Parallel Fringe Search

Lukas Mosimann & Christian Zeman

ETH Zürich

Design of Parallel and High-Performance Computing

December 16, 2013

Overview

Fringe Search

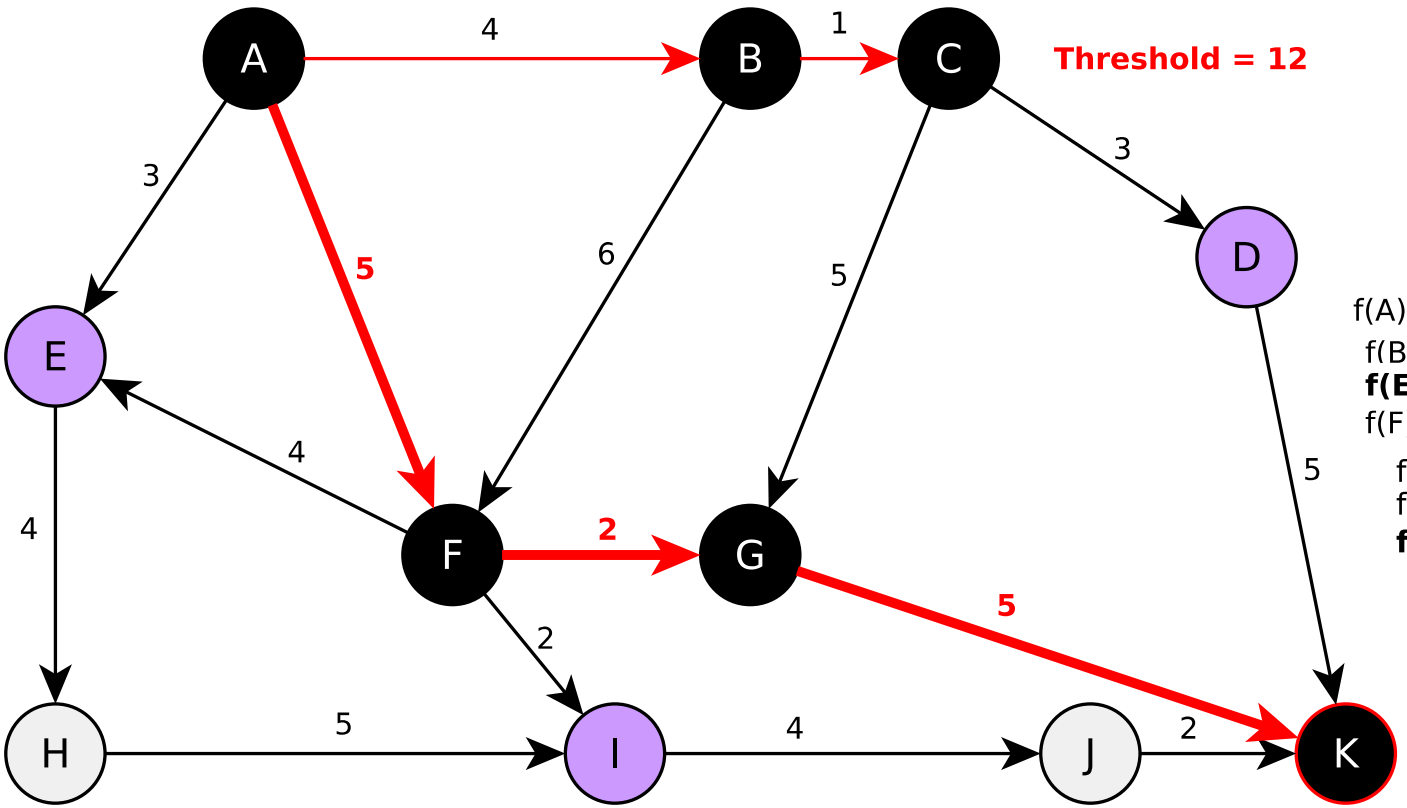

What we have done

- Serial implementation of fringe search (much faster than Boost A*)
- Parallel implementation with Open MP
  - 2 different locking concepts
  - Locks implemented using inline assembly (faster than Open MP locks)
- Benchmarking
  - Strong scaling
  - Weak scaling
  - Path quality
  - Threshold handling

Locking concept: Deadlock prevention

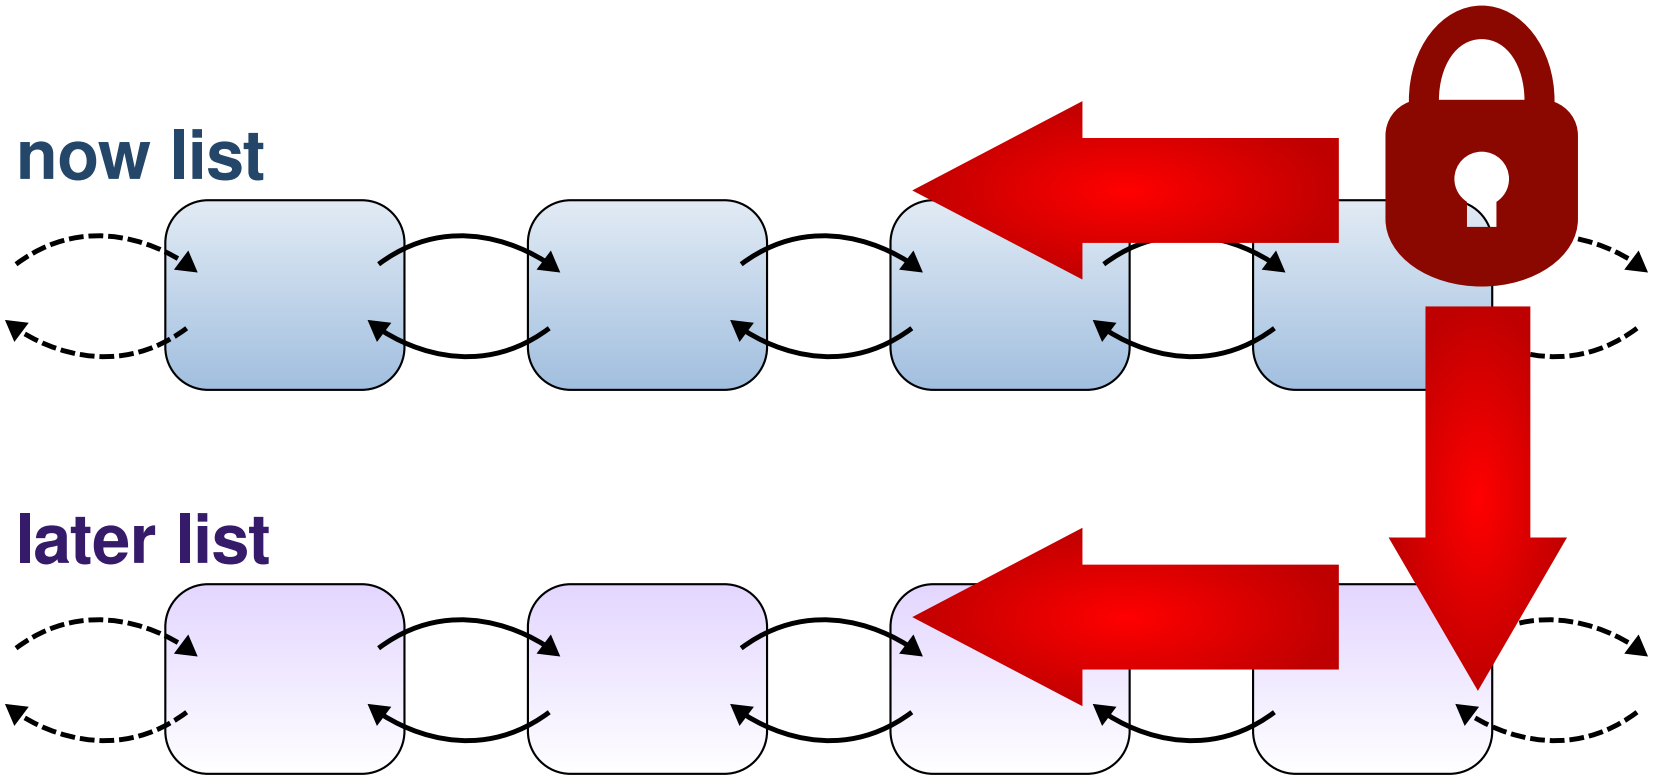

Locking concept: Insert node

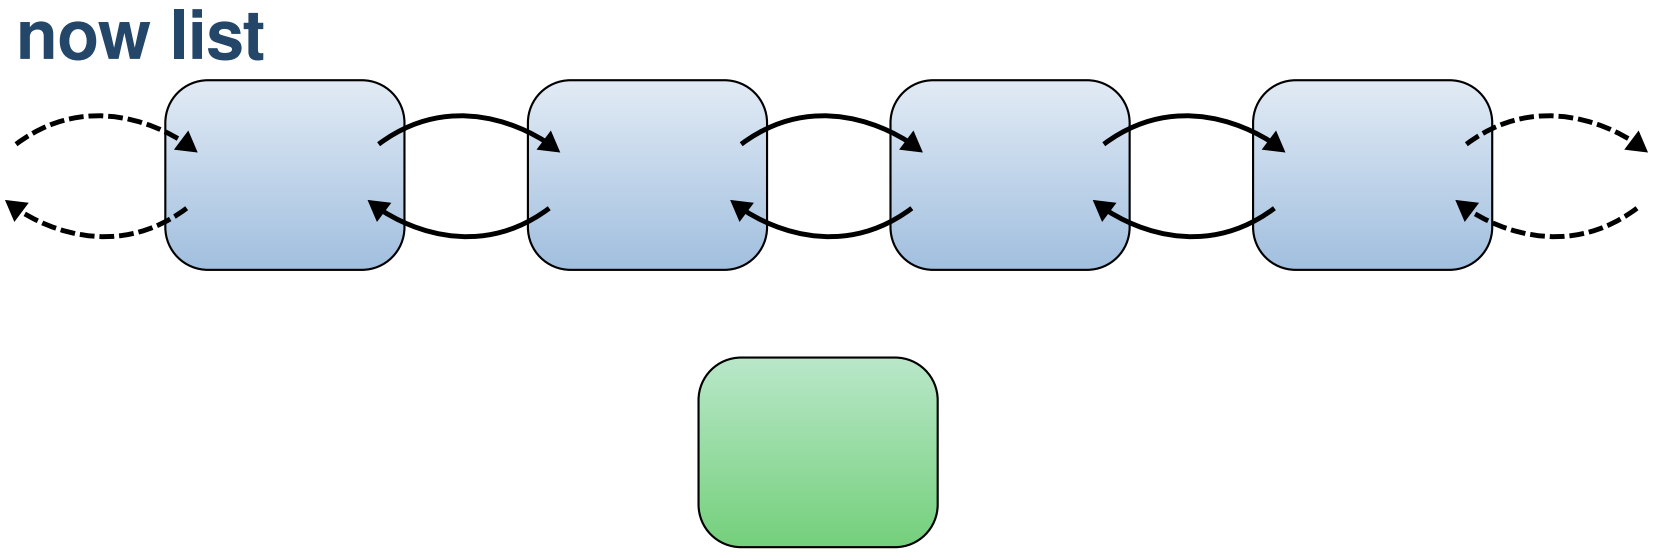

Locking concept: Insert node

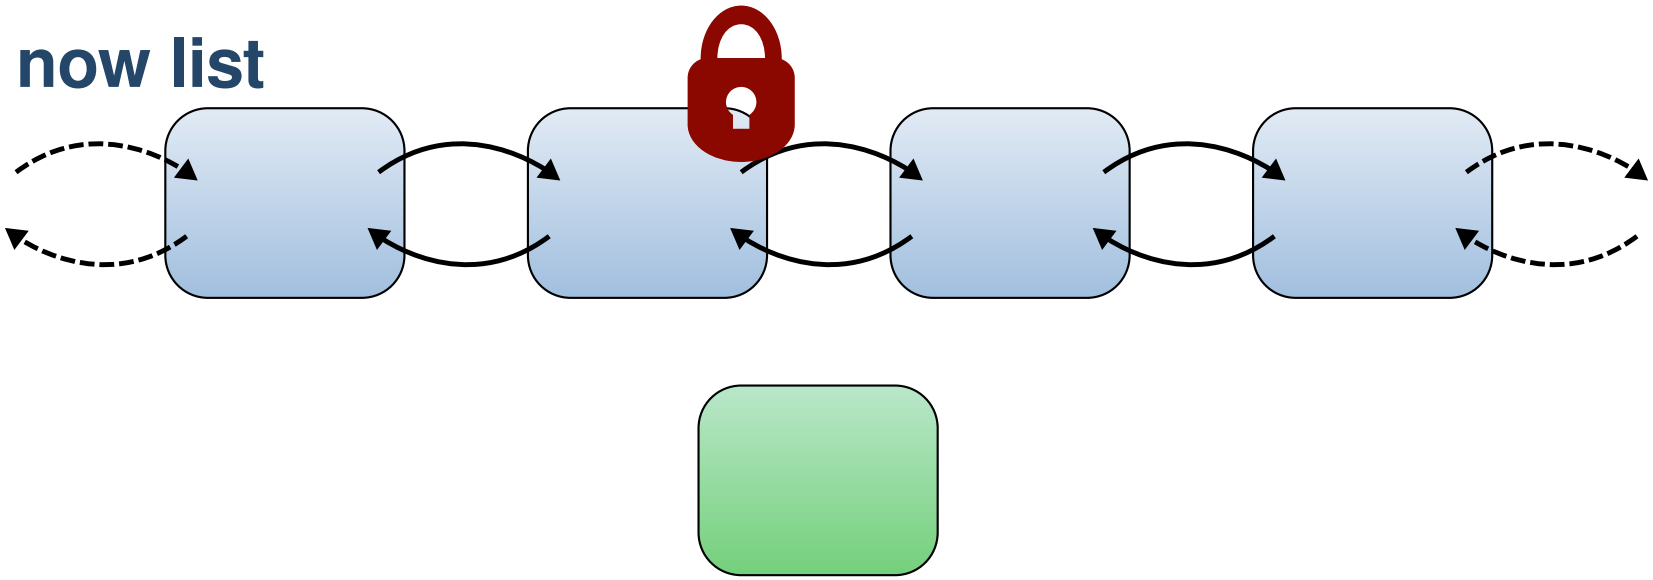

Locking concept: Insert node

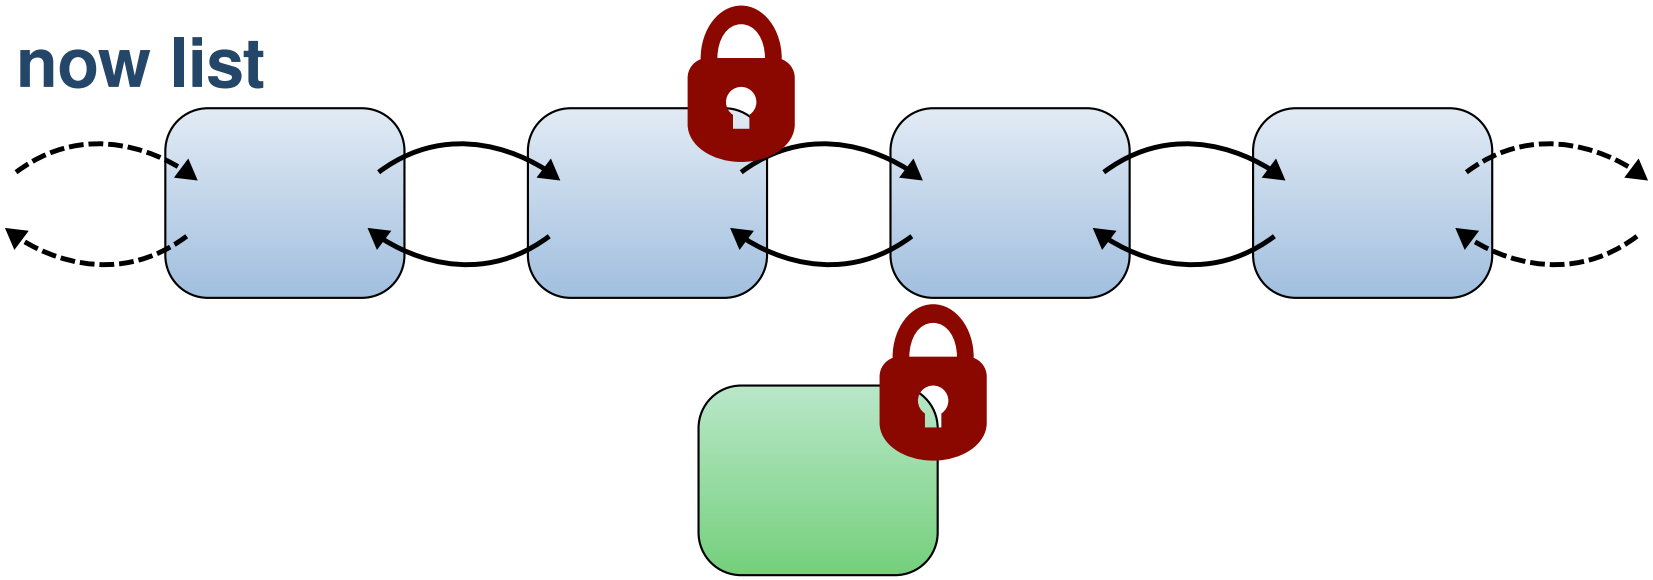

Locking concept: Insert node

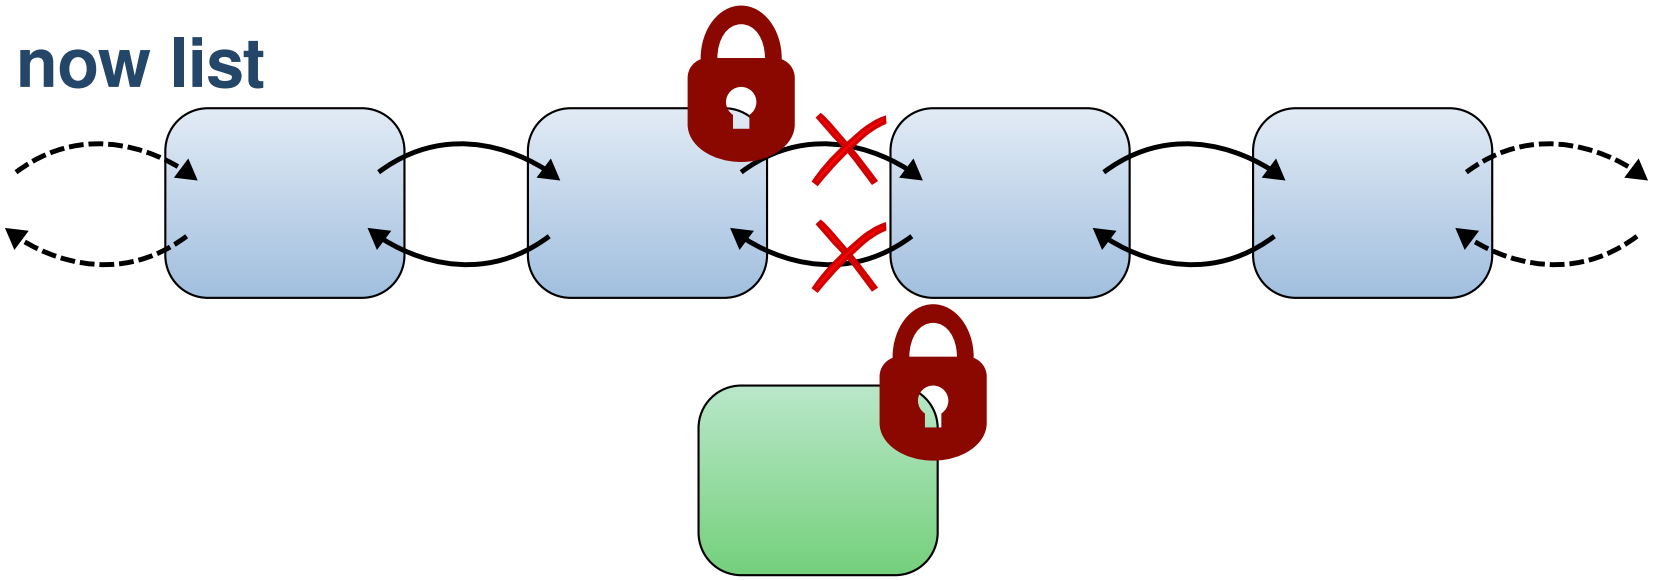

Locking concept: Insert node

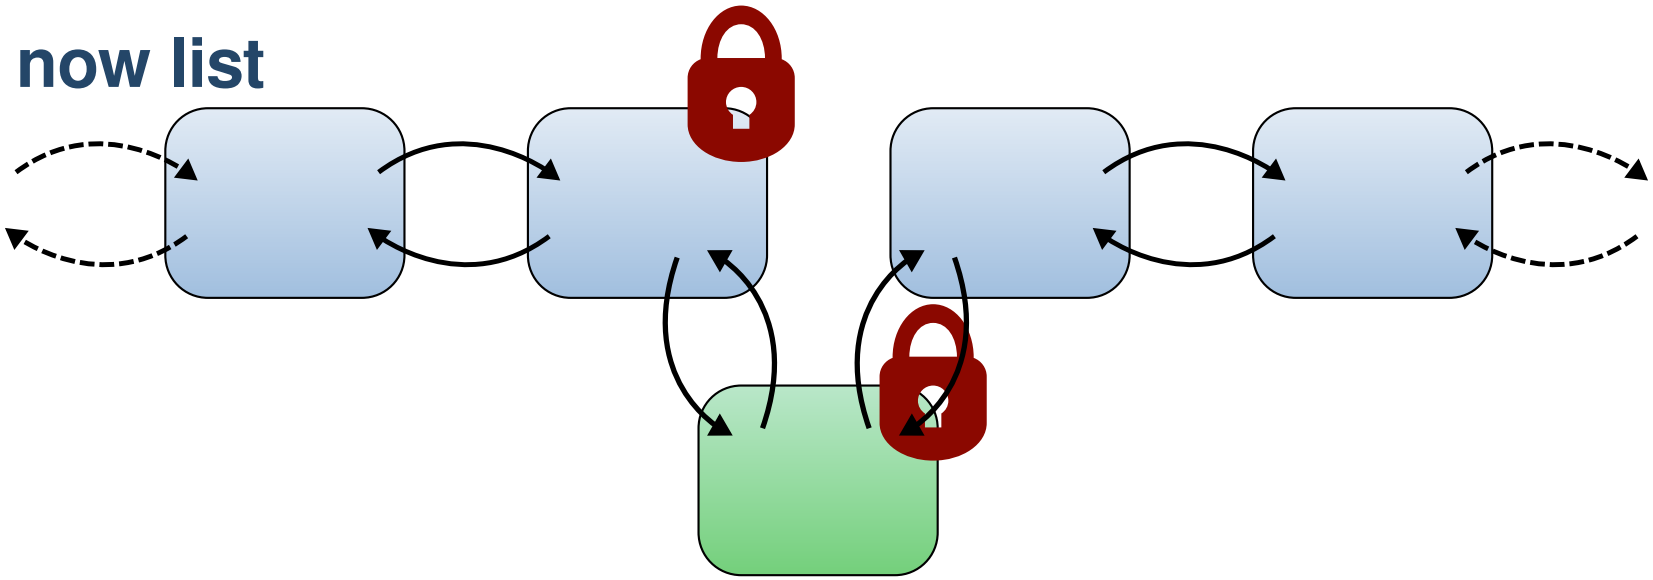

Locking concept: Insert node

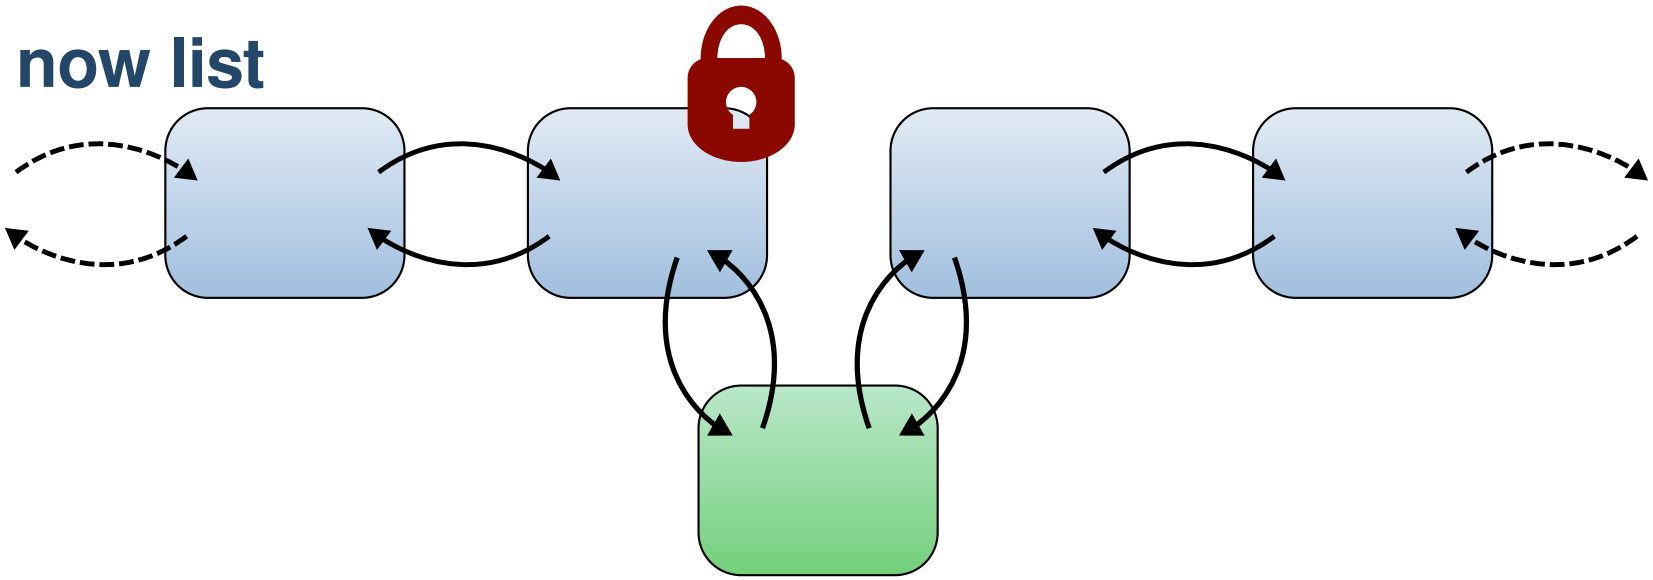

Locking concept: Insert node

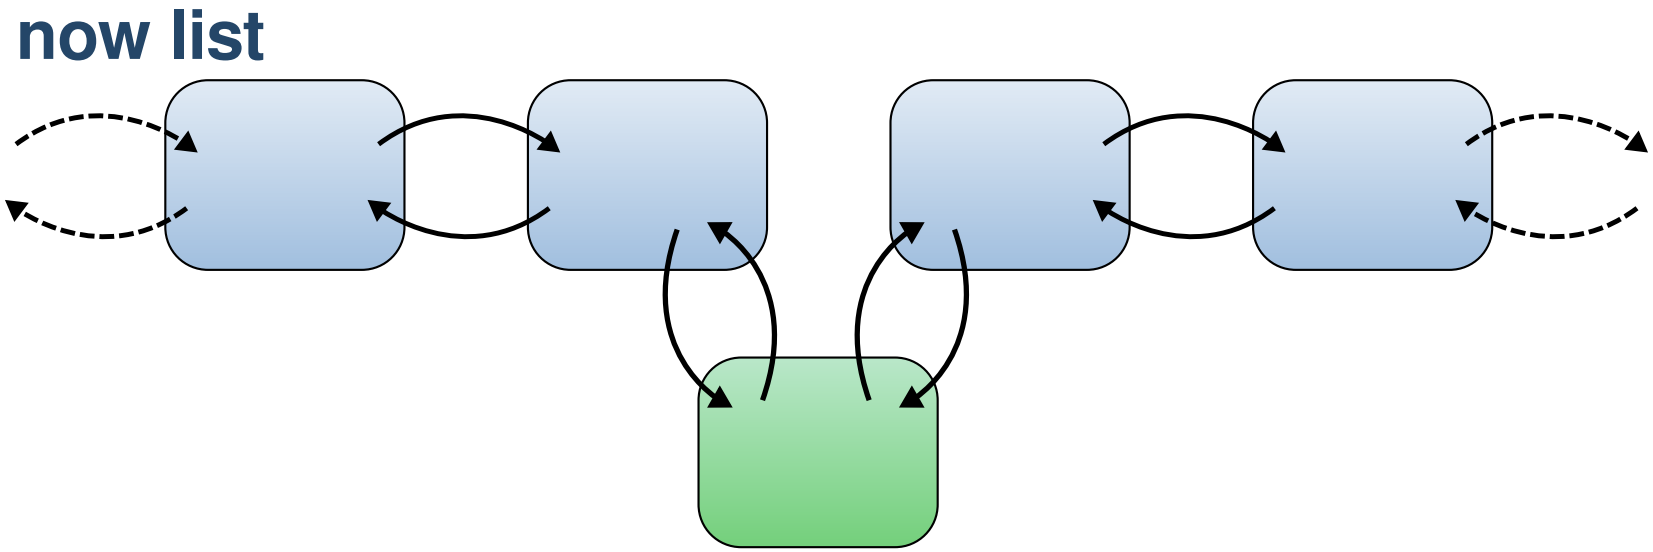

Locking concept: Remove node

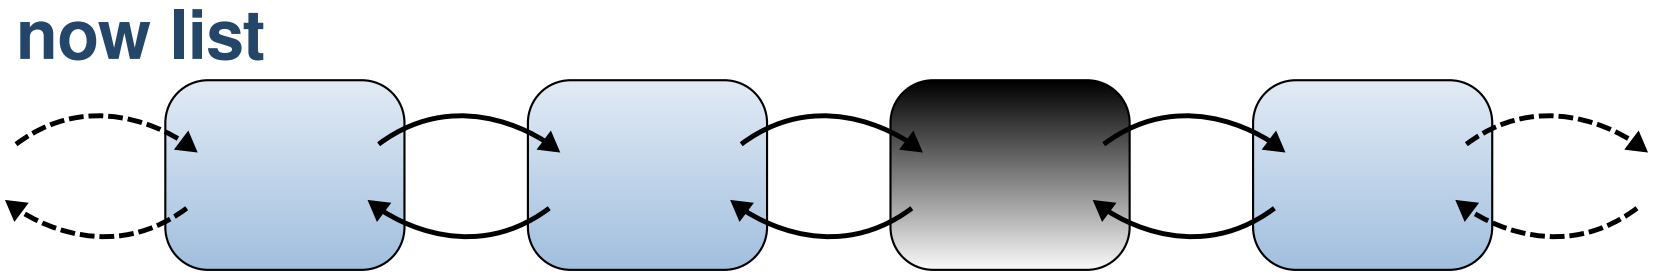

Locking concept: Remove node

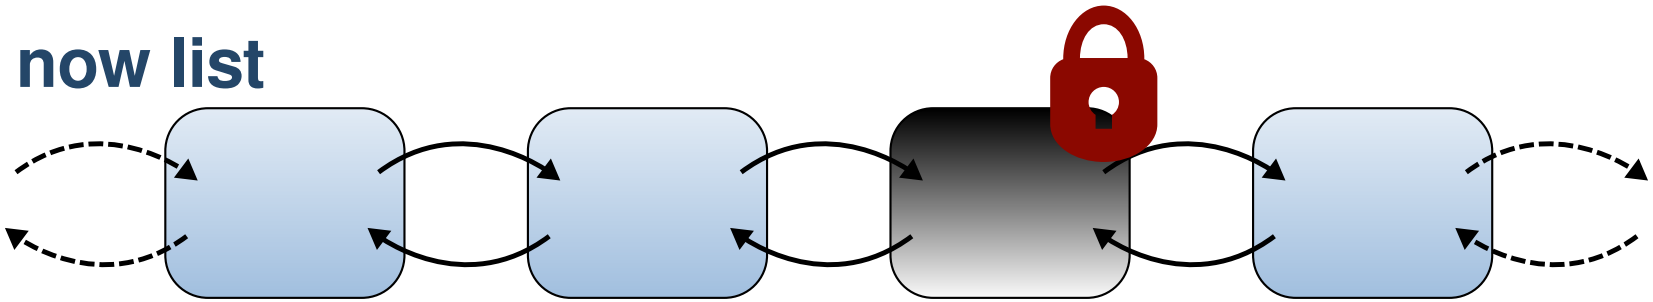

Locking concept: Remove node

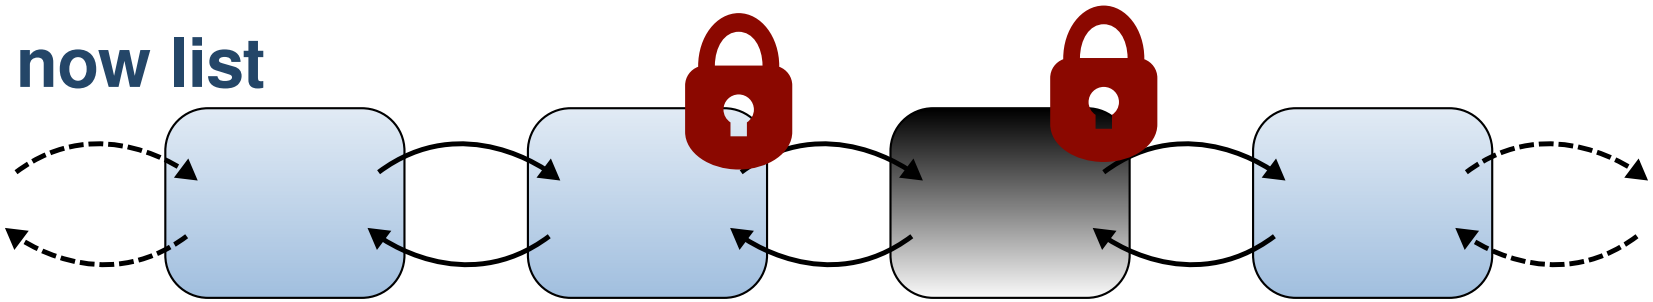

Locking concept: Remove node

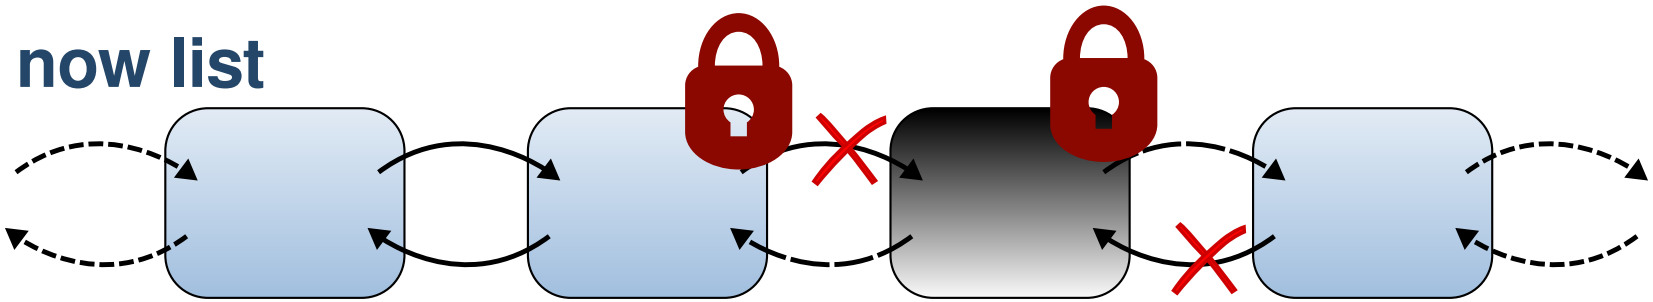

Locking concept: Remove node

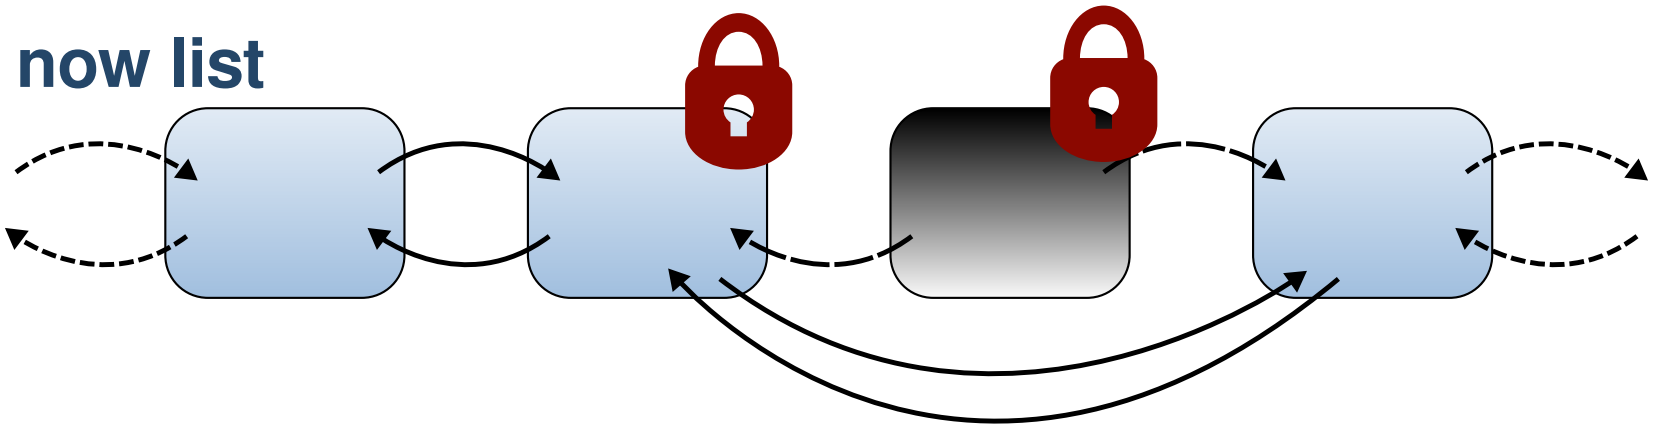

Locking concept: Remove node

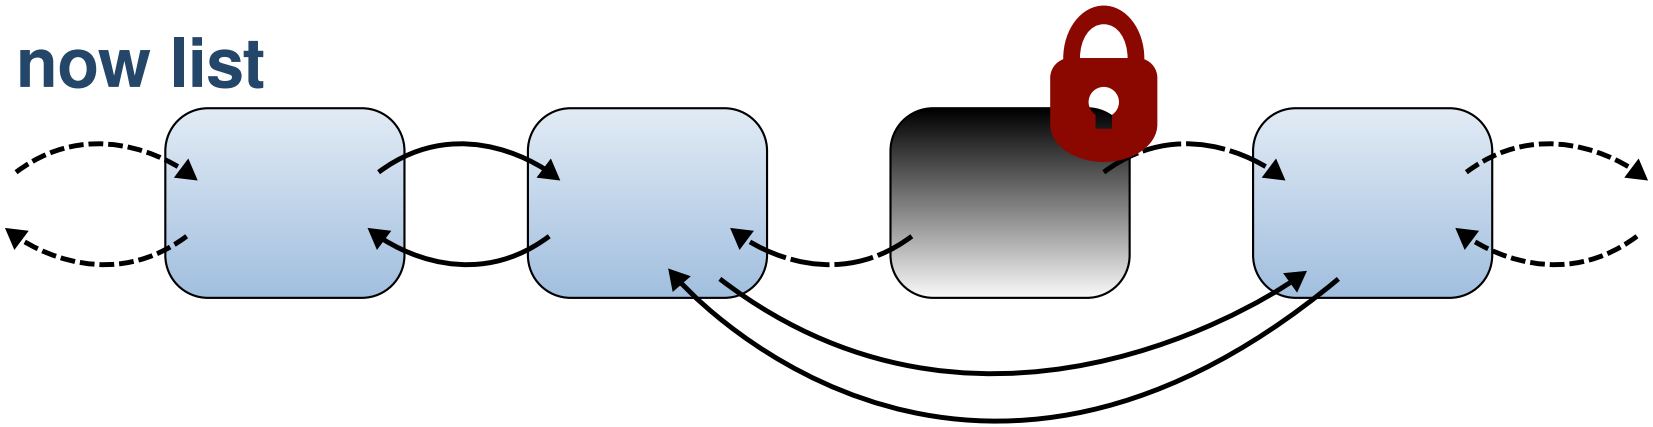

Locking concept: Remove node

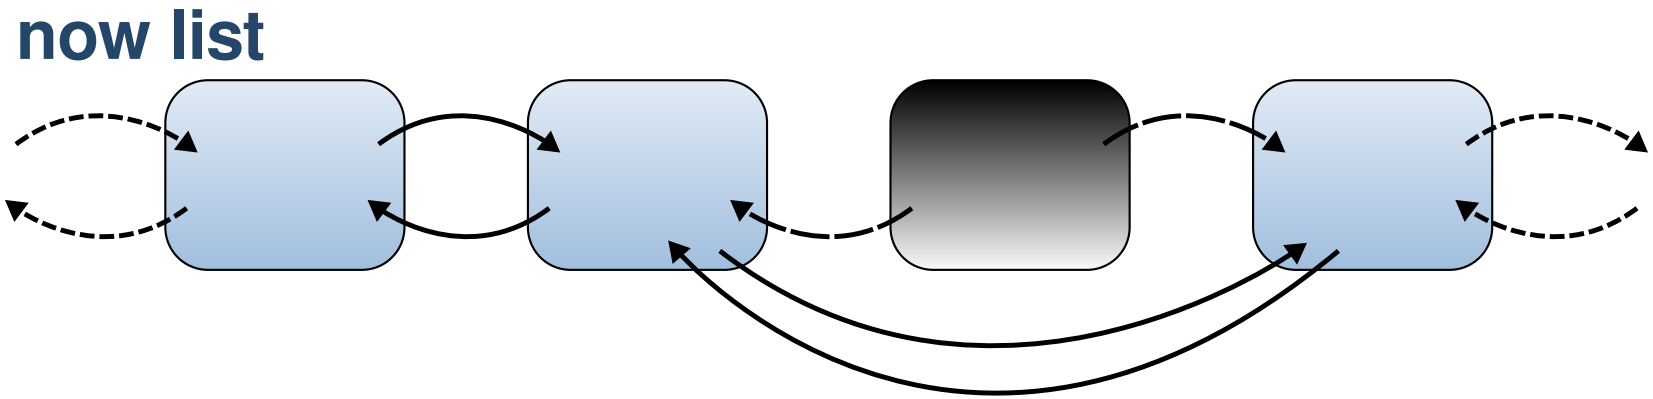

Locking concept: 2 concepts for removing nodes

Both concepts use optimistic locking for acquiring the locks.

Normal:

- Lock node and predecessor as shown before and remove it right away

Lazy locking:

- Don't lock anything and just mark the node as removed
- Other threads will clean up and remove it later

Benchmarking

Setup:

Each of the following boxplots was generated from data from **50 runs** on **1 node** of kanifushi.inf.ethz.ch.

Specifications of kanifushi.inf.ethz.ch:

- NUMA model with 32 CPUs on 4 nodes
- 8 CPUs per node
- Intel(R) Xeon(R) CPU E7- 4830 @ 2.13GHz
- per CPU: 32KB L1 cache, 256KB L2 cache
- per node: 24MB L3 cache, 16GB memory

The code is written in C++ / Open MP and it has been compiled with g++ v. 4.6.1 using O1 optimization.

Benchmarking

Graphs

- based on a regular grid with distance 1 and 8 edges per node
- each node is moved randomly (normal distribution, $\sigma = 0.3$)
- different obstacles (circle, crosses, etc.)

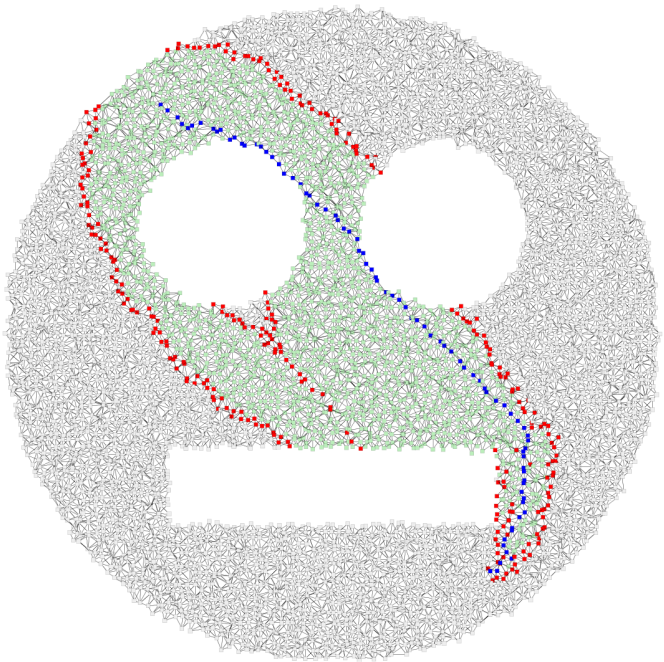

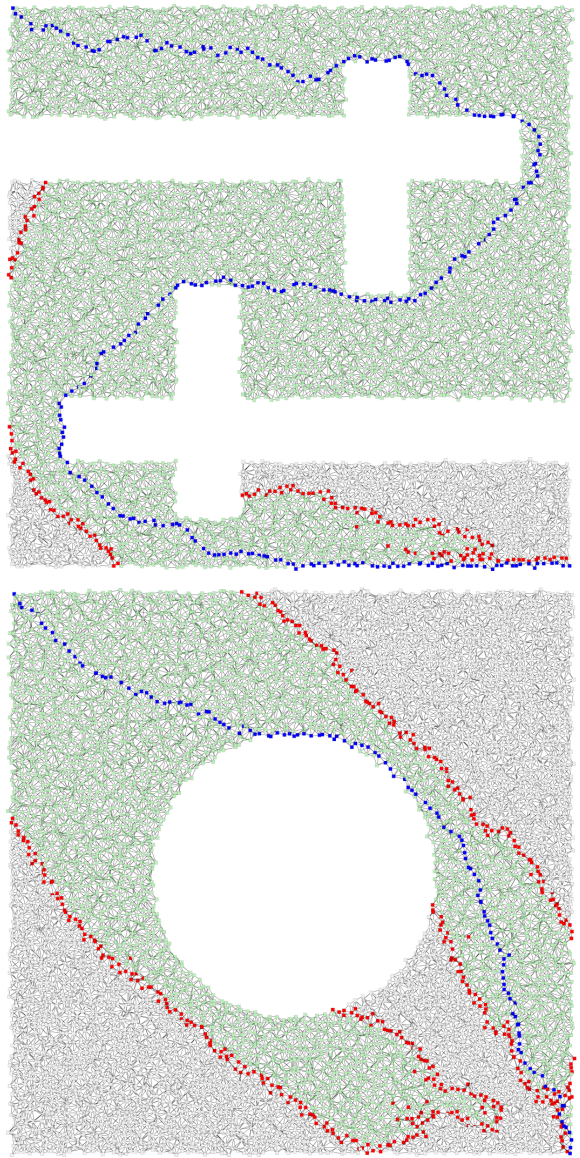

Benchmarking

Graphs

- based on a regular grid with distance 1 and 8 edges per node
- each node is moved randomly (normal distribution, $\sigma = 0.3$)
- different obstacles (circle, crosses, etc.)

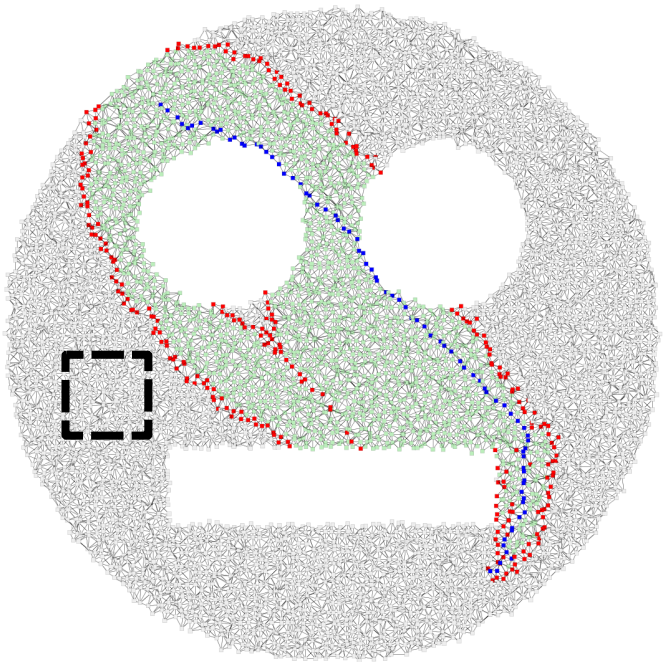

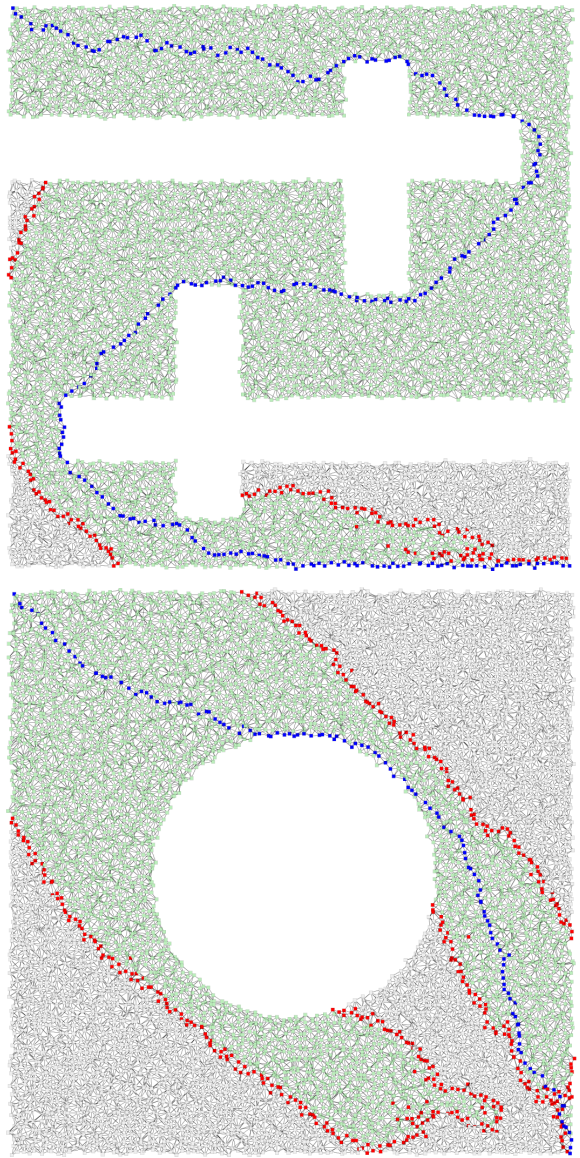

Benchmarking

Graphs

- based on a regular grid with distance 1 and 8 edges per node
- each node is moved randomly (normal distribution, $\sigma = 0.3$)
- different obstacles (circle, crosses, etc.)

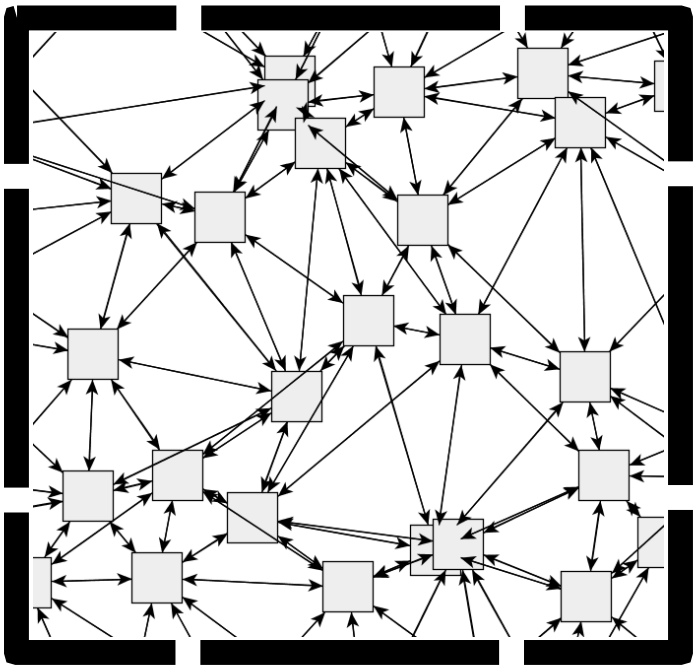

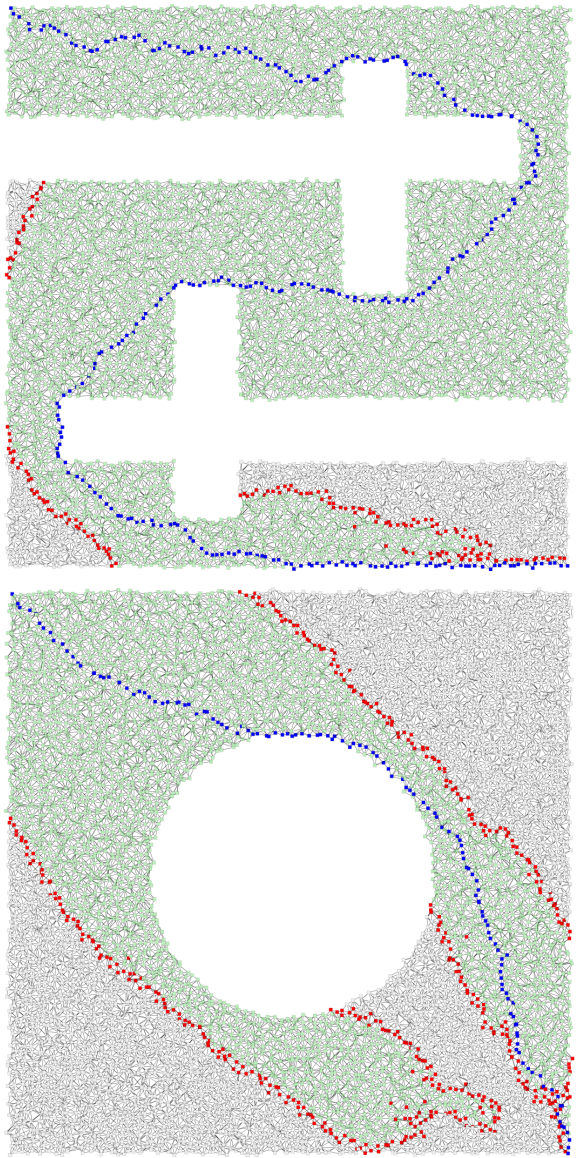

Benchmarking: Locks

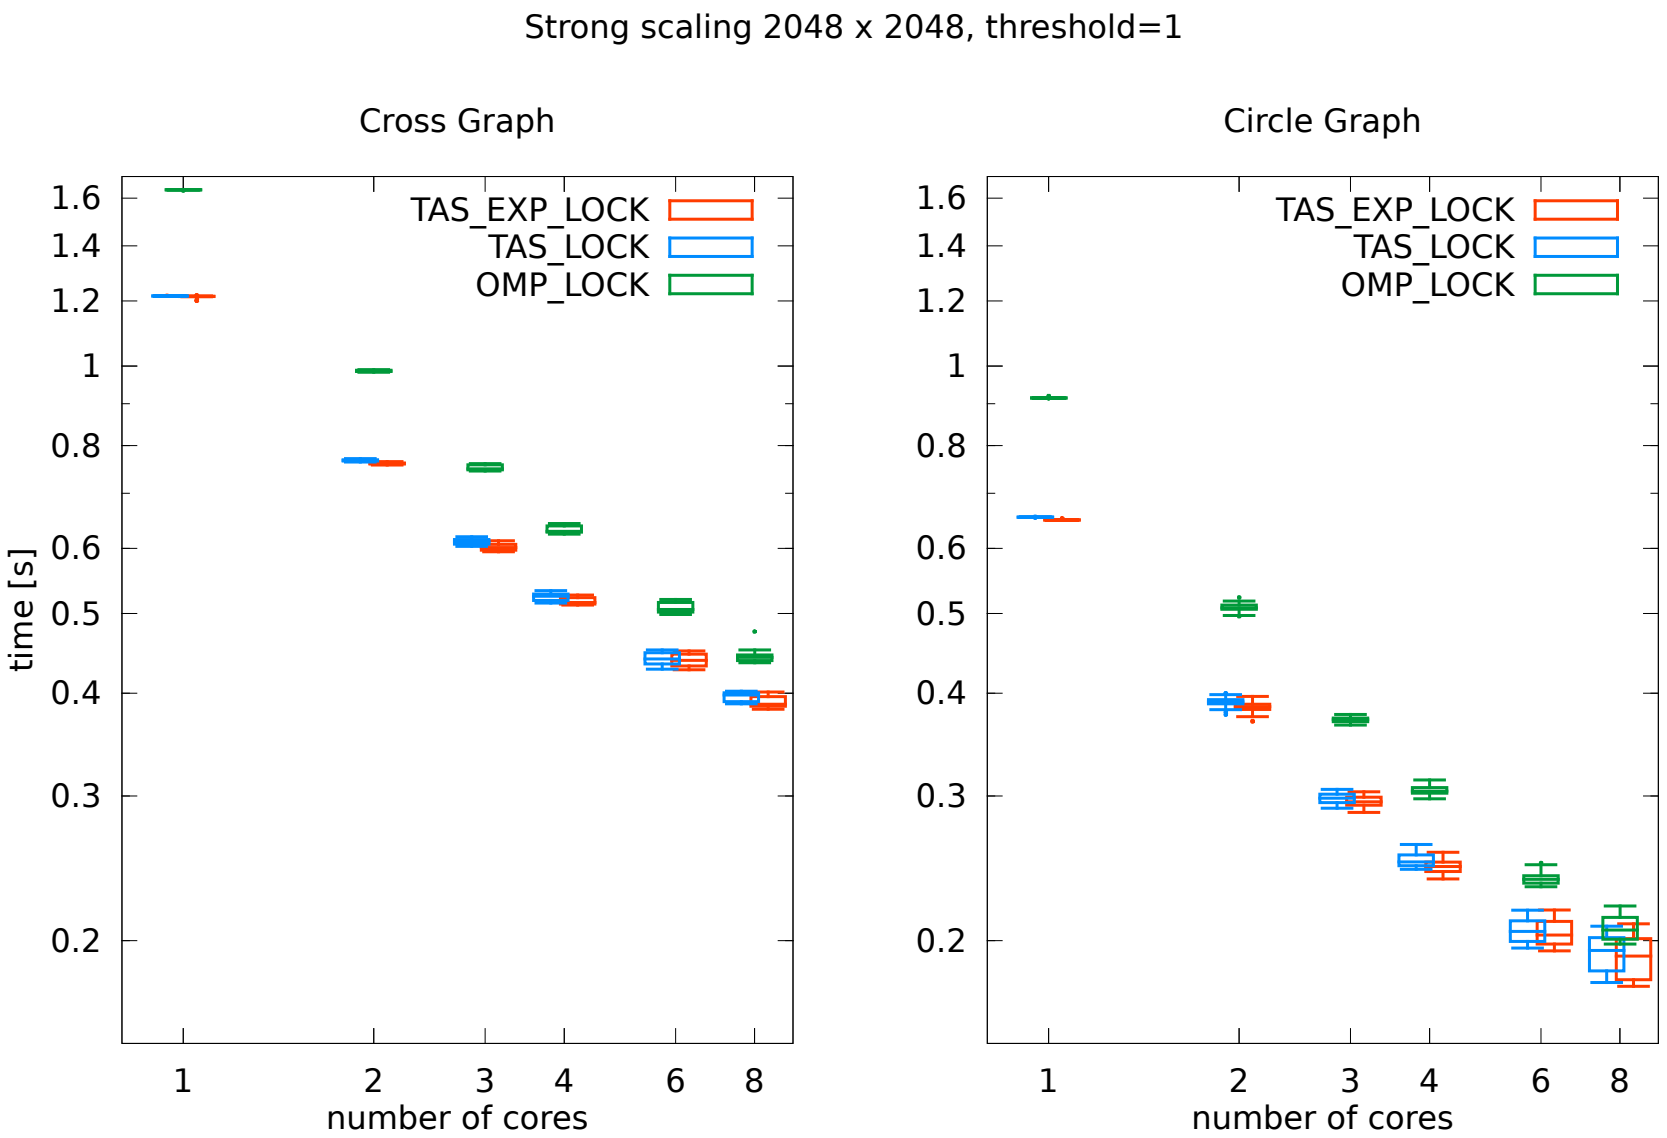

Benchmarking: Strong scaling

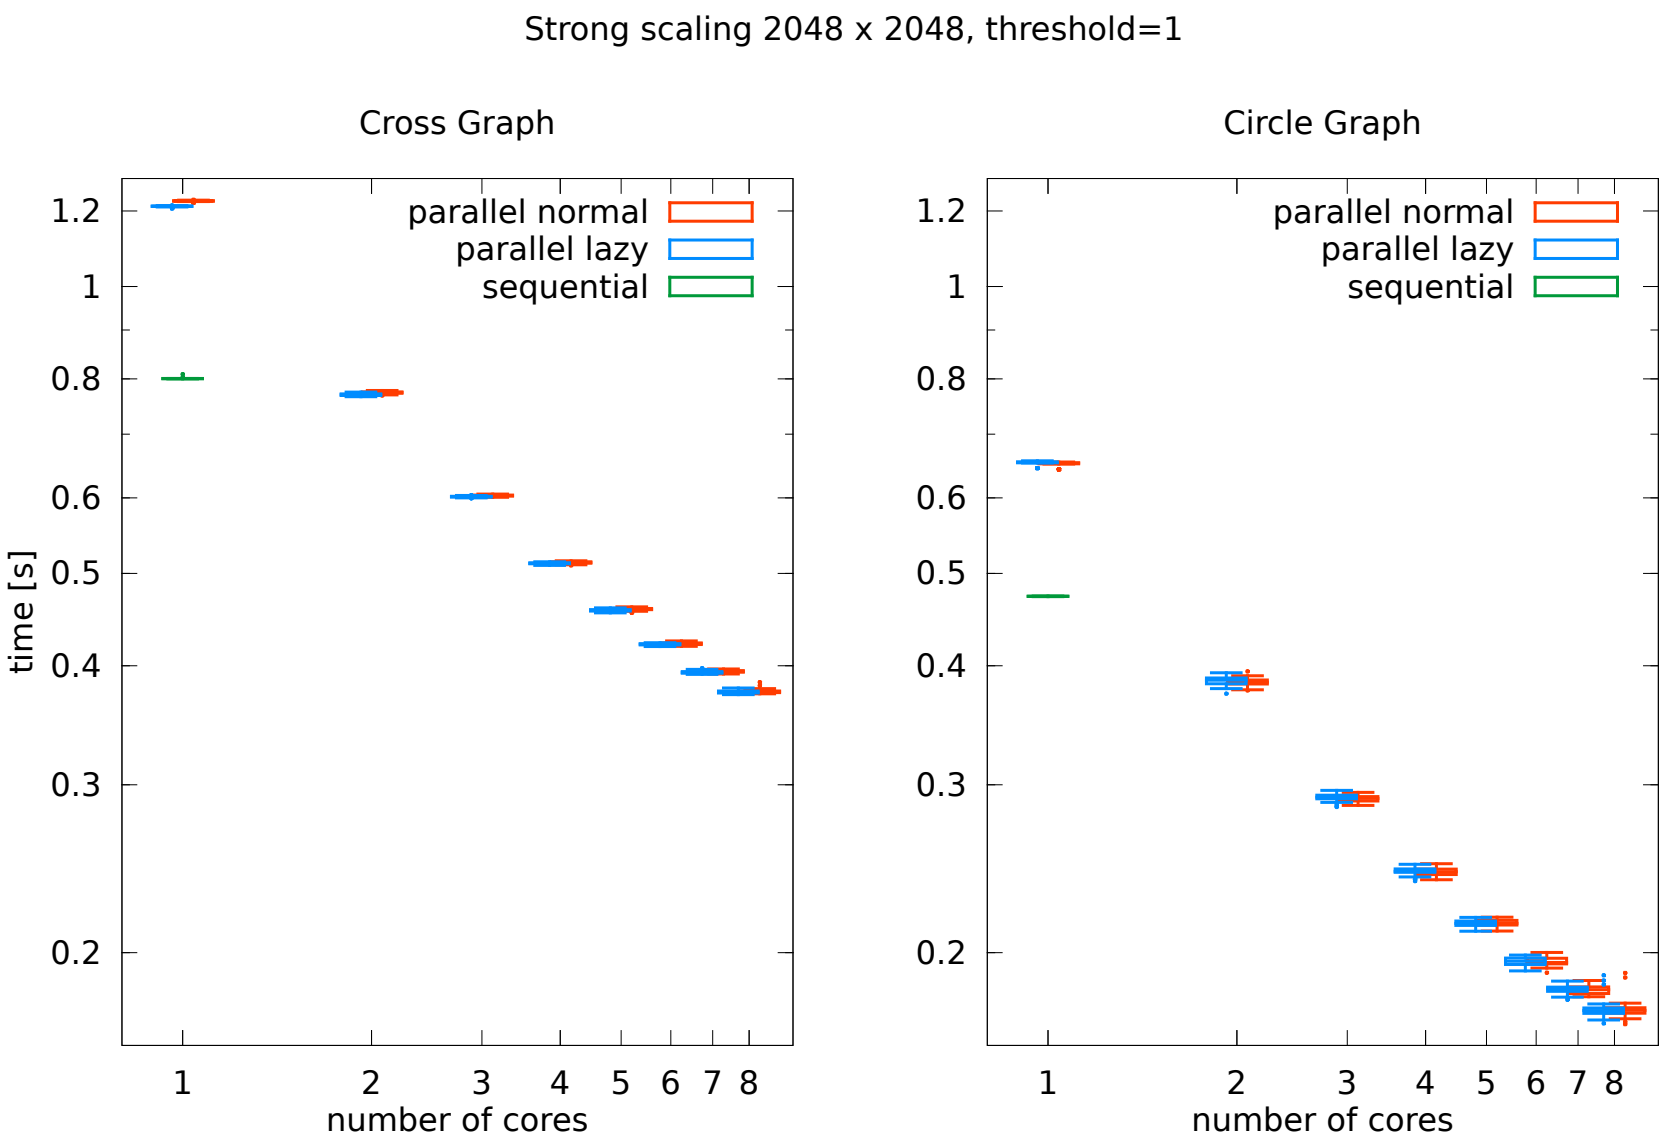

Benchmarking: Sequential scaling

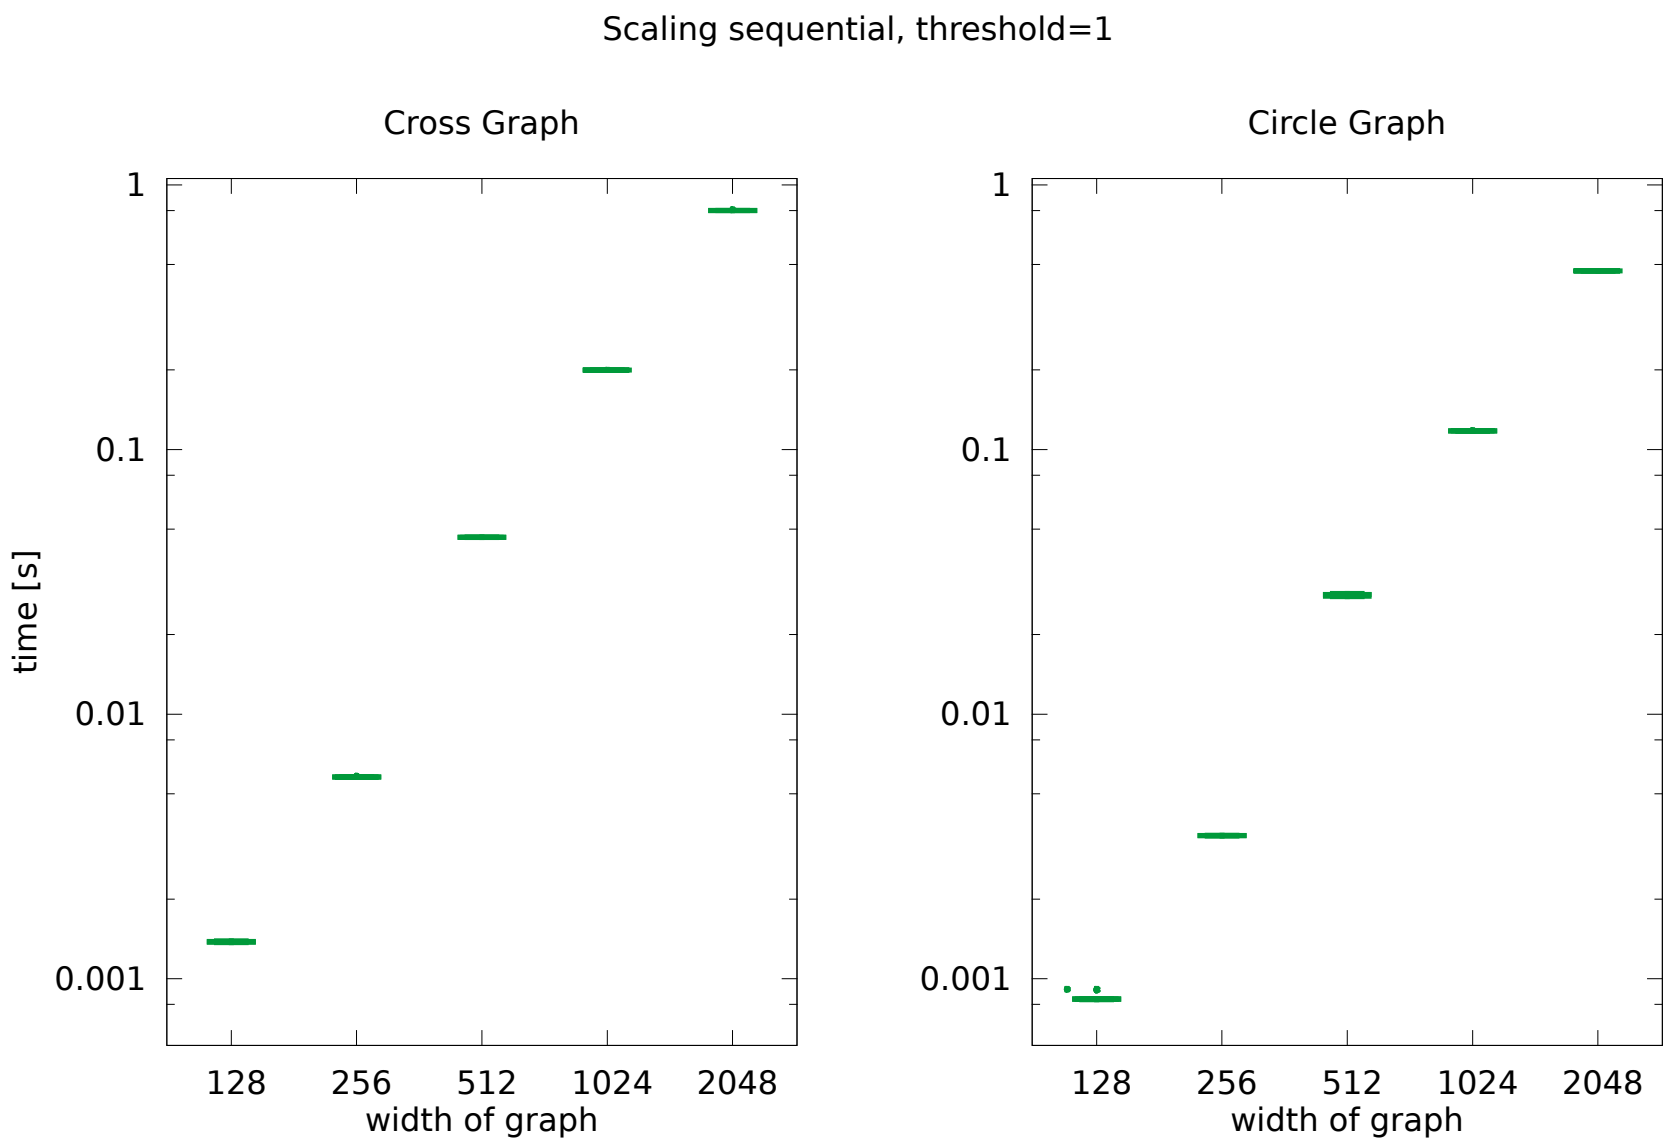

Benchmarking: Weak scaling

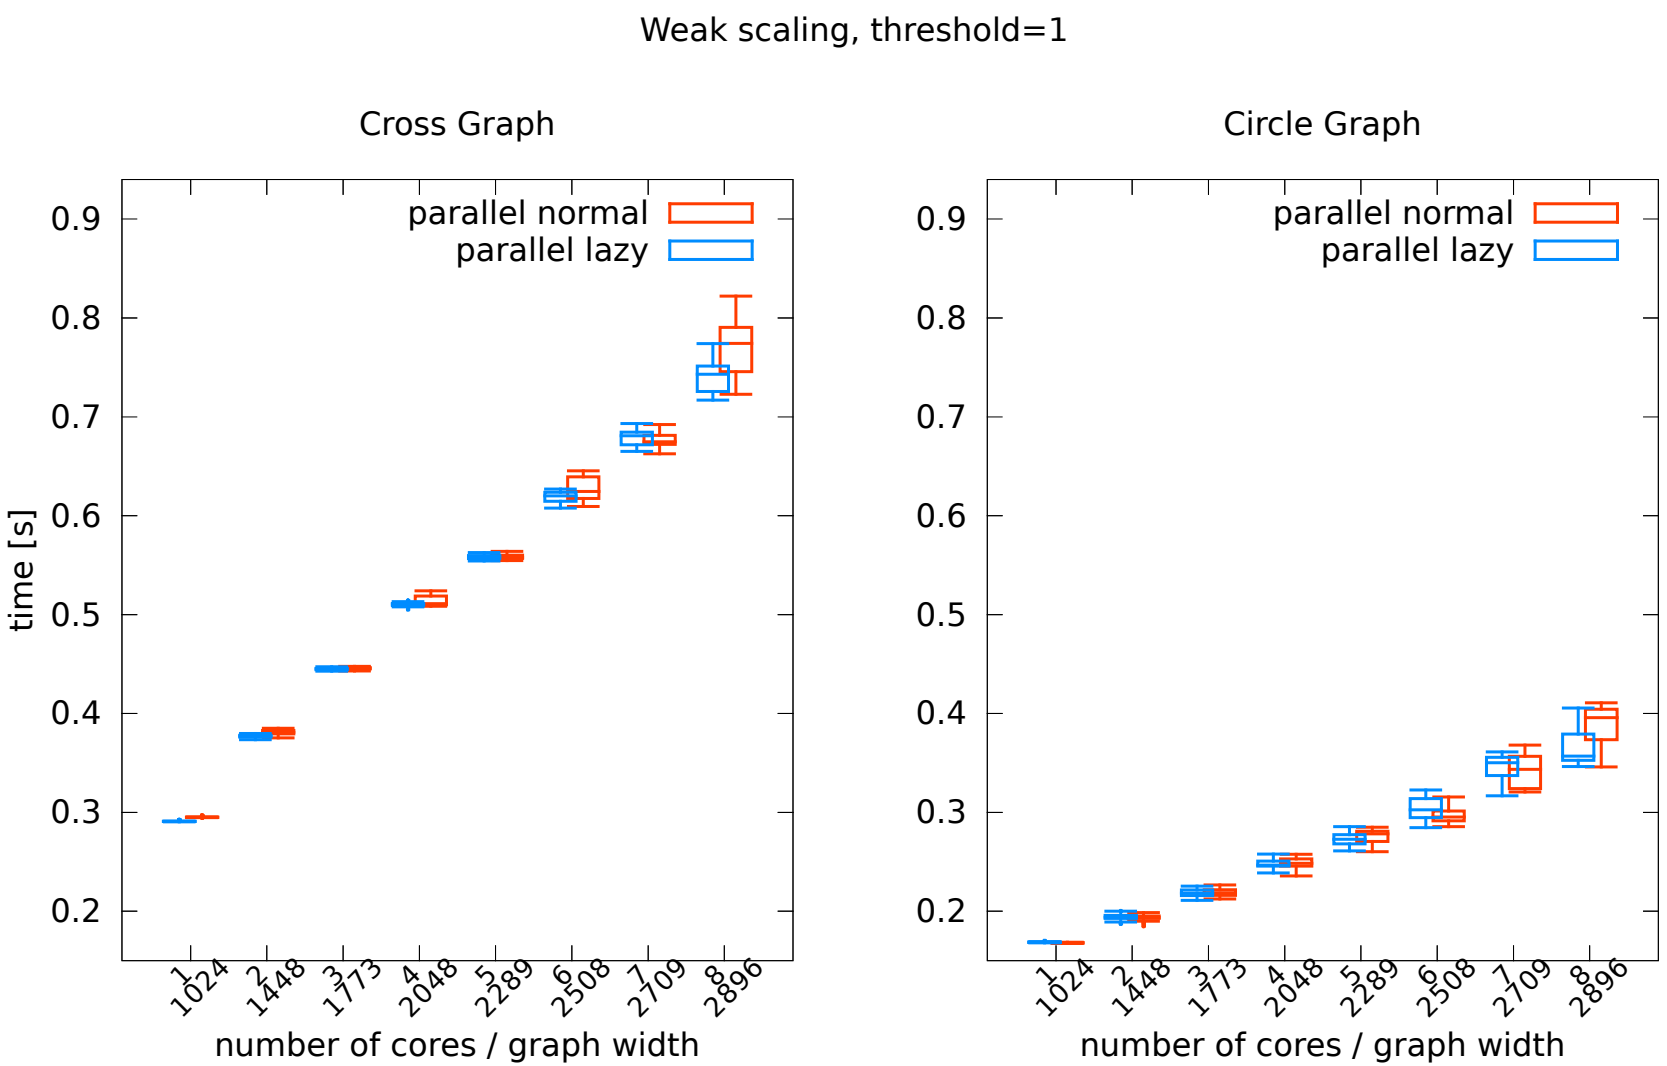

Benchmarking: Path length $\leftrightarrow$ # cores

compared to A* from Boost Graph Library

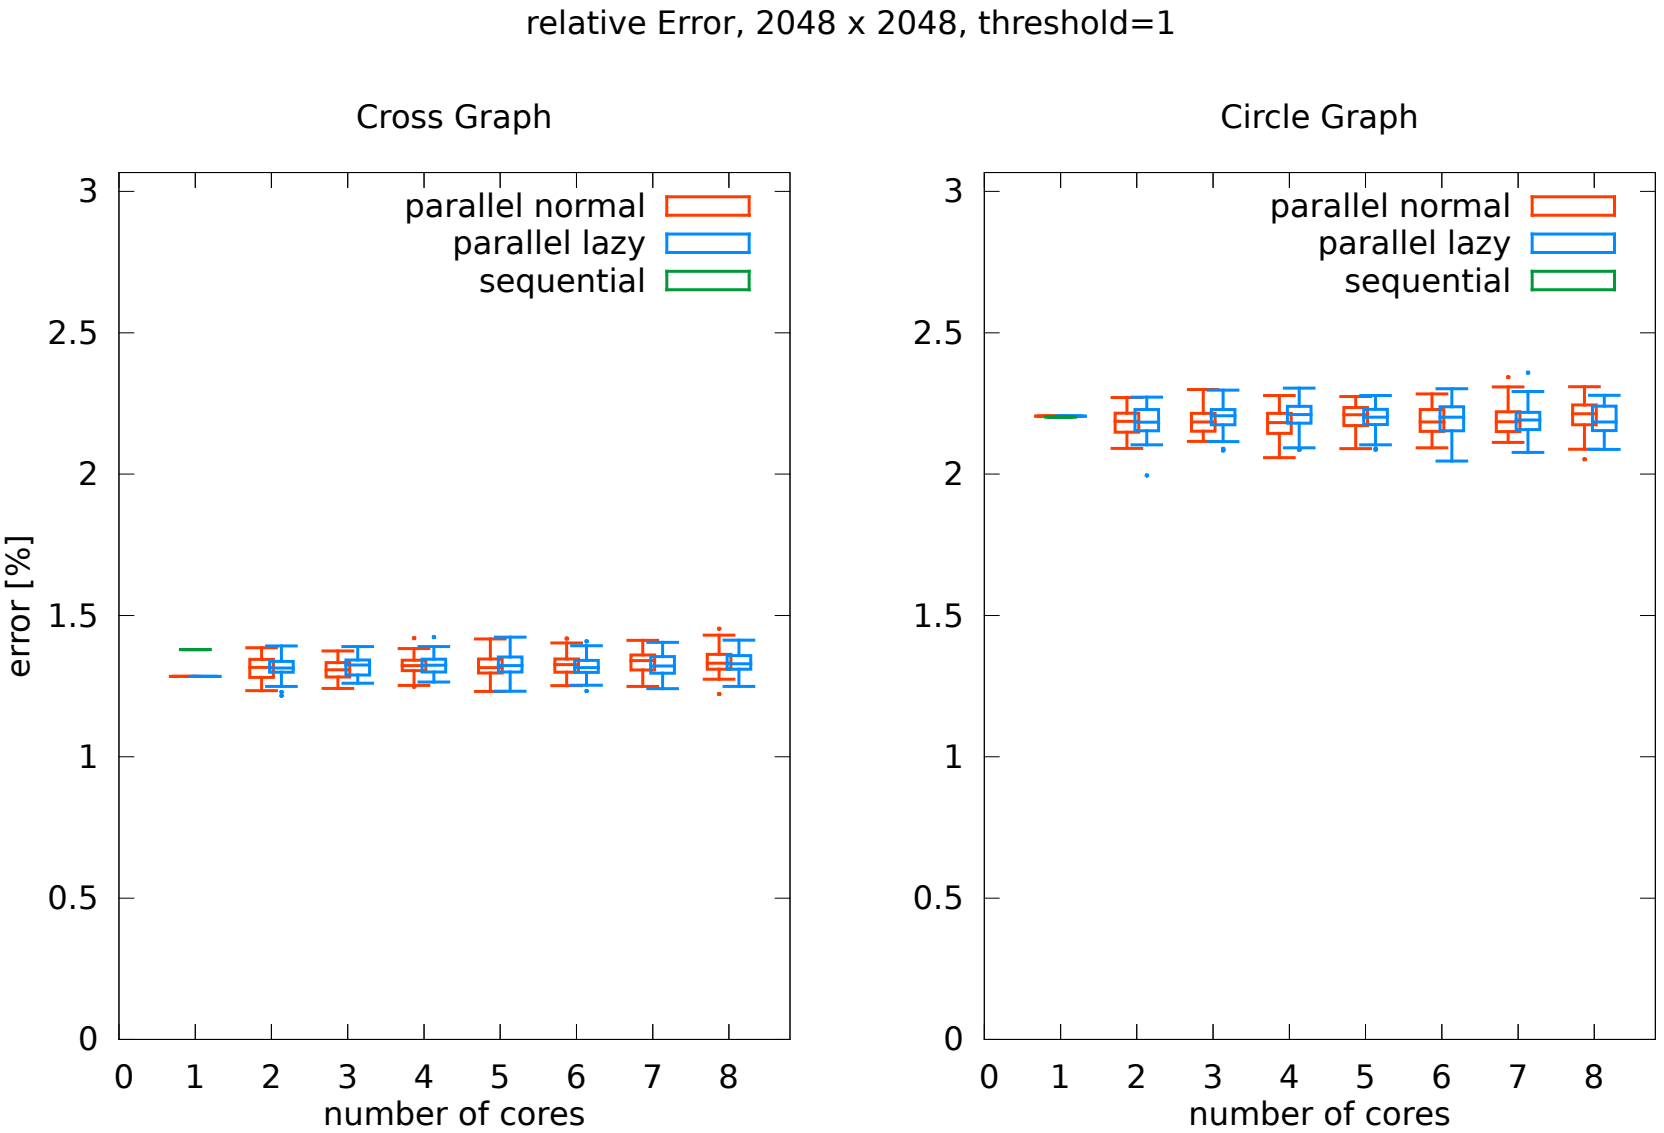

Benchmarking: Path length $\leftrightarrow$ Threshold update

Threshold update

- threshold += **0.1**
- threshold += **1**
- threshold += **10**

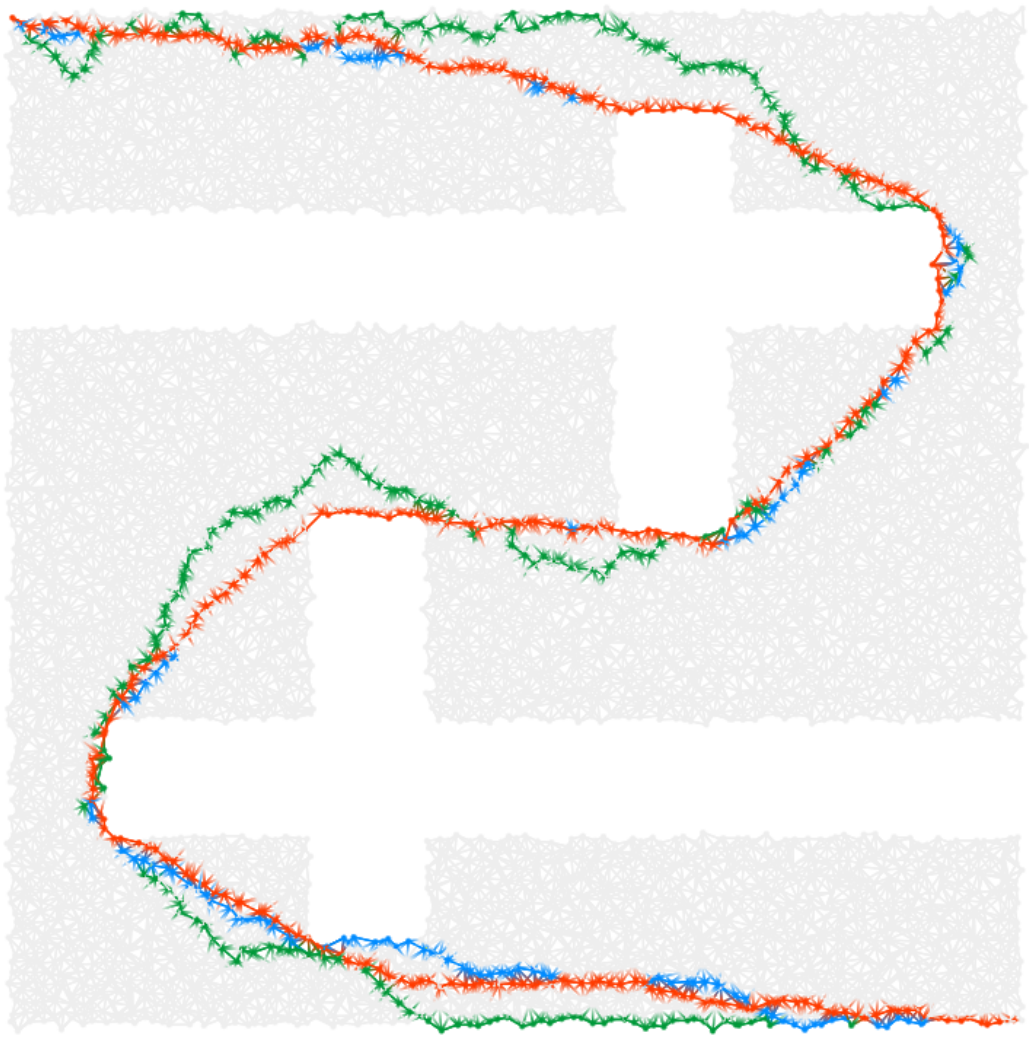

Benchmarking: Path length $\leftrightarrow$ Threshold update

Threshold update

- threshold += **0.1**
- threshold += **1**
- threshold += **10**

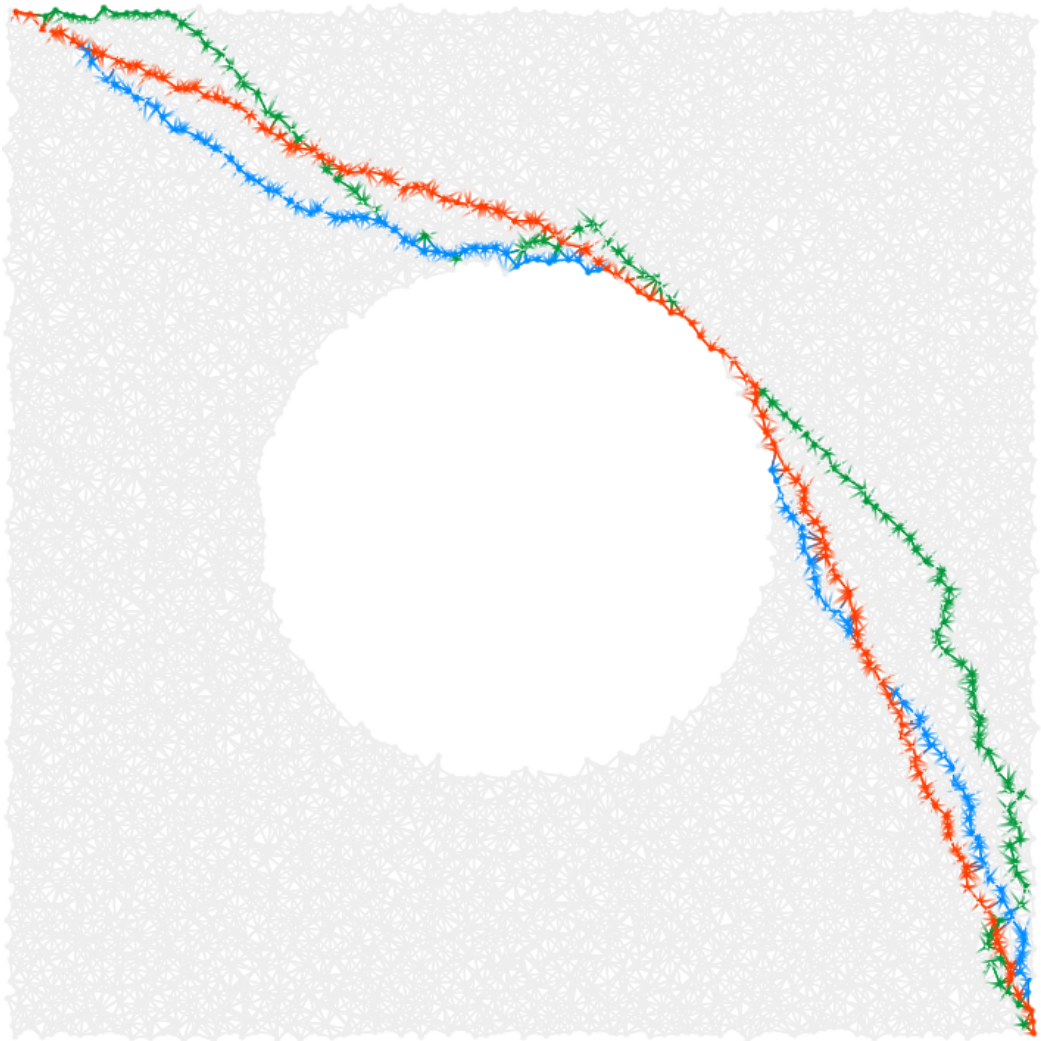

Benchmarking: Path length $\leftrightarrow$ Threshold update

Threshold update

- threshold += **0.1**
- threshold += **1**
- threshold += **10**

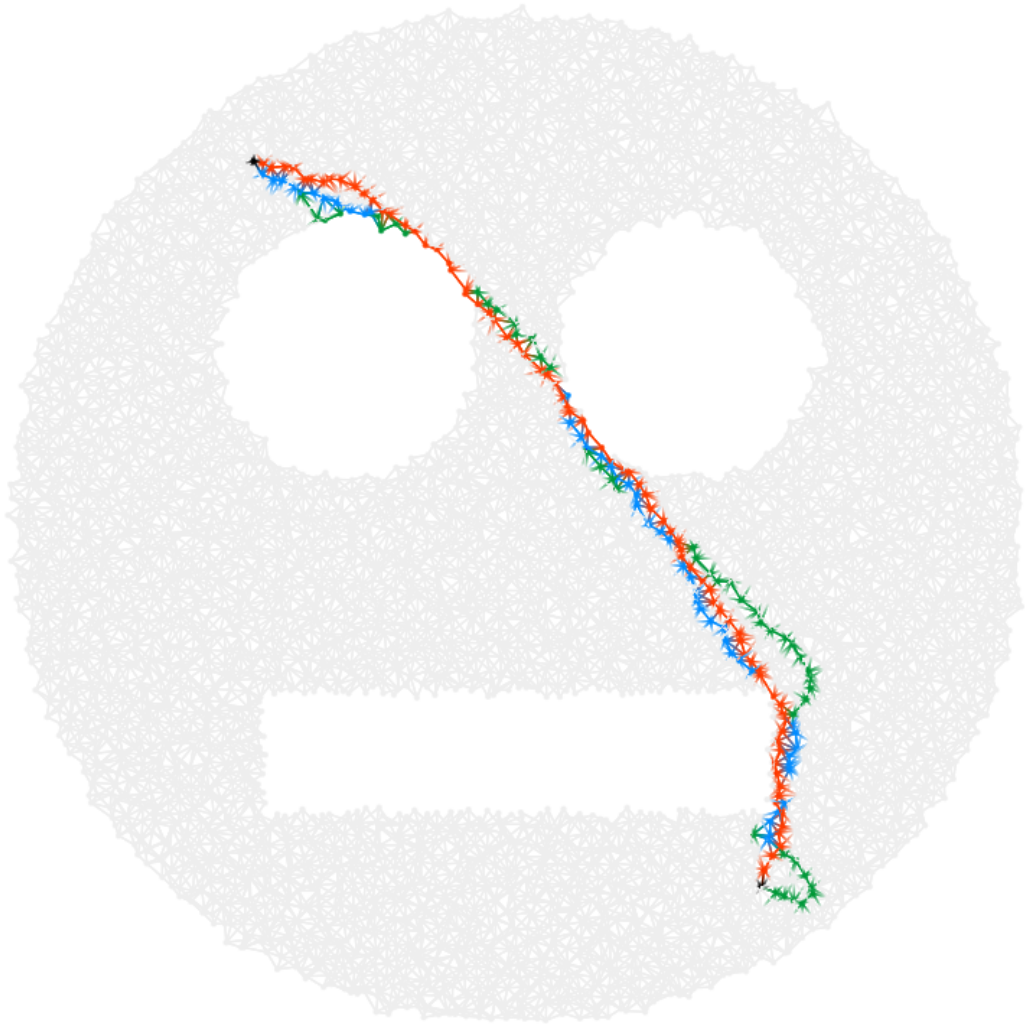

Benchmarking: Path length ↔ Threshold update

compared to A* from Boost Graph Library

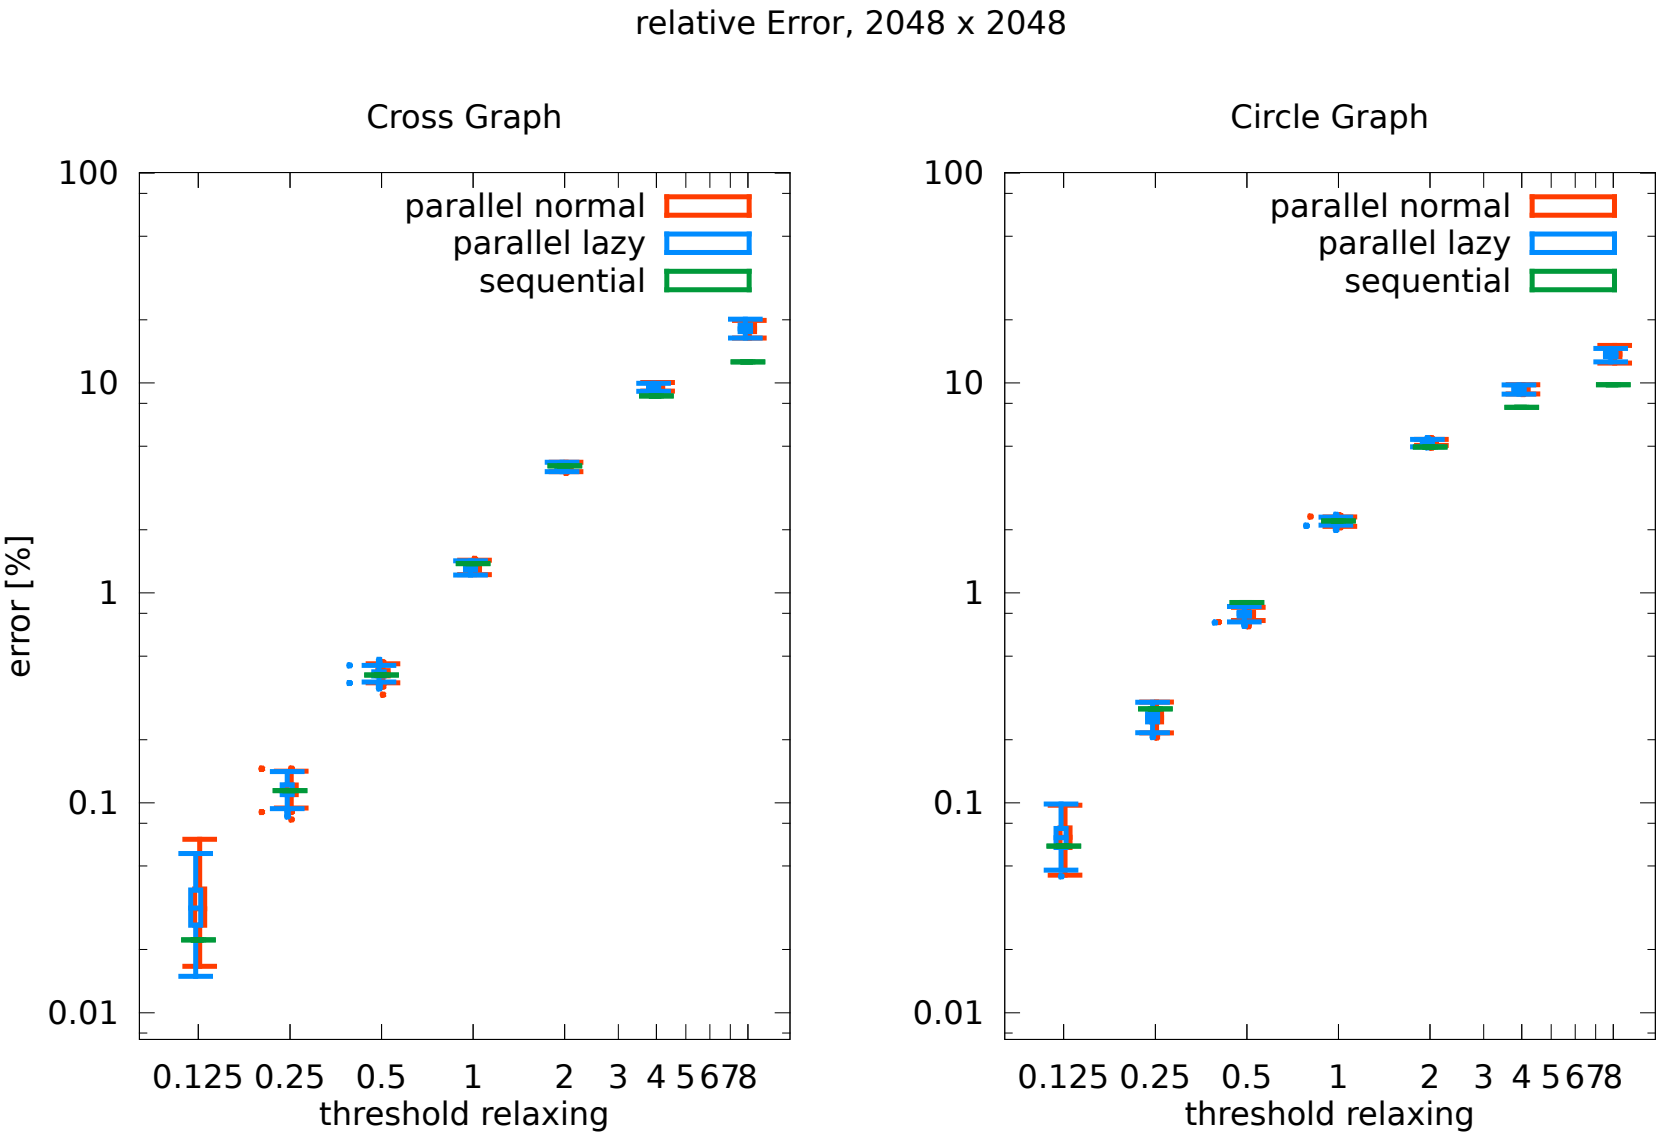

Benchmarking: Run time $\leftrightarrow$ Threshold update

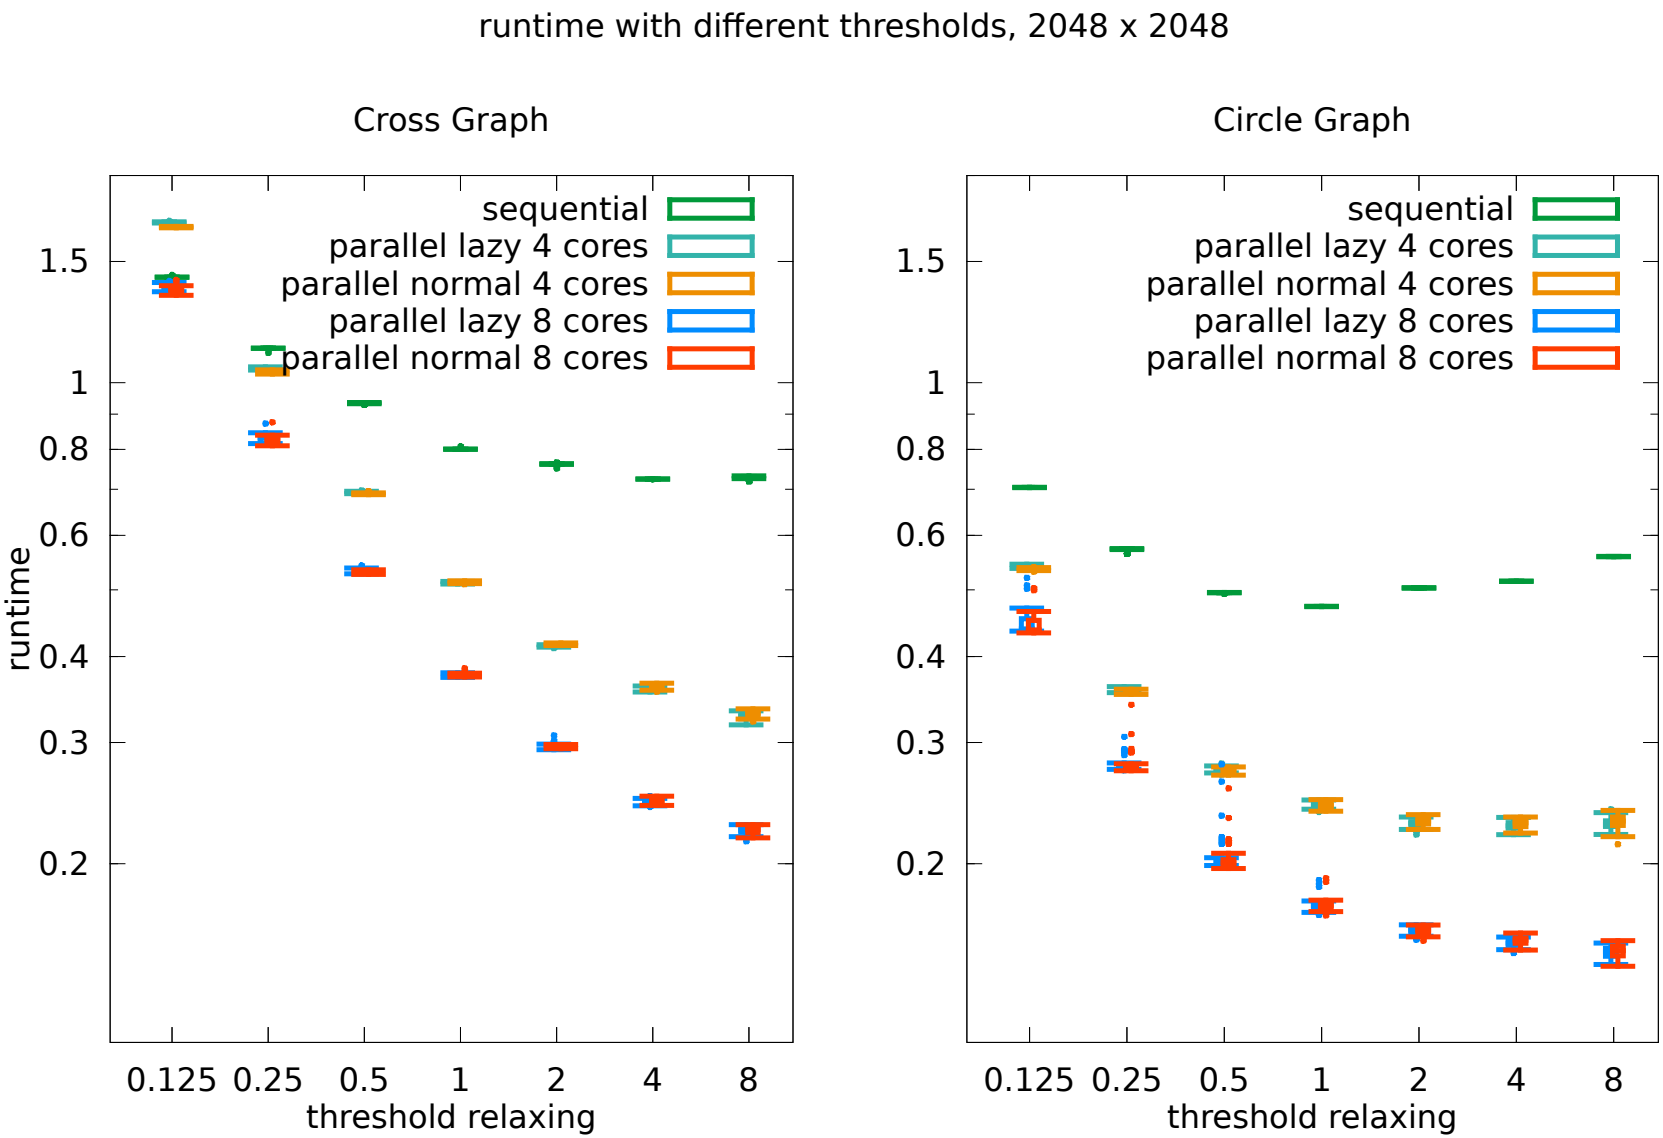

Conclusions

In general Fringe Search is a good single source shortest path algorithm, that can be well implemented in parallel.

- Path quality is not dependent of # cores
- Good strong scaling
- Weak scaling is not perfect
- quality $\leftrightarrow$ runtime trade-off can be tuned for desired result

References / Related Work

Yngvi Björnsson, Markus Enzenberger, Robert C. Holte and Jonathan Schaeffer (2005)

Fringe Search: Beating A* at Pathfinding on Game Maps

Proceedings of the 2005 IEEE Symposium on Computational Intelligence and Games (CIG05), Essex University, Colchester, Essex, UK, 4-6 April, 2005

Sandy Brand (2009)

Efficient obstacle avoidance using autonomously generated navigation meshes

Master Thesis (Delft University of Technology)

Sandy Brand and Rafael Bidarra (2012)

Multi-core scalable and efficient pathfinding with Parallel Ripple Search

Computer Animation and Virtual Worlds, Volume 23, Issue 2 2012, pp 73 – 85.

The End