHz
Cores	12 with 12 hardware threads	20 with 20 hardware threads
Memory	64 GB	64 GB
Interconnect	Cray Aries, Dragonfly	FDR Infiniband
Build	CMake 3.5.2, GCC 7.3.0	CMake 3.6.0, GCC 8.2.0
MPI	Cray MPICH 7.7.2	IntelMPI 2018.3

TABLE 2: Evaluation systems: Piz Daint with Cray XC50 nodes and Ivy Bridge cluster with two deca-core CPUs.

5.1 Microbenchmarks

We first evaluate the prefix scan algorithms with a set of microbenchmarks. We use an artificial operator with (1) a static execution time, where each operator application takes the same amount of time and (2) a dynamic configuration where the execution time is a random variable and for each time t , we use an exponential distribution with rate $\lambda = \frac{1}{t}$ to obtain a similar average running time. We use `std::mt19937`, a 32-bit Mersenne Twister PRNG from the C++ standard library, with a constant seed 1410 to ensure reproducible results. We scale it up on Piz Daint with varying number of data elements per CPU core. Whenever we compare static and work-stealing implementations, both solutions use random number generators in the same deterministic fashion to ensure that the comparison is scientifically valid. We evaluate (1) the scalability of prefix scan algorithms

and (2) effectiveness of work-stealing on large and generic problems with an unbalanced workload.

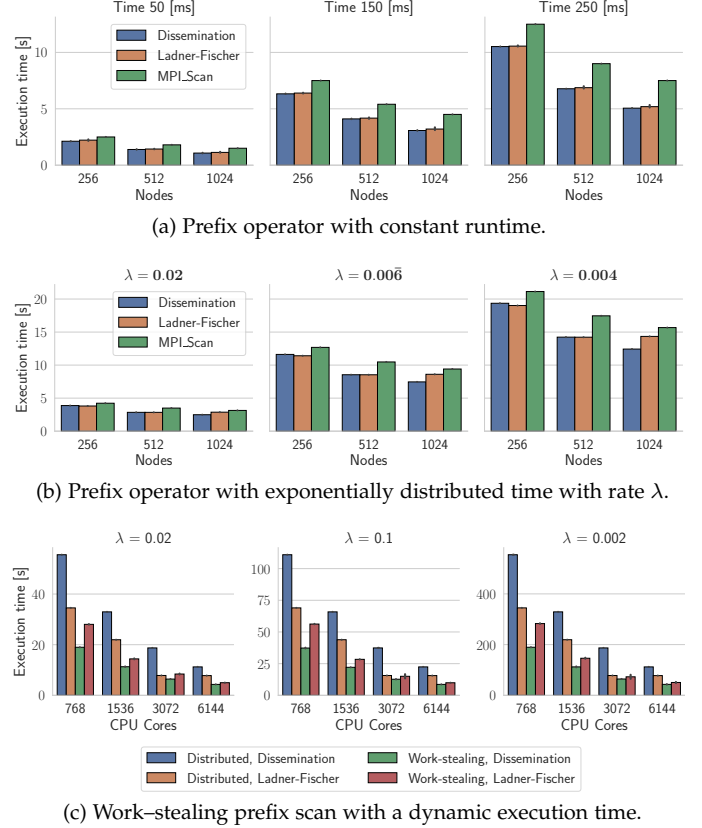


Fig. 9: Prefix scan algorithms on mock operators with a static and dynamic running time.

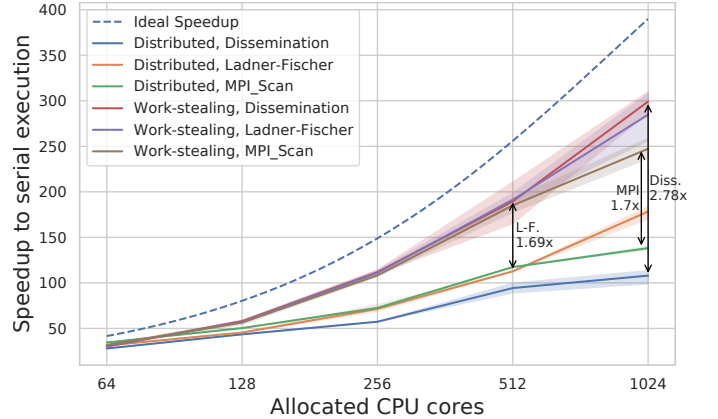


Fig. 10: The strong scaling of *full* image registration. Experimental results are indicated with solid lines whereas theoretical bound (4) discussed in Section 5.2 is dashed. Experiments conducted on Piz Daint for 4,096 images.

5.1.1 Inter-Node Scan

We test sensitivity of scan algorithms to unbalanced workloads and network congestion. We estimate the performance loss on different implementations of global scan by using a static, hierarchical prefix scan with one MPI rank and 12 threads per node. Figures 9a, 9b present results for static

and dynamic execution with 98304 data elements. Results show that scan algorithms perform differently on an ideally constant workload if the computation time plays a more important role than communication. Not surprisingly, the `MPI_Scan` performs worse than other prefix scan algorithms since it might be optimized for communication latency. Adding a controlled imbalance causes a performance drop and all prefix scan algorithms take on average twice more time. We can expect such slowdown in the image registration and other problems where load balance is an issue.

5.1.2 Work-stealing Scan

We evaluate the impact of our work-stealing on a generic prefix scan problem with an unbalanced workload. Figure 9c presents results for dynamic execution with 98304 data elements. Results show that our work-stealing provides substantial improvements when applied with the Ladner-Fischer scan whereas the performance of scan with the dissemination can be improved up to three times. The result is explained by chains of dependencies in task graphs of both algorithms. Dissemination and Ladner-Fischer scans represent distinct sequences of parallel computations, and the critical path is different in both algorithms as well, resulting in different impact of workload imbalance. The performance seems to be consistent across operators with varying execution time, and we see that work-stealing prefix scan improves the performance on a larger number of cores even if the distributed version stops to scale, as it is in the case of Ladner-Fischer from 3072 to 6144 cores.

5.2 Strong Scaling

We evaluate the strong scaling of the image registration on the Piz Daint system. As a baseline we choose the serial execution of a prefix scan that requires $N - 1$ operator applications for N deformations $\phi_{i,j}$. On a single core, this step takes on average 18422 seconds. Combined with the depth of a distributed prefix scan (Section 4.1, Eq. (2)), we obtain the upper performance bound for *scan registration*

$$\frac{N - 1}{D_{DS}(N, P)} = \frac{N - 1}{2 \cdot \frac{N}{P} - 1 + C_1 \log_2 P} \quad (3)$$

Figure 1 and Table 3 present parallel speedups and efficiency. In addition, in the Figure we compare with the upper performance bound although this can be achieved only on perfectly balanced workloads. For the *scan* computation, we observe that our work-stealing prefix scan up to 1.98x, 1.83x and 1.51x times over the dissemination, `MPI_Scan` and the Ladner-Fischer scan, respectively. Applying our load balancing brings the performance closer to the upper bound and prevents the `MPI_Scan` algorithm from stopping to scale on 1024 cores. The work-efficient Ladner-Fischer exhibits a substantial performance improvement on 1024 cores. This is explained by the fact that scan algorithms are differently affected by various workload imbalances and the scan hits the sweet-spot for this configuration. We don't observe any significant performance improvements over 512 cores on dissemination and `MPI_Scan` whereas the work-stealing version provides further improvements. In turn, the

dynamic prefix scan allows further scaling of long image series registration.

Minor slowdowns observed for some algorithms on 64 or 128 ranks can be an effect of measurement noise that is noticeably larger for the dynamic execution. At the same time, we observe significant performance improvements from 512 cores onwards. Both results align with the analysis presented in Section 3.5 where it has been shown that negative effects of an imbalanced operator have the highest impact when the local segment size is small. When the number of allocated cores is relatively small, and the data segment is in turn large, applying work-stealing on lower levels of hierarchical prefix scan cannot prevent all effects of unbalanced workload on the highest level, which is the entire node in our setup. Furthermore, in such setup, the computation is not dominated by the global scan which is especially sensitive to imbalanced workloads.

In addition, we study the performance of a *full registration* that includes the initial step of generating input deformations for the prefix scan (Section 2.3, function A). A full serial registration requires on average 37567 seconds of computation. The upper performance bound Eq. (3) is changed by adding N initial registration steps that are massively parallel, adding depth $\frac{N}{P}$ on P processes.

$$\frac{N + N - 1}{\frac{N}{P} + D_{DS}(N, P)} = \frac{2N - 1}{3 \cdot \frac{N}{P} - 1 + C_1 \log_2 P} \quad (4)$$

Figure 10 and Table 3 present results for the *full registration*. We observe that our work-stealing prefix scan improves the performance up to 2.78x, 1.7x and 1.69x times over the dissemination, `MPI_Scan` and the Ladner-Fischer scan, respectively. Similar to the *scan* registration, the work-stealing scan prevents stopping to scale over 512 cores when using dissemination and `MPI_Scan`. We observe a substantial improvement on dissemination prefix scan. Although more work is performed in first phase, the time spent in global scan phase decreases due to lower waiting times. Such result is not surprising since different scan algorithms are affected differently by load imbalance. Furthermore, our work-stealing performs better when more work is performed by each thread.

5.3 Hierarchical Prefix Scan

We study the performance effects of the introduction of a hierarchy of parallel workers, replacing the standard set-up of P MPI ranks with P' ranks equipped with T threads each, such that $P' \cdot T = P$ and present in Table 4 a comparison against serial execution and the distributed execution. We observe a significant difference in performance for both the dissemination and `MPI_Scan`. On the other hand, we observe performance degradation with the Ladner-Fischer in the sweet spot on 1024 cores. Hierarchical prefix scan leads to performance improvements thanks to a decreased cost of the inter-node synchronization in global phase, except of a single outlier with Ladner-Fischer scan on 1024 cores. There, the efficiency of a pure MPI solution suddenly increases to the point where it outperforms the dissemination prefix scan by over 60% and introducing hierarchy leads to performance degradation. We conclude that this specific setup is a sweetspot for the Ladner-Fischer scan where the

Cores	Scan Registration																	
	Dissemination			Distributed Ladner-Fischer			MPI_Scan			Dissemination			Work-stealing Ladner-Fischer			MPI_Scan		
	Time	S	\mathcal{E}	Time	S	\mathcal{E}	Time	S	\mathcal{E}	T	S	\mathcal{E}	Time	S	\mathcal{E}	Time	S	\mathcal{E}
64	979.26	18.81	0.29	765.61	24.06	0.38	779.82	23.62	0.37	780.89	23.59	0.37	805.37	22.87	0.36	800.13	23.02	0.36
128	677.38	27.20	0.21	404.00	45.60	0.36	552.22	33.36	0.26	455.43	40.45	0.32	473.43	38.91	0.30	466.09	39.53	0.31
256	550.90	33.44	0.13	338.42	54.44	0.21	398.26	46.26	0.18	286.52	64.30	0.25	260.95	70.60	0.28	288.25	63.91	0.25
512	340.76	54.06	0.11	283.56	64.97	0.13	298.32	61.75	0.12	171.60	107.35	0.21	186.88	98.58	0.19	193.11	95.40	0.19
1024	274.45	67.12	0.07	166.11	110.91	0.11	308.27	59.76	0.06	143.58	128.31	0.13	153.06	120.36	0.12	167.94	109.70	0.11

Cores	Full Registration																	
	Dissemination			Distributed Ladner-Fischer			MPI_Scan			Dissemination			Work-stealing Ladner-Fischer			MPI_Scan		
	Time	S	\mathcal{E}	Time	S	\mathcal{E}	Time	S	\mathcal{E}	Time	S	\mathcal{E}	Time	S	\mathcal{E}	Time	S	\mathcal{E}
64	1,345.51	27.92	0.44	1,203.96	31.20	0.49	1,088.50	34.51	0.54	1,247.10	30.12	0.47	1,222.26	30.74	0.48	1,179.33	31.86	0.50
128	864.01	43.48	0.34	826.54	45.45	0.36	745.69	50.38	0.39	648.47	57.93	0.45	657.33	57.15	0.45	666.51	56.36	0.44
256	655.41	57.32	0.22	529.13	71.00	0.28	518.36	72.47	0.28	336.08	111.78	0.44	338.54	110.97	0.43	347.25	108.19	0.42
512	400.27	93.86	0.18	333.43	112.67	0.22	320.10	117.36	0.23	202.59	185.43	0.36	196.81	190.88	0.37	203.36	184.74	0.36
1024	350.44	107.20	0.10	211.03	178.02	0.17	271.90	138.17	0.13	125.84	298.54	0.29	133.55	281.30	0.27	152.30	246.68	0.24

TABLE 3: Execution times, parallel speedups S and efficiency \mathcal{E} for (a) the standard, MPI-only distributed prefix scan, (b) ours hierarchical prefix scan with MPI ranks and work-stealing on OpenMP threads. Speedups are computed relative to the serial *scan* and *full* registration lasting 18422.17 and 37567.7, respectively.

Cores	Hierarchical Scan Registration								
	Dissemination			Distributed Ladner-Fischer			MPI_Scan		
	Time	S	S'	Time	S	S'	Time	S	S'
64	683.43	26.96	1.43	682.80	26.98	1.12	729.21	25.26	1.07
128	403.26	45.68	1.68	402.62	45.76	1.00	437.83	42.08	1.26
256	274.51	67.11	2.01	329.98	55.83	1.03	300.60	61.29	1.32
512	202.39	91.02	1.68	244.73	75.28	1.16	243.27	75.73	1.23
1024	162.53	113.35	1.69	215.74	85.39	0.77	175.55	104.94	1.76

TABLE 4: Execution times, parallel speedups S' and S'' with respect to serial and distributed execution for hierarchical prefix scan without work-stealing.

Cores	Distributed					
	Dissemination			Ladner-Fischer		
	Corehours	Energy [MJ]		Corehours	Energy [MJ]	
64	2.29x	1.34±0.12	1.89x	2.05x	1.18±0.14	1.66x
128	2.94x	1.76±0.1	2.49x	2.82x	1.16±0.13	1.64x
256	4.67x	2.34±0.19	3.3x	3.61x	1.74±0.13	2.45x
512	5.46x	3.3±0.36	4.65x	4.54x	2.85±0.13	4.01x
1024	9.55x	4.92±0.21	6.94x	5.75x	3.5±0.22	4.94x

Cores	Work-stealing					
	Dissemination			Ladner-Fischer		
	Corehours	Energy [MJ]		Corehours	Energy [MJ]	
64	1.98x	1.09±0.04	1.53x	1.99x	1.11±0.02	1.57x
128	2.18x	1.17±0.05	1.64x	2.2x	1.19±0.03	1.68x
256	2.5x	1.29±0.06	1.81x	2.67x	1.34±0.13	1.89x
512	3.88x	1.54±0.06	2.17x	2.93x	1.48±0.16	2.08x
1024	7.1x	2.21±0.1	3.12x	4.69x	2.28±0.15	3.22x

TABLE 5: The increase in computing time and energy cost for a *full* registration of 4,096 images on Piz Daint. Results are presented with sample standard deviation and compared against serial execution with 10.43 corehours and 0.71 MJ consumed.

global scan phase performs very well without significant delays.

5.4 Work and Energy

A prefix scan has to compensate for a reduction in depth by an increase in work. When combined with the load imbalance that further decreases parallel efficiency, the question has to be asked: how many resources are required to reduce the time of processing microscopy images? Precisely, we want to find answers to two research questions: (1) by how much does the imbalanced parallelization increase resource consumption, (2) does the work-stealing scan provide energy and compute time reduction in addition to improved performance?

To find answers, we design an experiment evaluating the resource and energy consumption of the *full* image registration. To obtain energy measurements on the Piz Daint, we use the Cray Resource Utilization Reporting (RUR) [32], a reliable source of job-wide energy measurements on Cray supercomputers [33]. The tool reports aggregated measurements of hardware energy counters for the entire node, including the node's energy consumption for the batch job duration (CPU cores, memory, other hardware) [34]. It does not include the interconnect shared between nodes, which is not a limitation because computation costs dominate communication costs in our problem. As a baseline, we select the energy consumption of the sequential registration procedure, similarly to the previous strong scaling experiment. Because the Cray RUR supports node-wide measurements only, we cannot use the value reported for the serial procedure straight away as it would include the energy consumption for the entire node where one core is allocated, and eleven are idle. Instead, we execute the same sequential registration on each of the twelve cores, we use the measured value as a summarized cost of twelve repetitions of the same problem and compute the average to represent the total energy cost of the serial version.

We measure the hardware resource consumption as the time required to obtain the result multiplied by the number of CPU cores allocated (corehours). As a baseline, we use the sequential execution with 8191 operator applications and 10.43 corehours consumed. Results based on the strong scaling experiment data are presented in Table 5.

Load Imbalance The presence of an imbalance leads to a significantly higher cost of parallel computations, with up to 9.55x and 5.75x more CPU allocation needed to finish the computations with the dissemination and Ladner-Fischer scans, respectively. The energy consumption increases by a factor of 6.94x and 4.94x, respectively, suggesting that a significant fraction of CPU time is spent on idle waiting that consumes little energy.

Work-stealing Applying our hierarchical and dynamic parallelization decreases the energy consumption up to 2.23x and 1.93x times when using the dissemination and Ladner-Fischer prefix scan, respectively. Additionally, our prefix scan requires up to 1.87x and 1.55x times less computing time. The decrease in energy consumption is higher than for hardware allocation time, as the work-stealing scan

still includes idle waiting times that cannot be eradicated entirely due to coarse-grained work chunks of image registration.

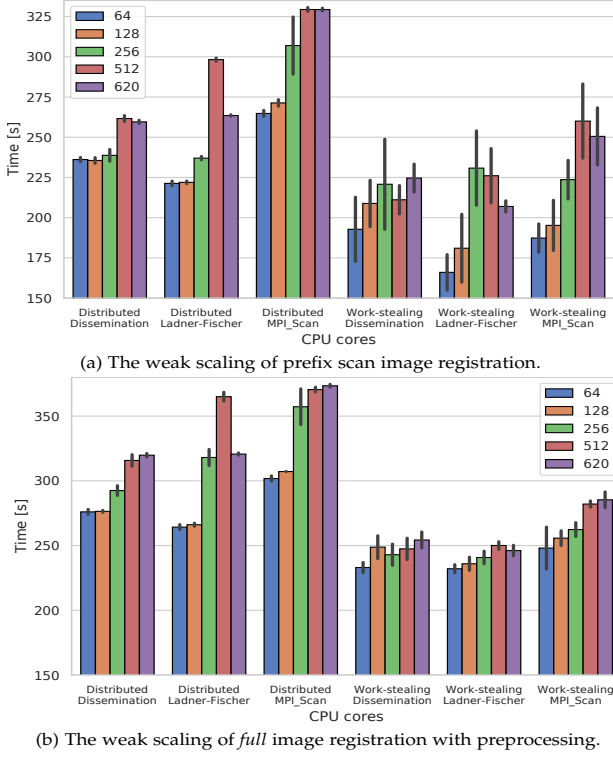


Fig. 11: The weak scaling of image registration on Ivy Bridge. Figures begin at value 150 for better readability.

5.5 Weak Scaling

An efficient weak scaling is an important goal since it allows to utilize additional hardware to process longer sequences of microscopy frames while keeping the analysis time practical. When increasing compute resources by the same factor as the problem size, the execution time cannot stay constant due to the logarithmic factor associated with prefix scan: $D_{DS}(k \cdot N, k \cdot P) = D_{DS}(N, P) + C_1 \log_2 k$.

We analyze the weak scaling for the scenario of 8 images per rank and scale it from 64 to 620 ranks on the Ivy Bridge system. Figures 11a and 11b present results for the prefix scan and the full registration procedure, respectively. While we observe an increased execution time for static prefix scan algorithms, the work-stealing procedure helps to mitigate the effects of a logarithmic increase in depth. This benefit is especially visible in the full registration, where for both dissemination and Ladner–Fischer prefix scan the dynamically balanced version scales to a larger number of cores with a minor change in execution time.

6 CONCLUSIONS

In this paper, we proposed a prefix scan parallelization strategy for the registration of a long series of electron microscopy images. This work is first to consider a prefix scan problem where the optimization focus moves from communication latency to a computationally expensive and

highly load-imbalanced operator. To overcome scaling difficulties and slowdowns of prefix scans on imbalanced computations, we apply the work-efficient Ladner–Fischer prefix scan and provide a novel node-local work-stealing procedure that can be applied to any prefix scan parallel computation. We show that work-stealing improves the performance of prefix scan on imbalanced workloads up to $2.1x$ times while decreasing CPU allocations and energy consumption up to $1.87x$ and $2.23x$ times, respectively. As a result, an analysis of arbitrarily long microscopy series is now possible thanks to a dynamic prefix scan strategy that keeps scaling with increasing hardware resources. We identified a performance gap in the MPI collectives, and with our dynamic and work-stealing scan algorithm, we provide a generic solution that enables efficient distributed parallelization of imbalanced scan computations.

ACKNOWLEDGMENTS

The authors would like to thank Professor Sarah Haigh, University of Manchester, for providing the TEM data. We would also like to acknowledge the following organizations for providing us with access to their supercomputers: the Swiss National Supercomputing Centre (CSCS), and the Aachen Institute for Advanced Study in Computational Engineering Science (AICES). This work was in parts supported by the Schweizerische Nationalfonds zur Förderung der wissenschaftlichen Forschung (SNF, Swiss National Science Foundation) through Project 170415 and programme Ambizione (grant PZ00P2168016). B. Berkels was funded in part by the Excellence Initiative of the German Federal and State Governments through grant GSC 111.

REFERENCES

- [1] S. Chatterjee, G. E. Blelloch, and M. Zagha, “Scan primitives for vector computers,” in *Proceedings of the 1990 ACM/IEEE Conference on Supercomputing*, ser. Supercomputing ’90. IEEE Computer Society Press, 1990, pp. 666–675.
- [2] G. E. Blelloch, “Prefix sums and their applications,” School of Computer Science, Carnegie Mellon University, Tech. Rep. CMU-CS-90-190, Nov. 1990.
- [3] T. H. Cormen, C. Stein, R. L. Rivest, and C. E. Leiserson, *Introduction to Algorithms*, 2nd ed. McGraw-Hill Higher Education, 2001.
- [4] G. E. Blelloch, *Vector Models for Data-parallel Computing*. Cambridge, MA, USA: MIT Press, 1990.
- [5] (2019) Piz daint. [Online]. Available: <https://www.cscs.ch/computers/piz-daint/>
- [6] B. Berkels, P. Binev, D. A. Blom, W. Dahmen, R. C. Sharpley, and T. Vogt, “Optimized imaging using non-rigid registration,” *Ultramicroscopy*, vol. 138, pp. 46 – 56, 2014.
- [7] M. Snir, “Depth-size trade-offs for parallel prefix computation,” *J. Algorithms*, vol. 7, no. 2, pp. 185–201, Jun. 1986.
- [8] H. Zhu, C.-K. Cheng, and R. Graham, “On the construction of zero-deficiency parallel prefix circuits with minimum depth,” *ACM Trans. Des. Autom. Electron. Syst.*, vol. 11, no. 2, pp. 387–409, Apr. 2006.
- [9] G. E. Blelloch, “Scans as primitive parallel operations,” *IEEE Trans. Comput.*, vol. 38, no. 11, pp. 1526–1538, Nov. 1989.
- [10] R. P. Brent and H. T. Kung, “The chip complexity of binary arithmetic,” in *Proceedings of the Twelfth Annual ACM Symposium on Theory of Computing*, ser. STOC ’80. New York, NY, USA: ACM, 1980, pp. 190–200.
- [11] Ömer Eğecioğlu, C. K. Koc, and A. J. Laub, “A recursive doubling algorithm for solution of tridiagonal systems on hypercube multiprocessors,” *Journal of Computational and Applied Mathematics*, vol. 27, no. 1, pp. 95 – 108, 1989, special Issue on Parallel Algorithms for Numerical Linear Algebra.

- [12] P. M. Kogge and H. S. Stone, "A parallel algorithm for the efficient solution of a general class of recurrence equations," *IEEE Trans. Comput.*, vol. 22, no. 8, pp. 786–793, Aug. 1973.
- [13] W. D. Hillis and G. L. Steele, Jr., "Data parallel algorithms," *Commun. ACM*, vol. 29, no. 12, pp. 1170–1183, Dec. 1986.
- [14] R. E. Ladner and M. J. Fischer, "Parallel prefix computation," *J. ACM*, vol. 27, no. 4, pp. 831–838, Oct. 1980.
- [15] M. P. I. Forum, "MPI: A Message-Passing Interface Standard Version 3.0," 09 2012.
- [16] P. Sanders and J. L. Träff, *Parallel Prefix (Scan) Algorithms for MPI*. Springer Berlin Heidelberg, 2006, pp. 49–57.
- [17] C. P. Kruskal, L. Rudolph, and M. Snir, "The power of parallel prefix," *IEEE Transactions on Computers*, vol. C-34, no. 10, pp. 965–968, Oct 1985.
- [18] H. Meijer and S. G. Akl, "Optimal computation of prefix sums on a binary tree of processors," *International Journal of Parallel Programming*, vol. 16, no. 2, pp. 127–136, Apr 1987.
- [19] Ömer Eğecioğlu and Çetin Kaya Koç, "Parallel prefix computation with few processors," *Computers & Mathematics with Applications*, vol. 24, no. 4, pp. 77 – 84, 1992.
- [20] S. Maleki, A. Yang, and M. Burtscher, "Higher-order and tuple-based massively-parallel prefix sums," in *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 539–552. [Online]. Available: <https://doi.org/10.1145/2908080.2908089>
- [21] P. Sanders, J. Speck, and J. L. Träff, *Full Bandwidth Broadcast, Reduction and Scan with Only Two Trees*. Springer Berlin Heidelberg, 2007, pp. 17–26.
- [22] P. Sanders, J. Speck, and J. L. Träff, "Two-tree algorithms for full bandwidth broadcast, reduction and scan," *Parallel Computing*, vol. 35, no. 12, pp. 581 – 594, 2009.
- [23] S. Maleki, M. Musuvathi, and T. Mytkowicz, "Parallelizing dynamic programming through rank convergence," *SIGPLAN Not.*, vol. 49, no. 8, pp. 219–232, Feb. 2014.
- [24] T. Gradl, A. Spörl, T. Huckle, S. J. Glaser, and T. Schulte-Herbrüggen, "Parallelising matrix operations on clusters for an optimal control-based quantum compiler," in *Euro-Par 2006 Parallel Processing*, W. E. Nagel, W. V. Walter, and W. Lehner, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 751–762.
- [25] K. Waldherr, T. Huckle, T. Auckenthaler, U. Sander, and T. Schulte-Herbrüggen, "Fast 3d block parallelisation for the matrix multiplication prefix problem," in *High Performance Computing in Science and Engineering, Garching/Munich 2009*, S. Wagner, M. Steinmetz, A. Bode, and M. M. Müller, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 39–50.
- [26] T. Auckenthaler, M. Bader, T. Huckle, A. Spörl, and K. Waldherr, "Matrix exponentials and parallel prefix computation in a quantum control problem," *Parallel Computing*, vol. 36, no. 5, pp. 359 – 369, 2010, parallel Matrix Algorithms and Applications.
- [27] S. Wang, Y. Bai, and G. Pekhimenko, "Scaling back-propagation by parallel scan algorithm," *CoRR*, vol. abs/1907.10134, 2019.
- [28] Y. Zou and S. Rajopadhye, "Scan detection and parallelization in "inherently sequential" nested loop programs," in *Proceedings of the Tenth International Symposium on Code Generation and Optimization*, ser. CGO '12. New York, NY, USA: Association for Computing Machinery, 2012, p. 74–83.
- [29] A. B. Yankovich, B. Berkels, W. Dahmen, P. Binev, S. I. Sanchez, S. A. Bradley, A. Li, I. Szlufarska, and P. M. Voyles, "Picometre-precision analysis of scanning transmission electron microscopy images of platinum nanocatalysts," *Nature Communications*, vol. 5, Jun. 2014.
- [30] (2019) Quocmesh library. [Online]. Available: <http://numod.ins.uni-bonn.de/software/quocmesh/>
- [31] L. Nguyen, T. Hashimoto, D. N. Zakharov, E. A. Stach, A. P. Rooney, B. Berkels, G. E. Thompson, S. J. Haigh, and T. L. Burnett, "Atomic-scale insights into the oxidation of aluminum," *ACS Applied Materials & Interfaces*, vol. 10, no. 3, pp. 2230–2235, 2018, pMID: 29319290.
- [32] A. Barry, "Resource utilization reporting," in *Proceedings of the Cray User Group Meeting 2013 (CUG 2013)*, May 2013.
- [33] A. J. Younge, R. E. Grant, J. H. Laros, M. Levenhagen, S. L. Olivier, K. Pedretti, and L. Ward, "Small scale to extreme: Methods for characterizing energy efficiency in supercomputing applications," *Sustainable Computing: Informatics and Systems*, vol. 21, pp. 90–102, 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2210537918301021>
- [34] A. Hart, H. Richardson, J. Doleschal, T. Ilsche, M. Bielert, and M. Kappel, "User-level power monitoring and application performance on cray xc30 supercomputers," *Proceedings of the Cray User Group (CUG)*, 01 2014.



Marcin Copik is a PhD student at ETH Zurich. His research interests include high-performance computing, serverless computing and performance modeling.



Tobias Grosser is an Associate Professor at University of Edinburgh, where he works in the Compiler and Architecture Design Group. His research interests are in compilation, programming language design, and effective performance programming.



Torsten Hoefler is a Professor at ETH Zurich, where he leads the Scalable Parallel Computing Lab. His research aims at understanding performance of parallel computing systems ranging from parallel computer architecture through parallel programming to parallel algorithms.



Paolo Bientinesi is a Professor in High-Performance Computing at Umeå University and the director of the High Performance Computing Center North (HPC2N). His research interests include numerical linear algebra, tensor operations, performance modelling & prediction, computer music, and the automatic generation of algorithms and code.



Benjamin Berkels is a Juniorprofessor for Mathematical Image and Signal Processing at AICES, RWTH Aachen. His research interests include variational and joint methods for image registration and segmentation, and medical image processing.