Parallel Prefix Algorithms for the Registration of Arbitrarily Long Electron Micrograph Series

Marcin Copik, Paolo Bientinesi (advisor), Benjamin Berkels (advisor)

Motivation

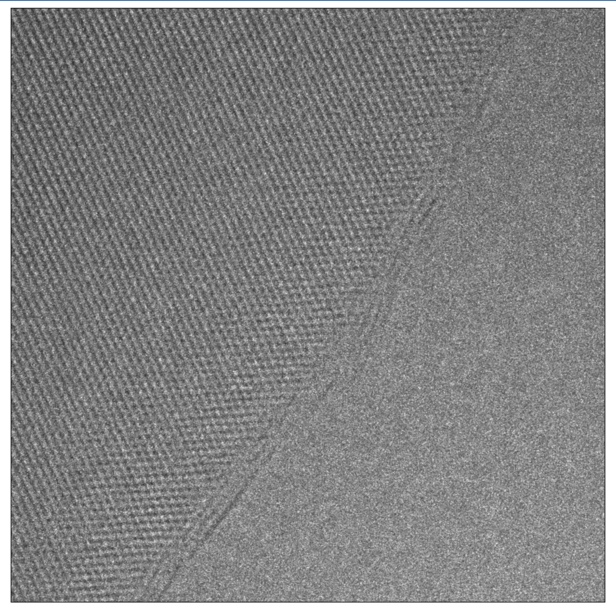

The process of aluminum oxidization, acquired with a transmission electron microscope. Courtesy of the University of Manchester.

- The interaction of electrons with the specimen is mapped to a display device
- The quality of images is limited by changes induced by electrons and movement of the sample
- Solve both problems by acquiring a series of low dose frames and processing it to extract information
- Image registration is computationally expensive - can we parallelize that process?

Image registration

Goal: find rigid deformations $\phi_{0,i}$ for frames $f_1, \ldots, f_n$

$$f_0 \approx f_i \circ \phi_{0,i}$$

Neighboring frames: apply registration function **A**

$$\phi_{i,i+1} = \mathbf{A}(f_i, f_{i+1})$$

Non–neighboring frames: use registration function **B**

$$\phi_{i,j} = \mathbf{B}(\phi_{i,k}, \phi_{k,j})$$

Prefix sum registration: with a binary operator $\odot_B$

$$\phi_{0,i} = \phi_{0,1} \odot_B \phi_{1,2} \odot_B \cdots \odot_B \phi_{i-1,i}$$

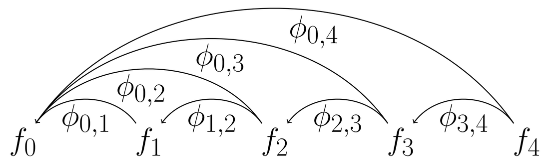

Parallel image registration

- A **trivially parallelizable** preprocessing step on neighboring pairs of frames

$$f_0, f_1, \ldots, f_n \rightarrow \phi_{0,1}, \phi_{1,2}, \ldots, \phi_{n-1,n}$$

- A prefix sum registers each frame to the first image, producing a new sequence

$$\phi_{0,1}, \phi_{0,2}, \ldots, \phi_{0,n}$$

a parallel prefix sum is necessary

Parallel prefix sum

- Multiple algorithms for a parallel prefix sum exist
- A decrease in span is achieved by performing more work
- Scan algorithms with $\log_2 N$ span are optimal, but they can not be work–efficient

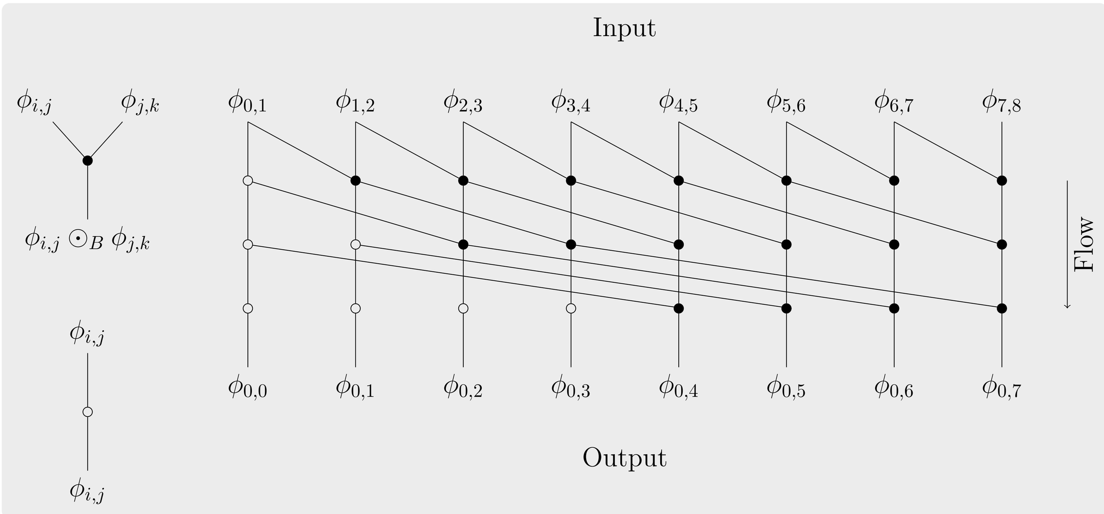

Kogge–Stone(Hills–Steele) parallel prefix sum

$S(N) = \log_2 N$

Image registration operator

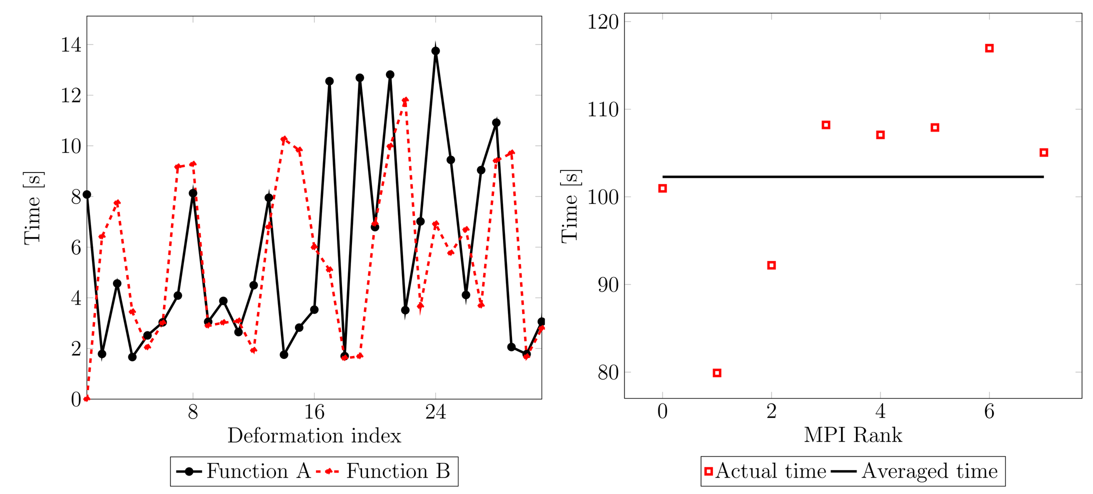

(1) Registration times on 64 images (2) Function B on 128 deformations on 8 MPI ranks

Existing parallel prefix sum algorithms

- Are simple and not computationally intensive
- Have a stable and deterministic execution time

In our problem, the image registration operator

- Is complex and computationally intensive (1)
- Has an unpredictable and unbalanced execution time (2)

Parallel prefix sum for image registration

- Communication cost $\ll$ computation cost
  $\Longrightarrow$ do not optimize for communication
- Reduce computation time
  $\Longrightarrow$ minimize span
- Variations in execution time are not known a priori
  $\Longrightarrow$ distribute data equally between MPI ranks

Distributed prefix sum

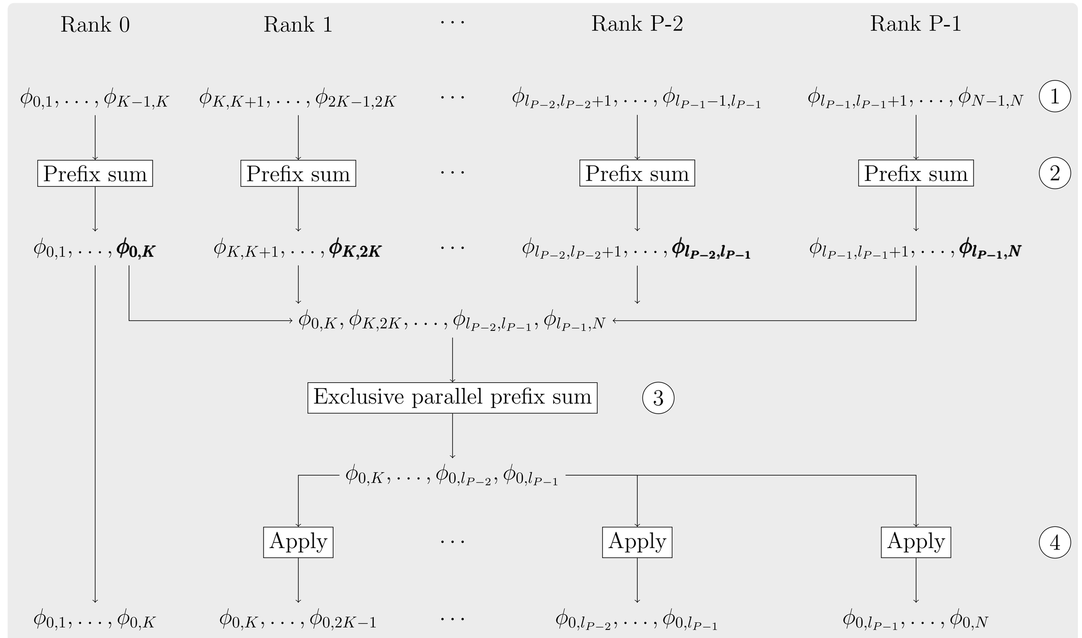

1. Rank J obtains a sequence of data $\phi_{I_J,I_J+1}, \phi_{I_J+1,I_J+2}, \ldots, \phi_{I_{J+1}-1,I_{J+1}}$ with $K = \frac{N}{P}$ elements, where $I_J = J \cdot K$ is a left index for rank J
2. Rank J performs a local prefix sum, reducing sequence to $\phi_{I_J,I_{J+1}}$

$$S(N,P) = \frac{N}{P} - 1$$

3. A global exclusive prefix sum on reduced values $\phi_{I_J,I_{J+1}}$ is performed, rank J obtains a result ϕ_{0,I_J}

$$S(N,P) = C_1 \log_2 P$$

4. The global result ϕ_{0,I_J} and local values ϕ_{I_J,I_J+i} are combined to form final results ϕ_{0,I_J+i}

$$S(N,P) = \frac{N}{P}$$

The span of a prefix sum for N images on P ranks

$$S(N,P) = \frac{N}{P} + C_1 \log_2 P + \frac{N}{P} + C_2$$

The attainable speedup is bounded from above by

$$SP(N,P) = \frac{S(N,1)}{S(N,P)} = \frac{N-1}{\frac{2N}{P} + C_1 \log_2 P + C_2}$$

Strong scaling

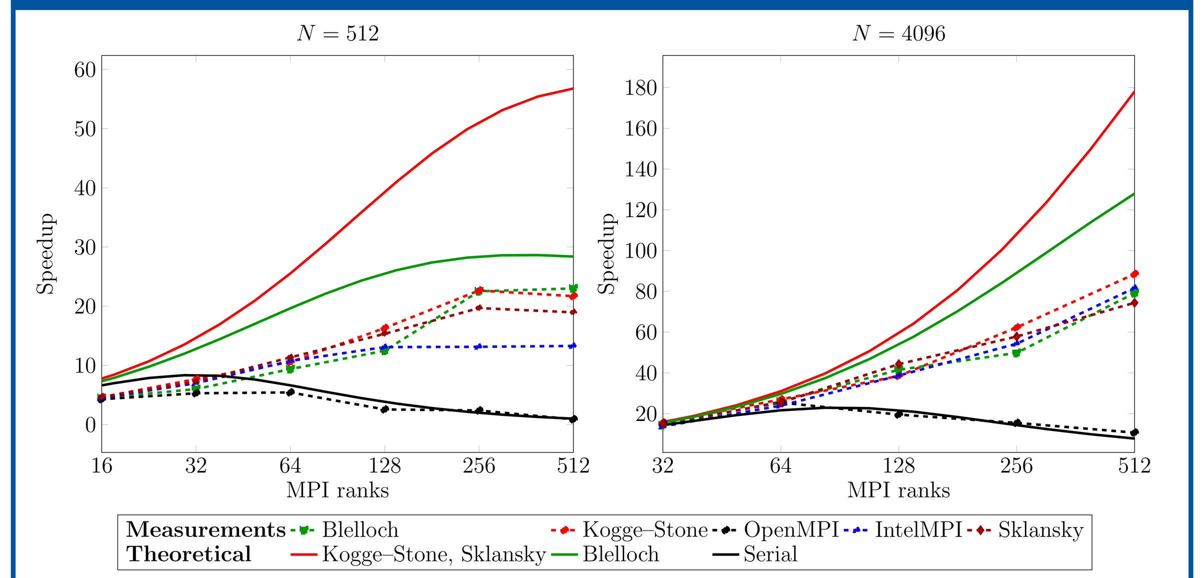

Theoretical and measured speedups are represented with solid and dashed lines, respectively.

Weak scaling

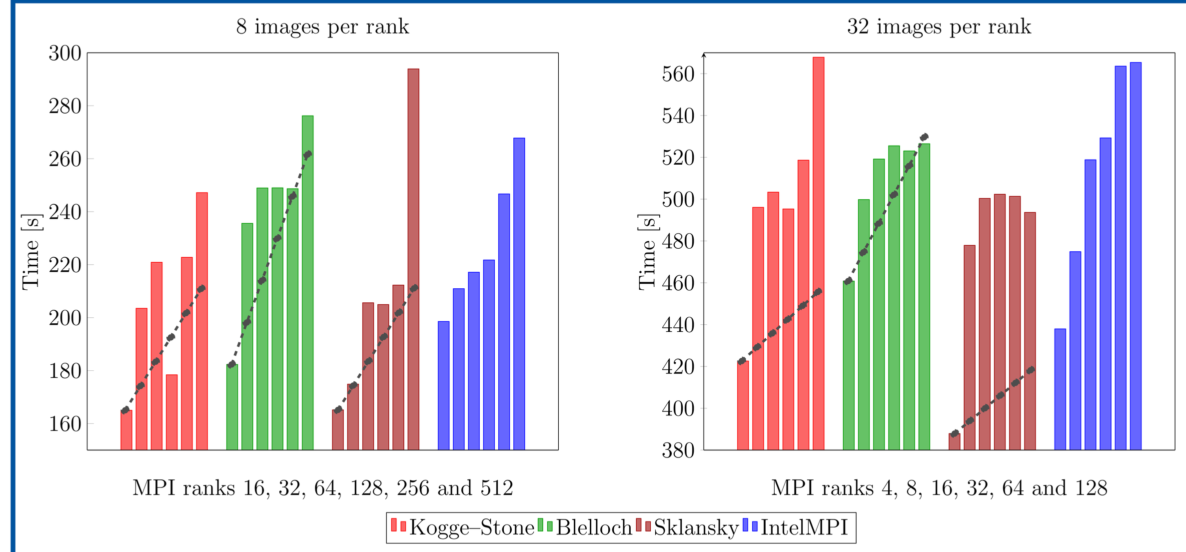

Linear weak scaling is not attainable, since

$$S(k \cdot N, k \cdot P) = S(N,P) + \log_2 k \cdot C_1$$

Multithreading

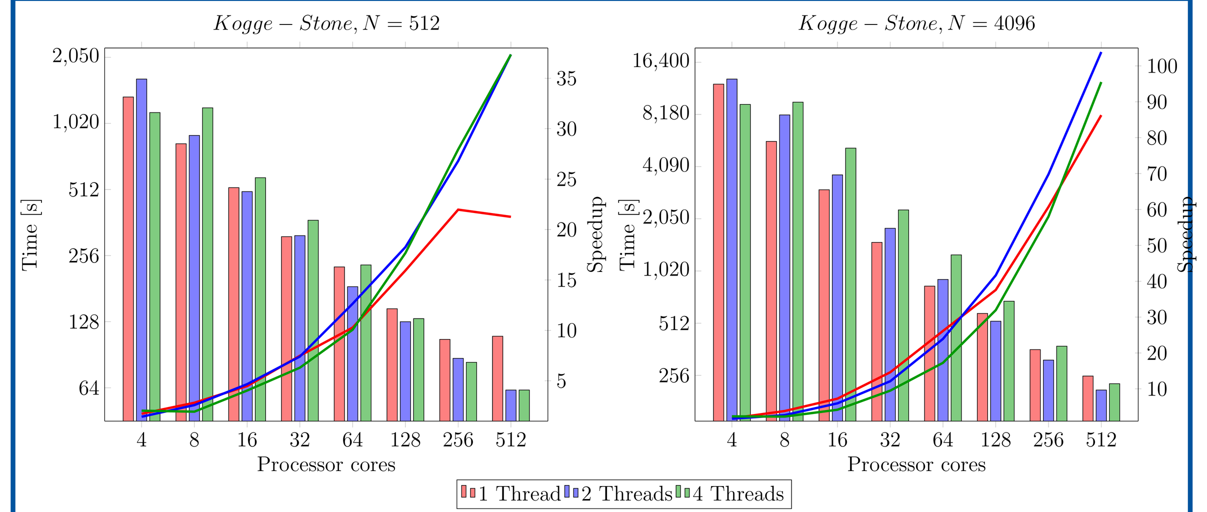

Allocate $\frac{P}{T}$ ranks with T threads per each MPI process.

Summary

- The benchmarking has been performed on Intel Xeon E5-2680 v2 IvyBridge processors with GCC 5.3
- Parallelization is limited by the nature of prefix sum
- Measured performance does not meet theoretical predictions because of the unbalanced workload
- Custom scan implementations outperform `MPI_Scan` and `MPI_Scan` from OpenMPI 1.10 and IntelMPI 2017.1
- The performance is improved by shifting resources to operator parallelization
- The image registration procedure is a novel example of a parallel prefix sum with non–trivial sum operator

References

- Benjamin Berkels et al. (2014). “Optimized imaging using non-rigid registration”. In: *Ultramicroscopy* 138, pp. 46–56
- G. E. Blelloch (Nov. 1989). “Scans As Primitive Parallel Operations”. In: *IEEE Trans. Comput.* 38.11, pp. 1526–1538

Marcin Copik
E-Mail: mcopik@gmail.com

High Performance and Automatic Computing
Schinkelstr. 2, 52062 Aachen, GERMANY
www.hpac.rwth-aachen.de
Telefon: +49 241 80 99 131, Fax: +49 241 80 628498

HPAC High Performance and Automatic Computing
RWTH AACHEN UNIVERSITY