Developing and Evaluating Methods for Distributing Observations in DART

Kamil Yousuf

Mentors: Helen Kershaw, Marlee Smith

Summer Internships in Parallel Computational Science

Data Assimilation Research Section

July 30, 2024

This material is based upon work supported by the National Center for Atmospheric Research, which is a major facility sponsored by the National Science Foundation under Cooperative Agreement No. 1852977.

Data Assimilation Research Testbed (DART)

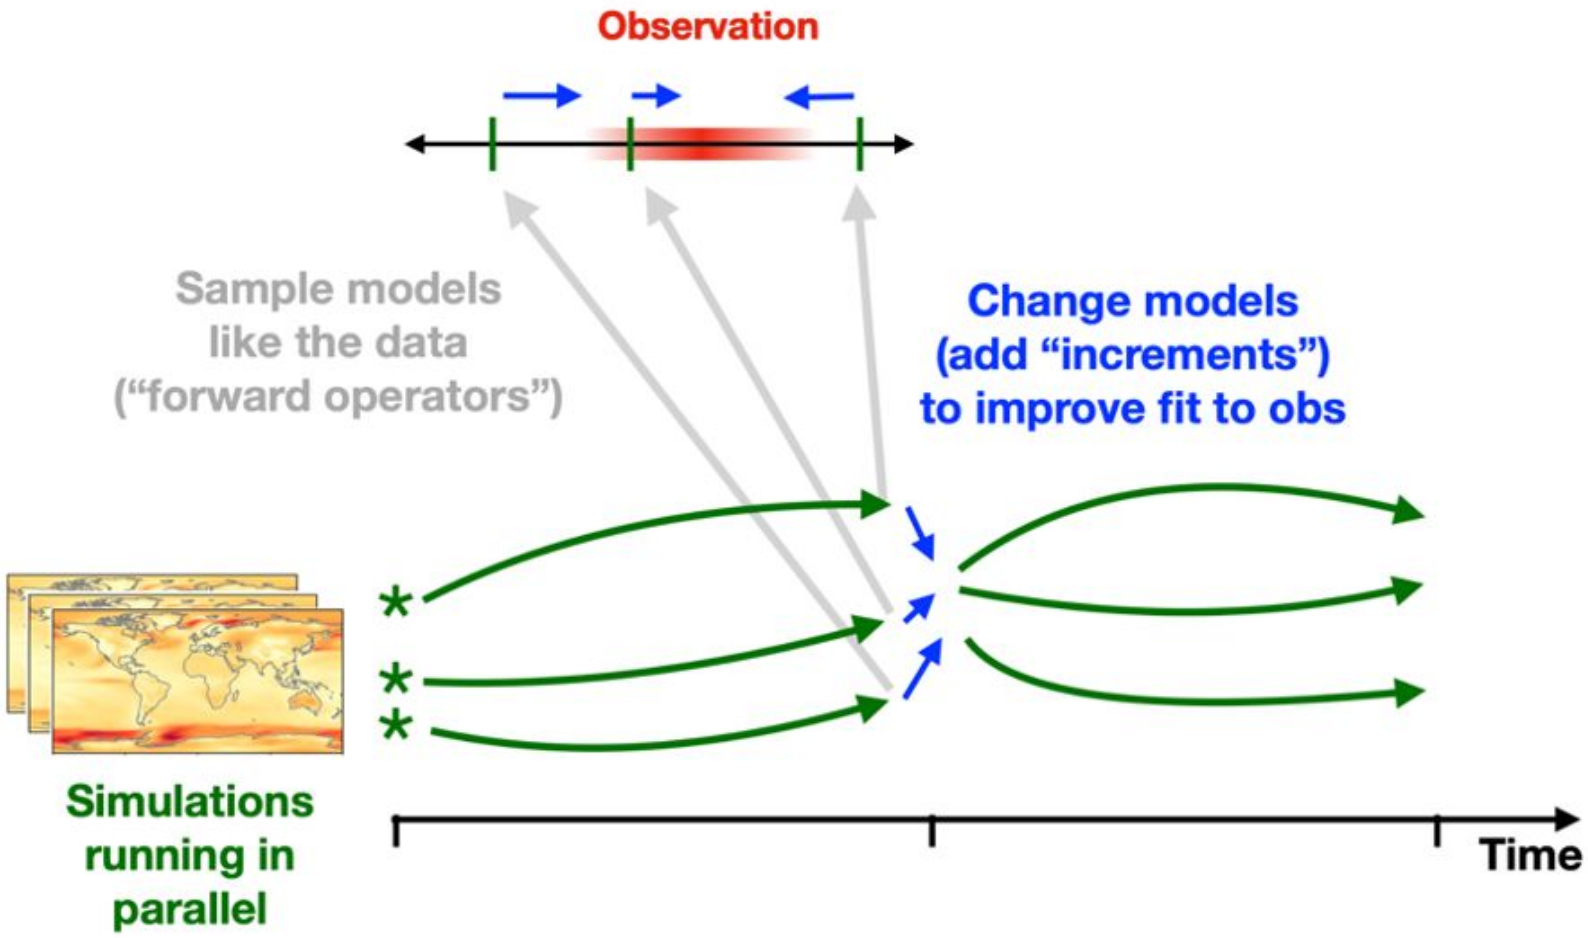

Data Assimilation Research Testbed (DART)

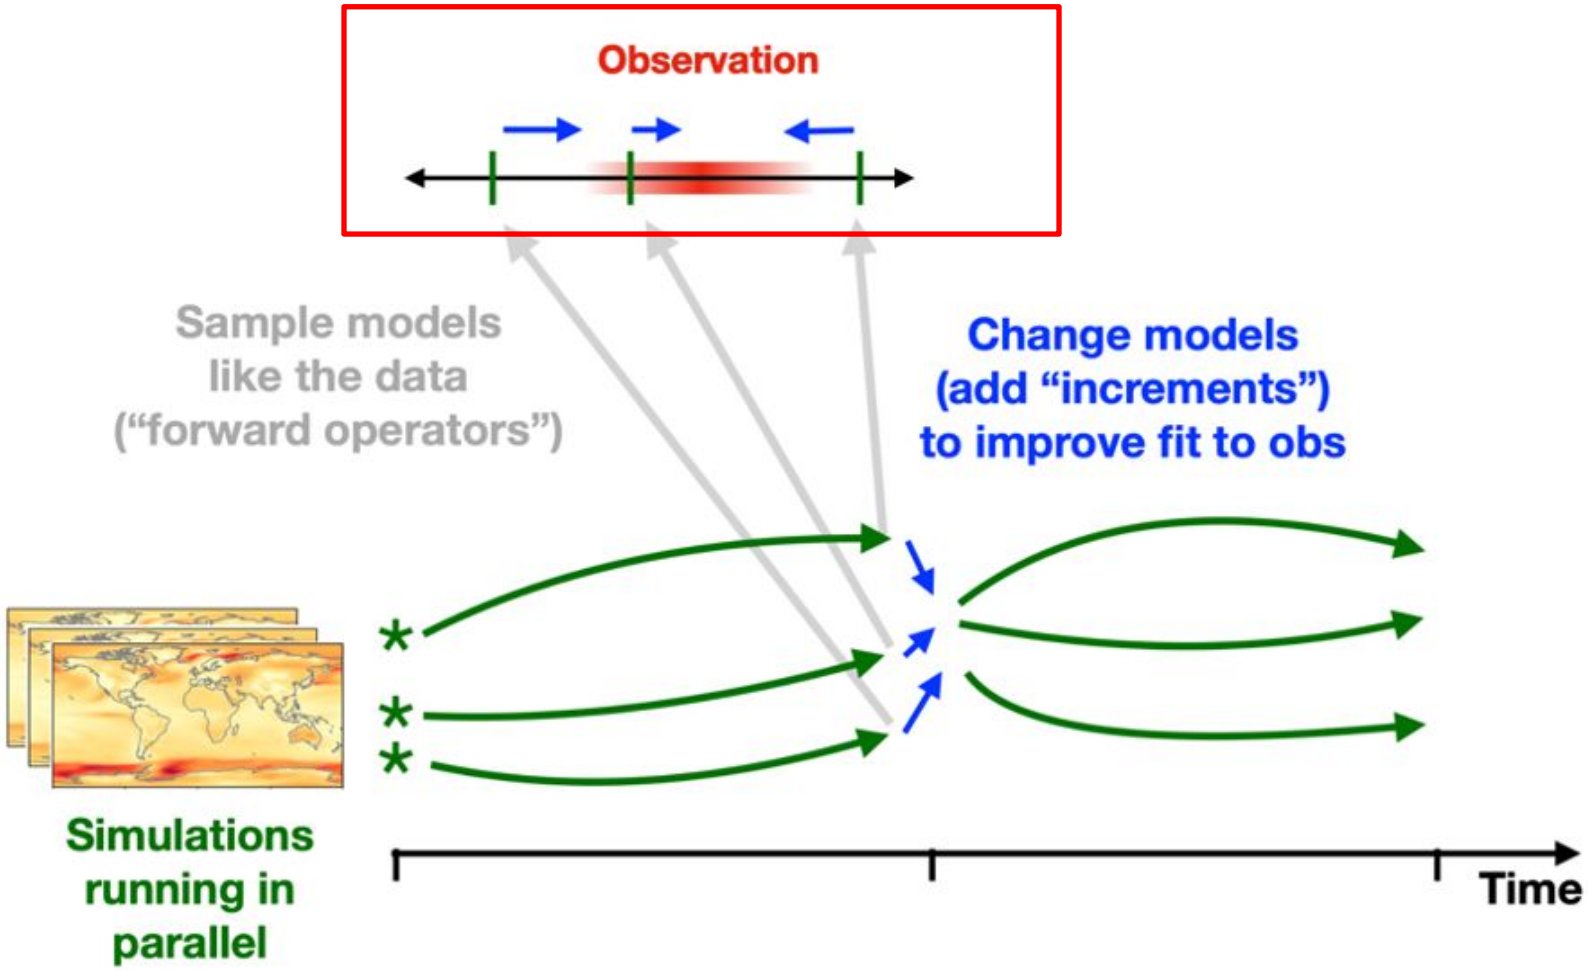

Observations: what are they?

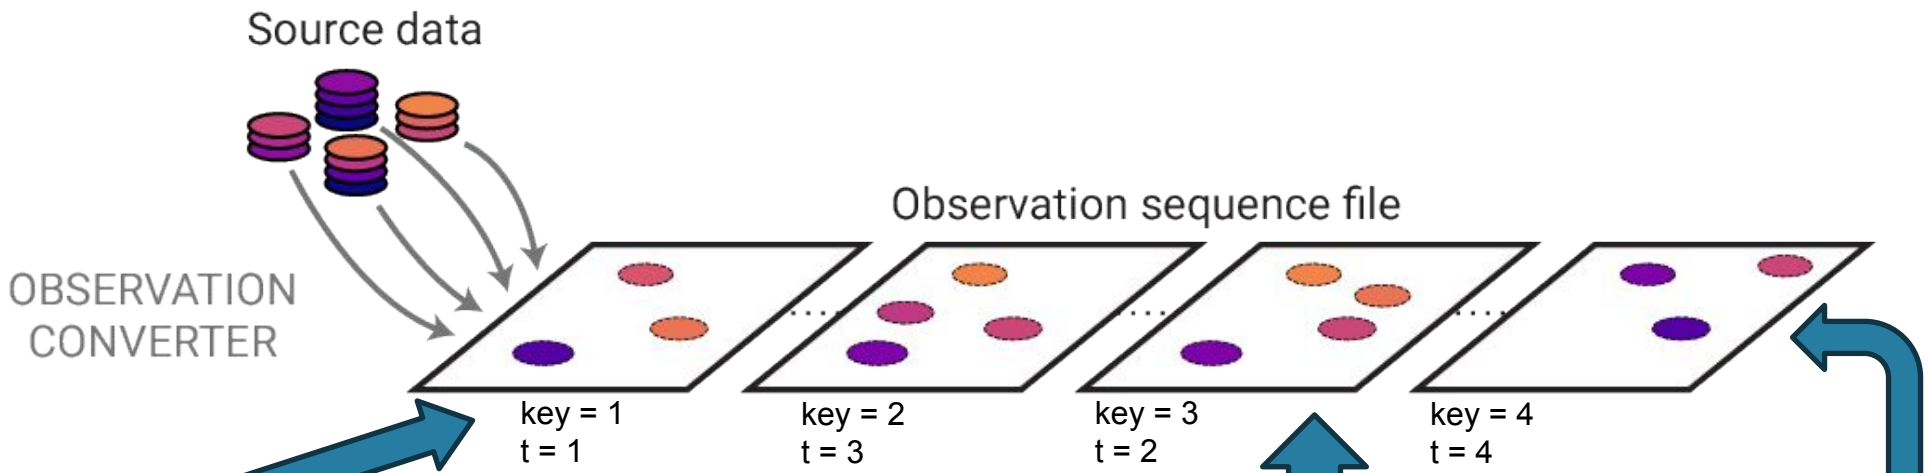

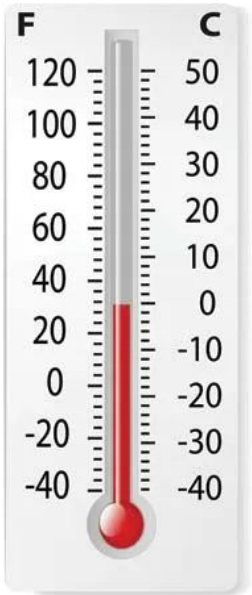

Observations: what are they?

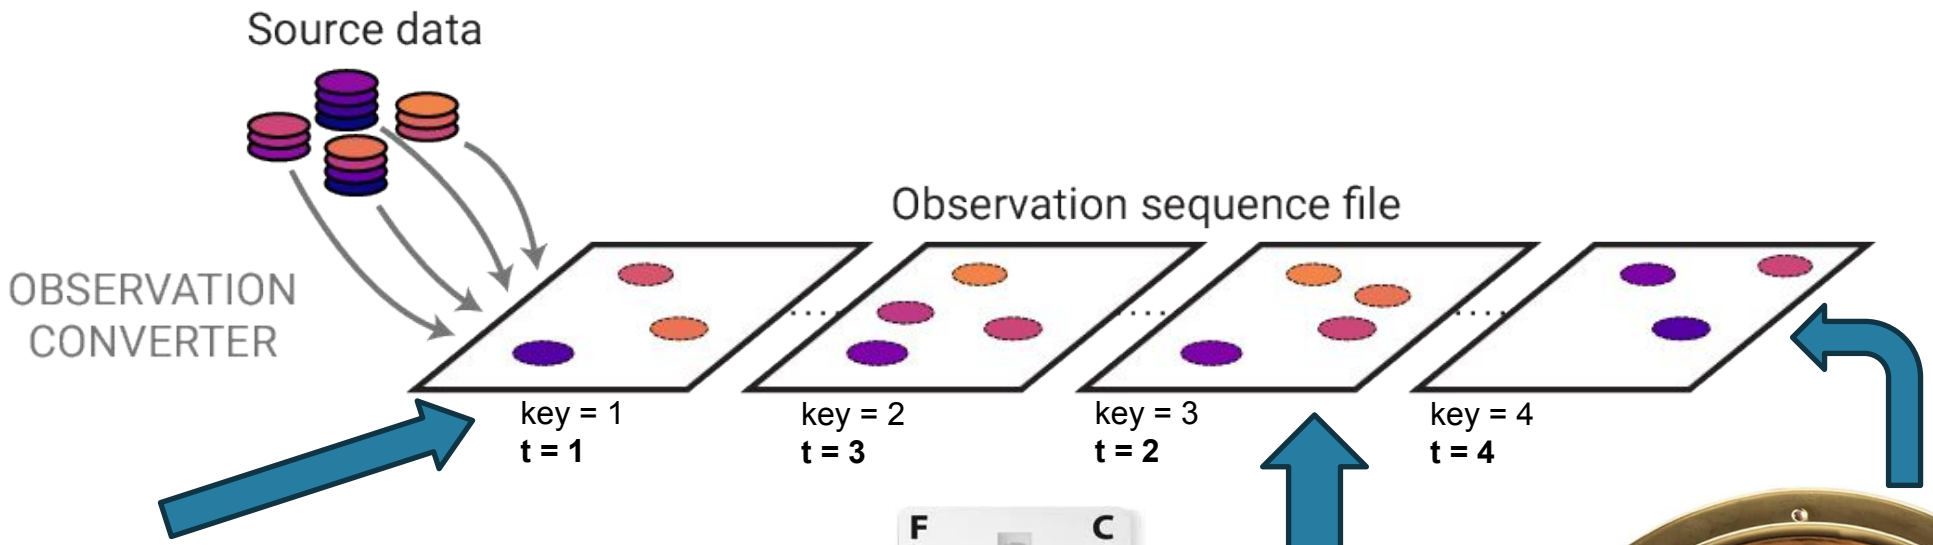

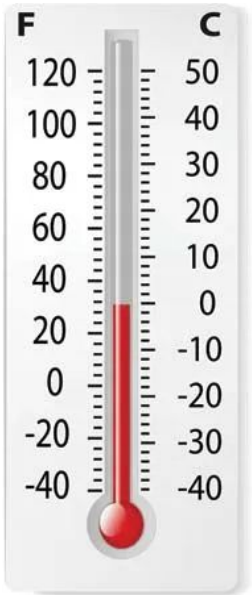

Observations: On disk vs. in memory representation?

On disk:

In memory:

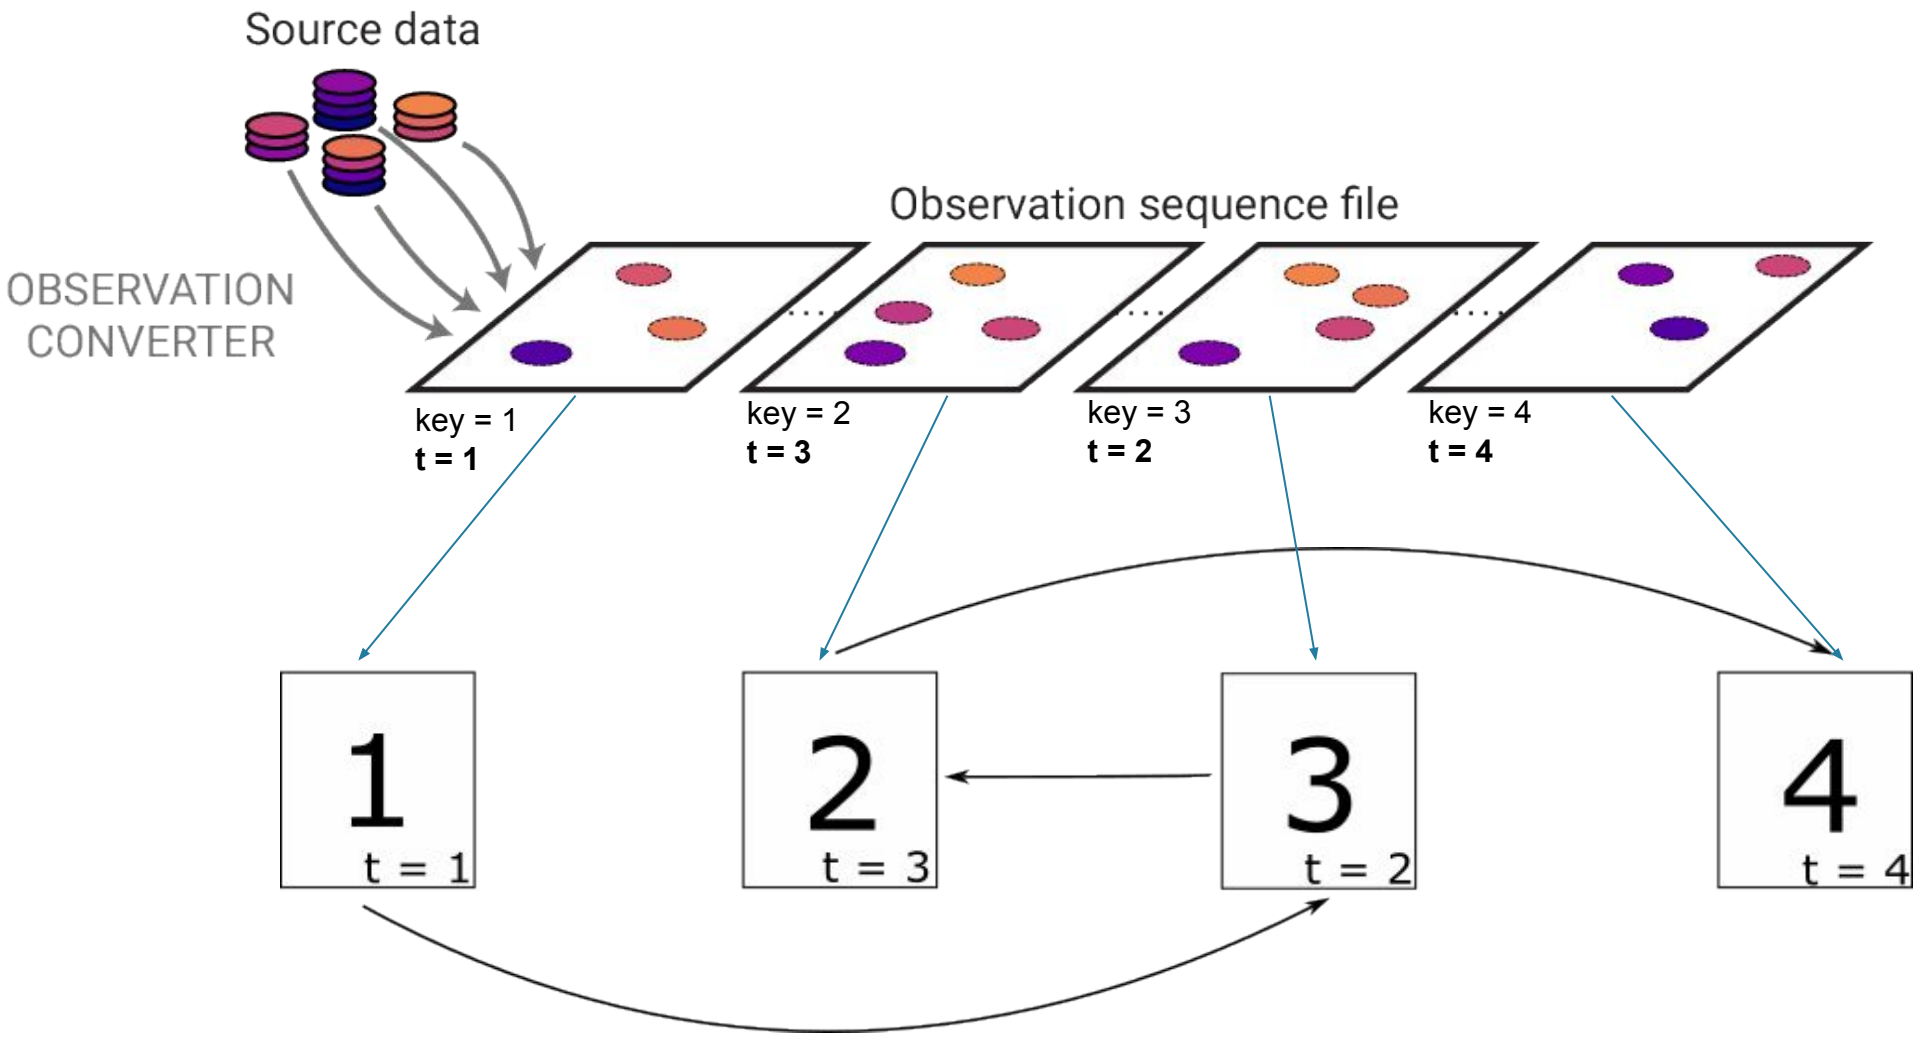

Observations: Time windows

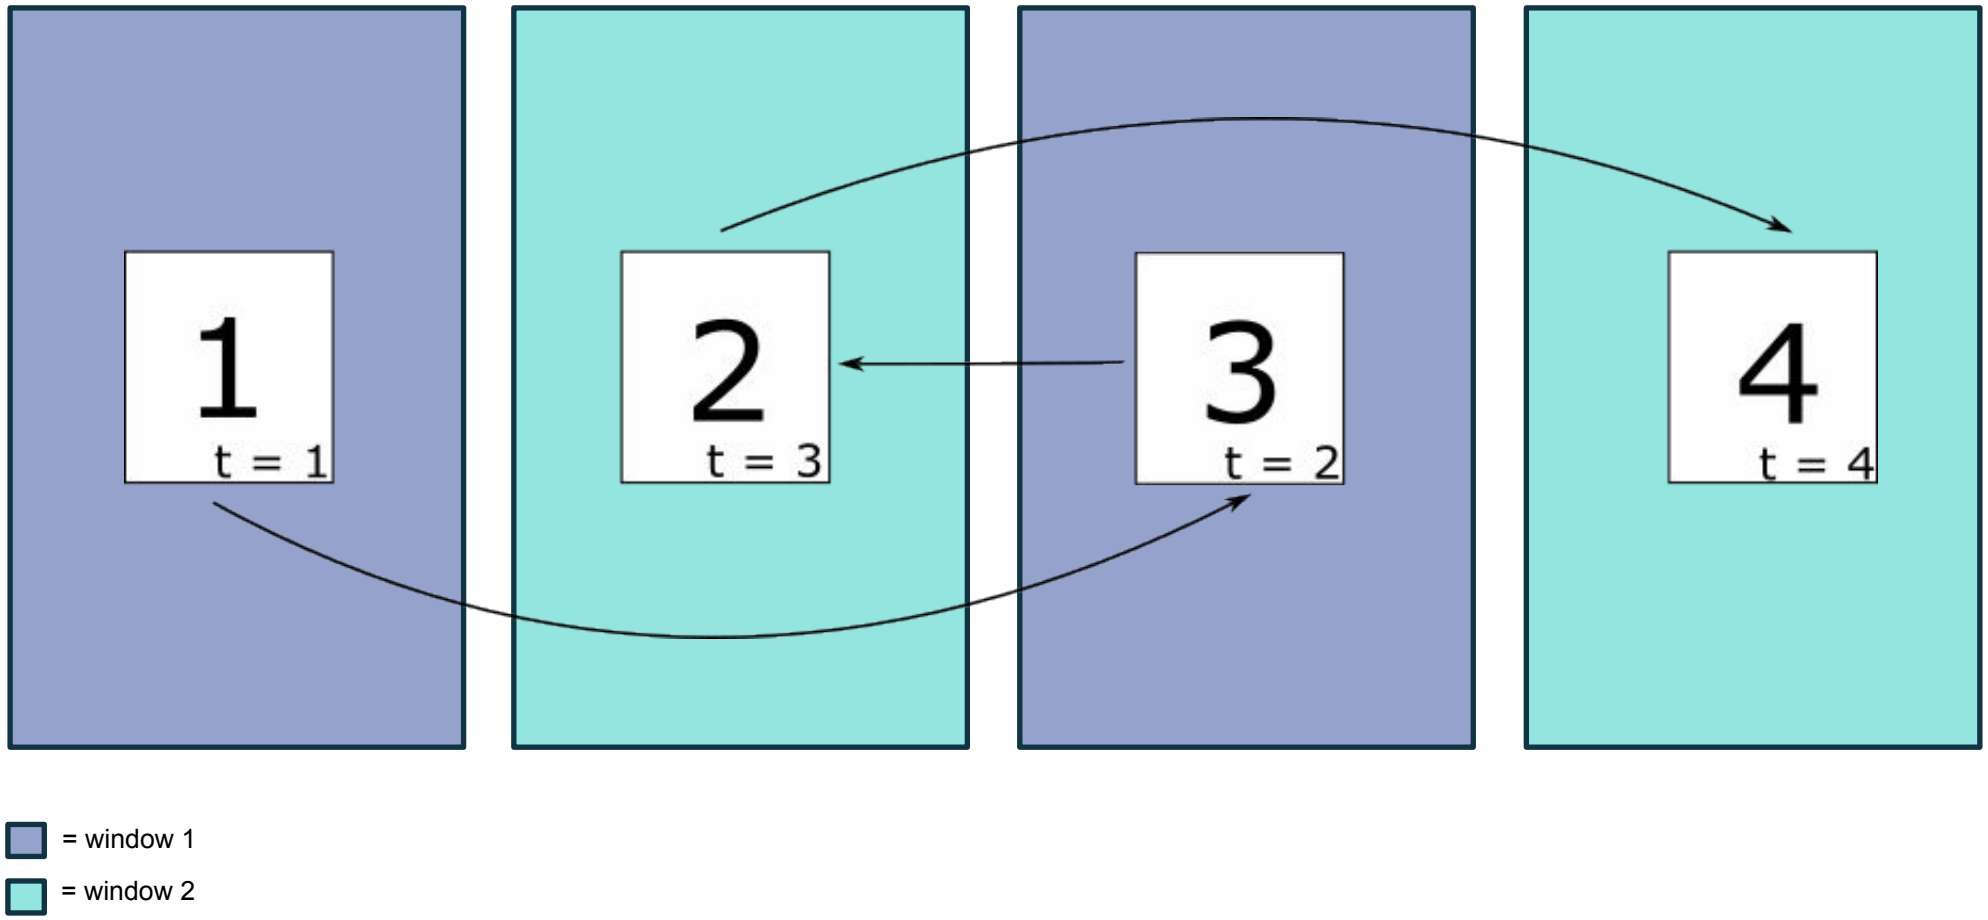

Current Method

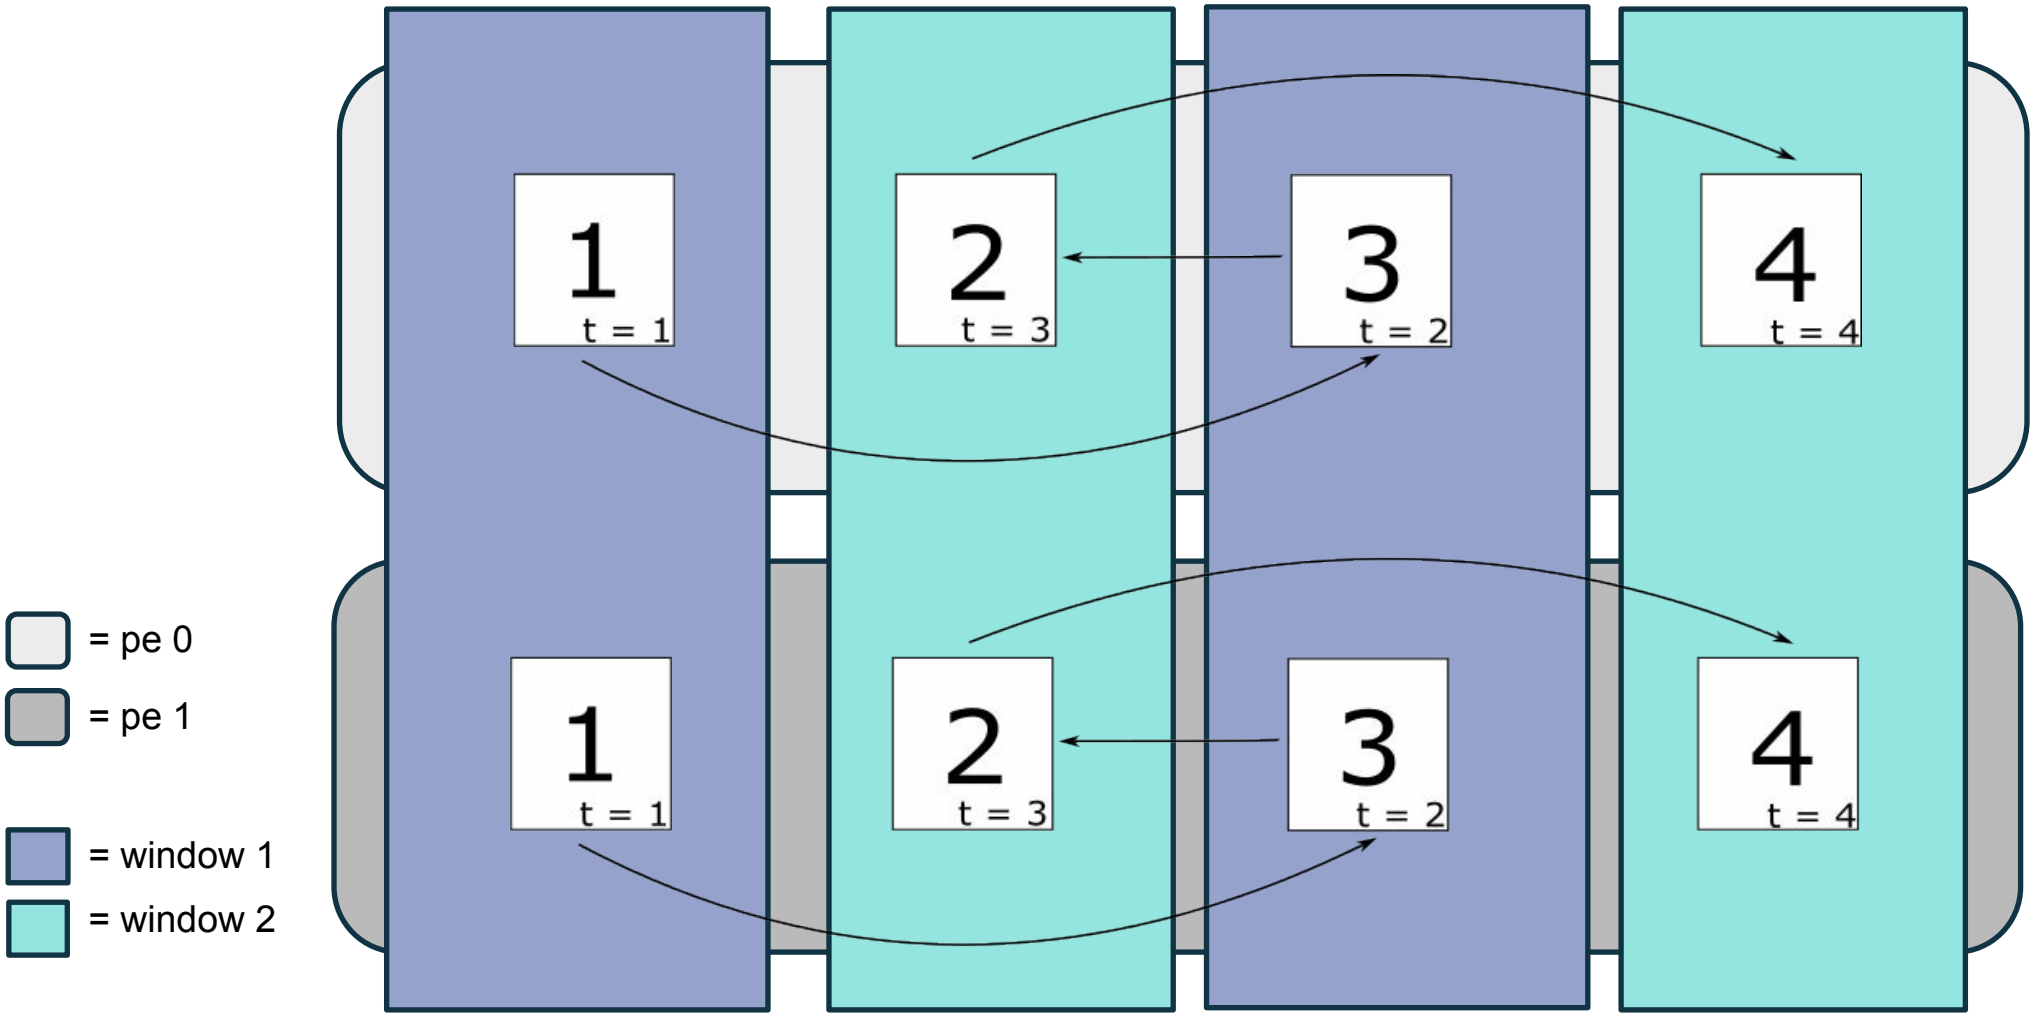

Problems with Current Method

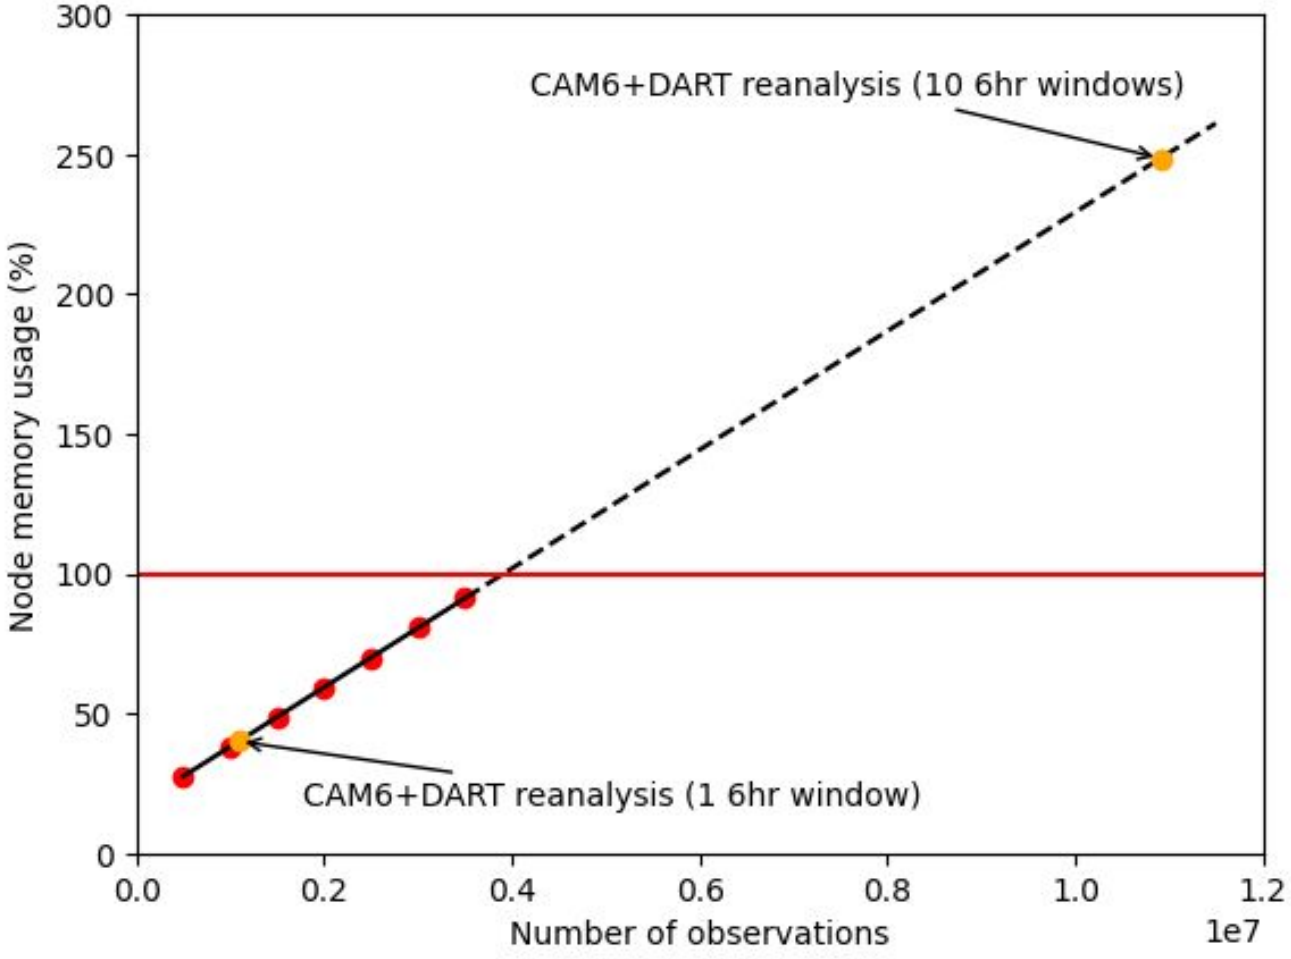

Project Goals

1. Distribute observations across all processes

2. Ensure distribution of observations does not affect performance
   - Reduce communication to a minimum!
   - **Only when absolutely necessary!**

Gather-Sort-Scatter

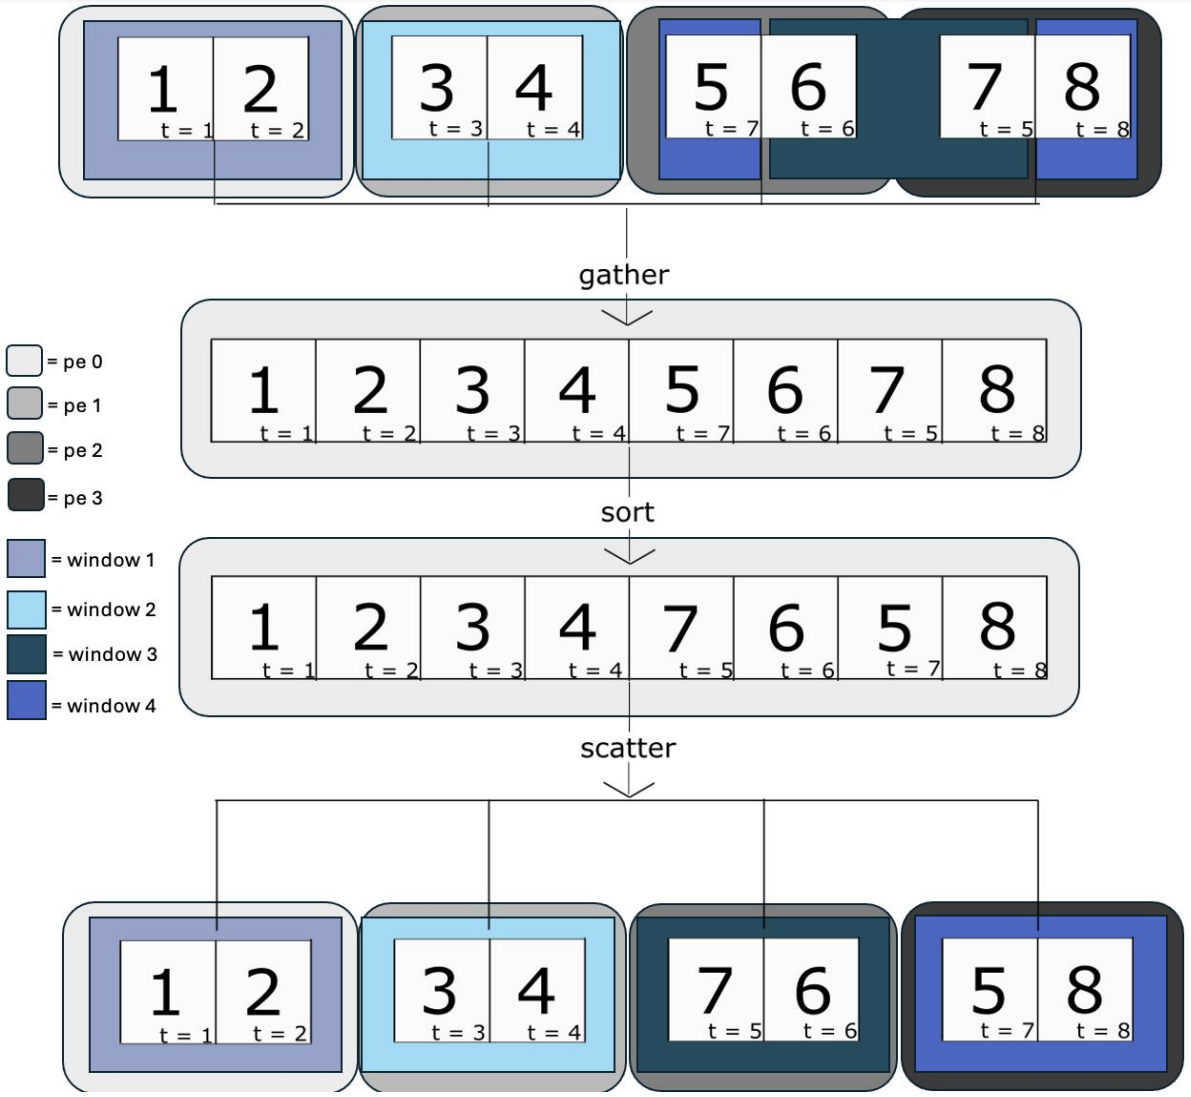

Gather-Sort-Scatter

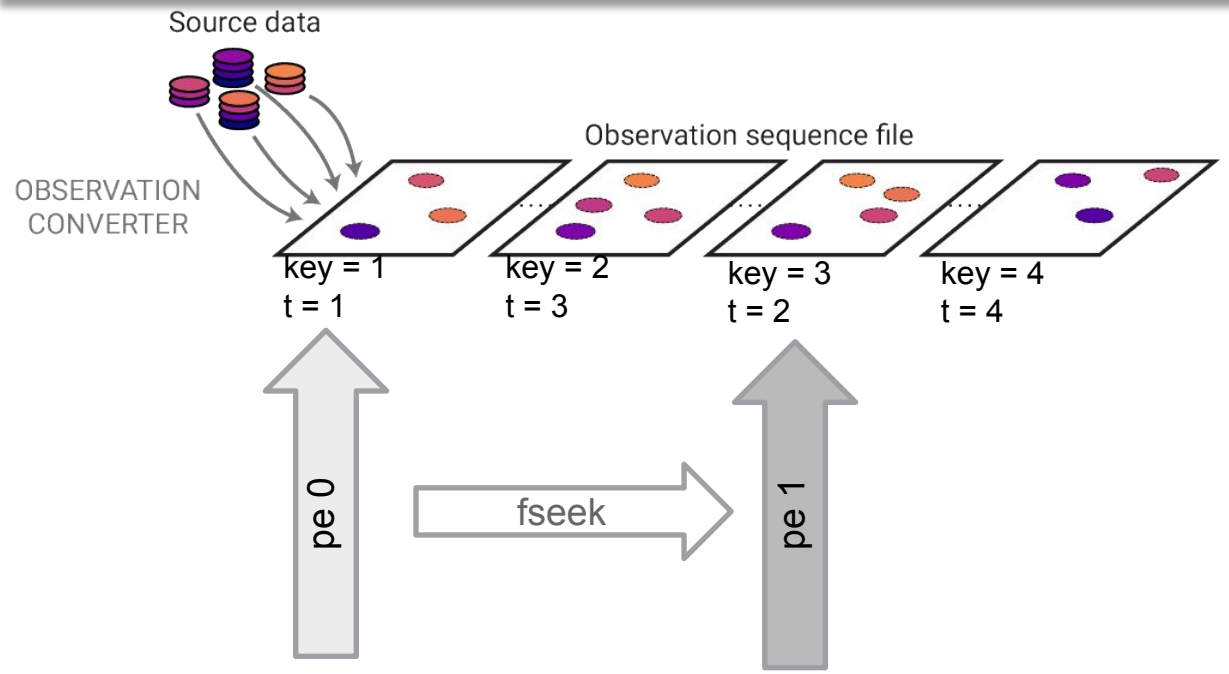

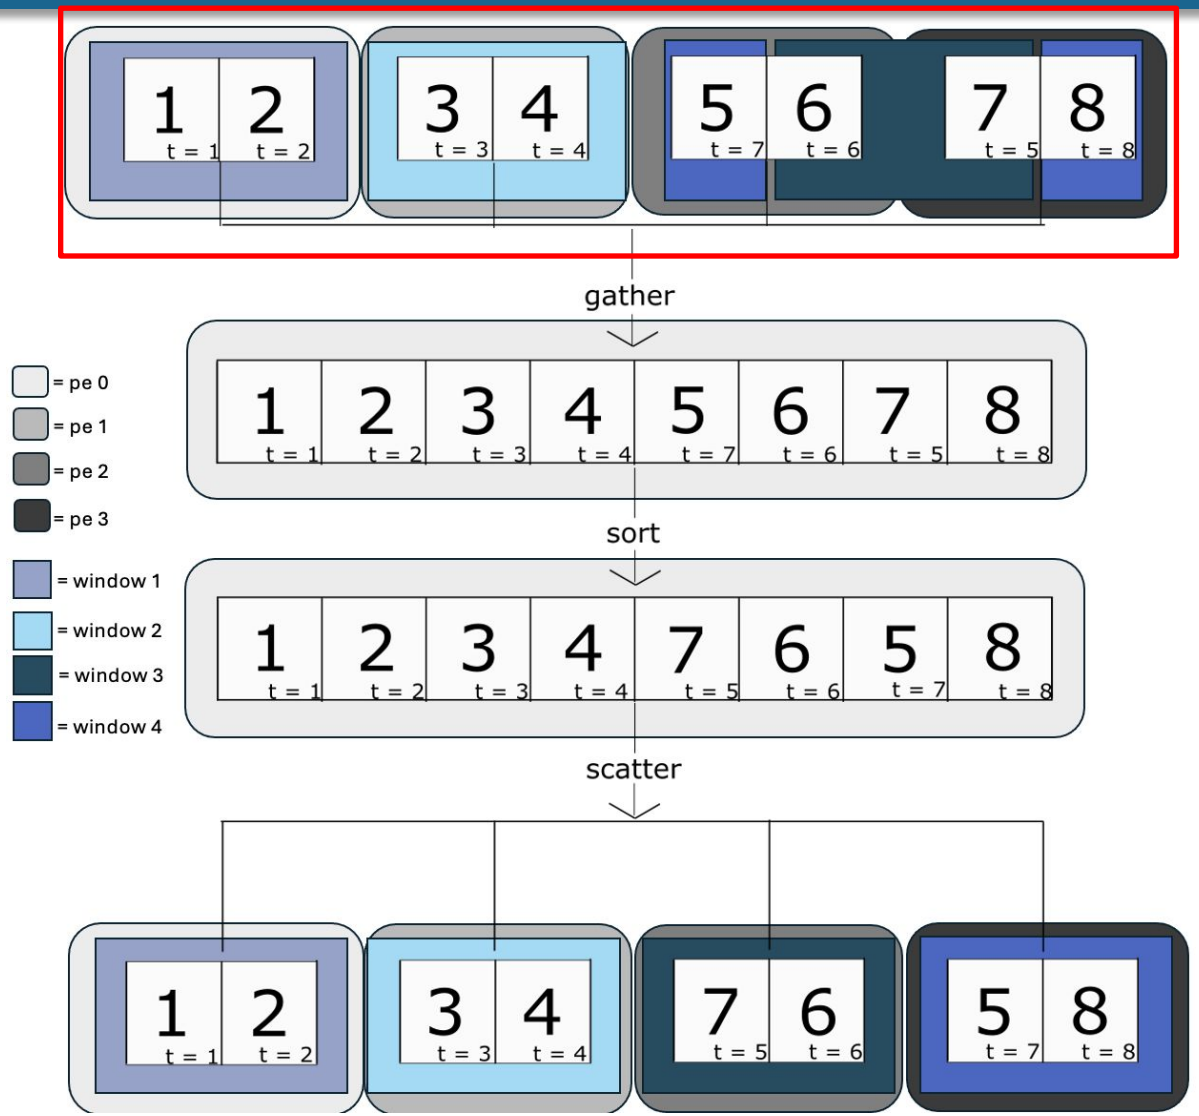

Gather-Sort-Scatter

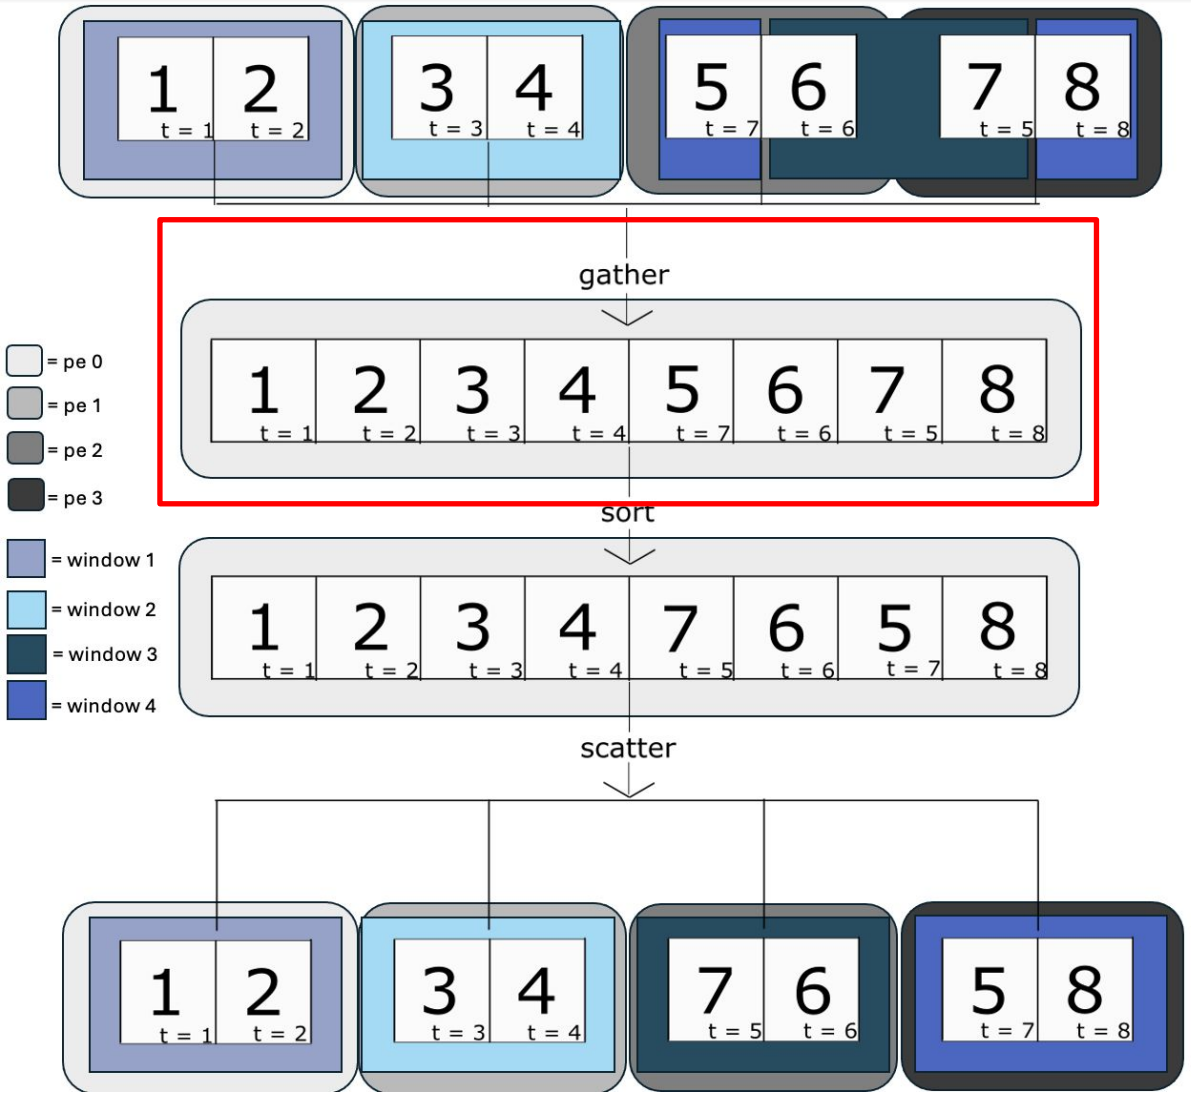

Gather-Sort-Scatter

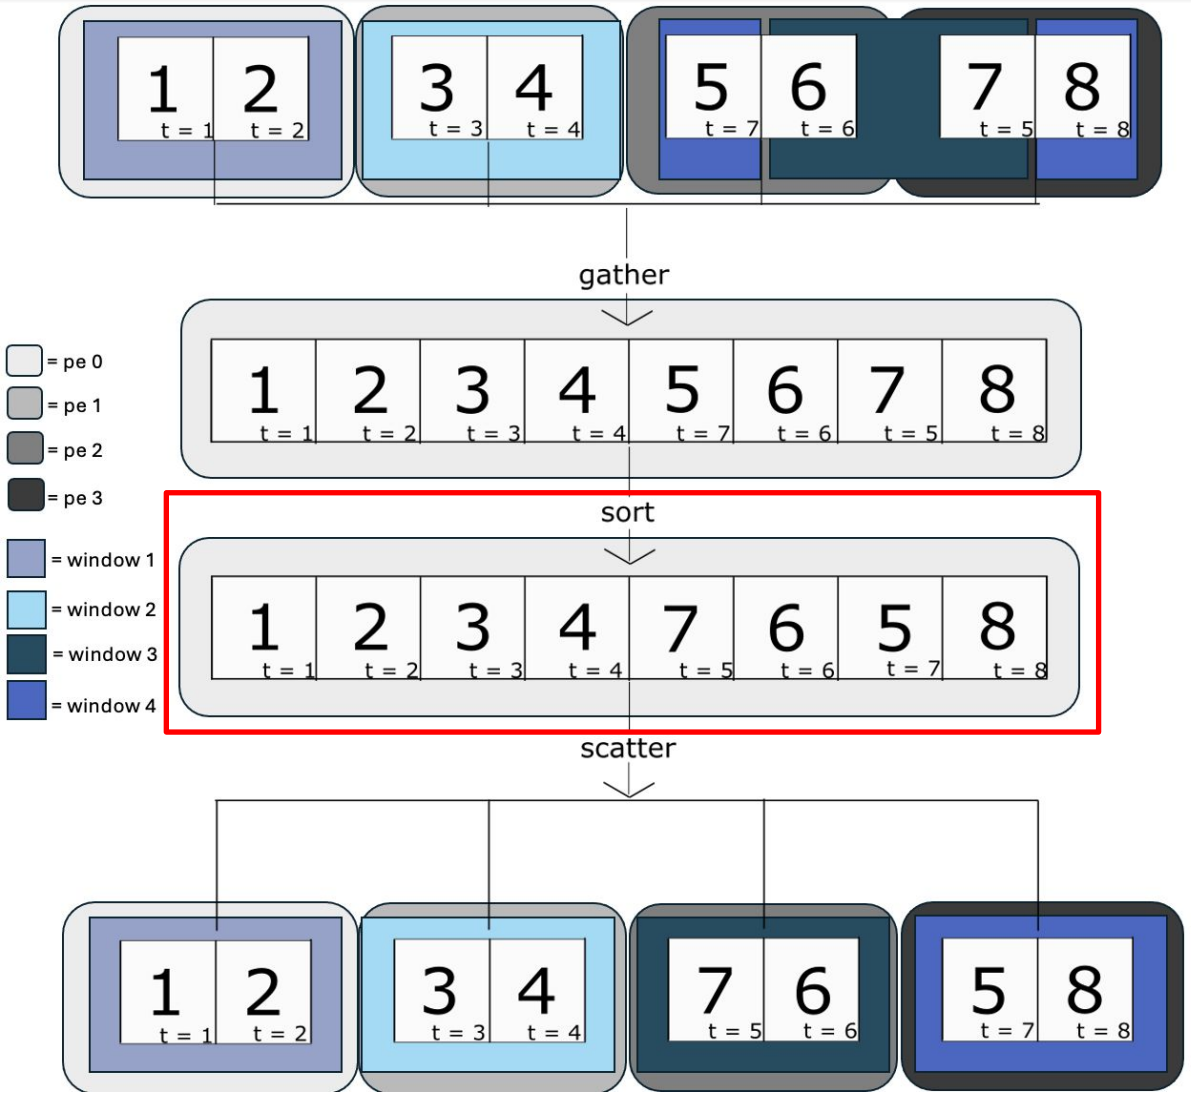

Gather-Sort-Scatter

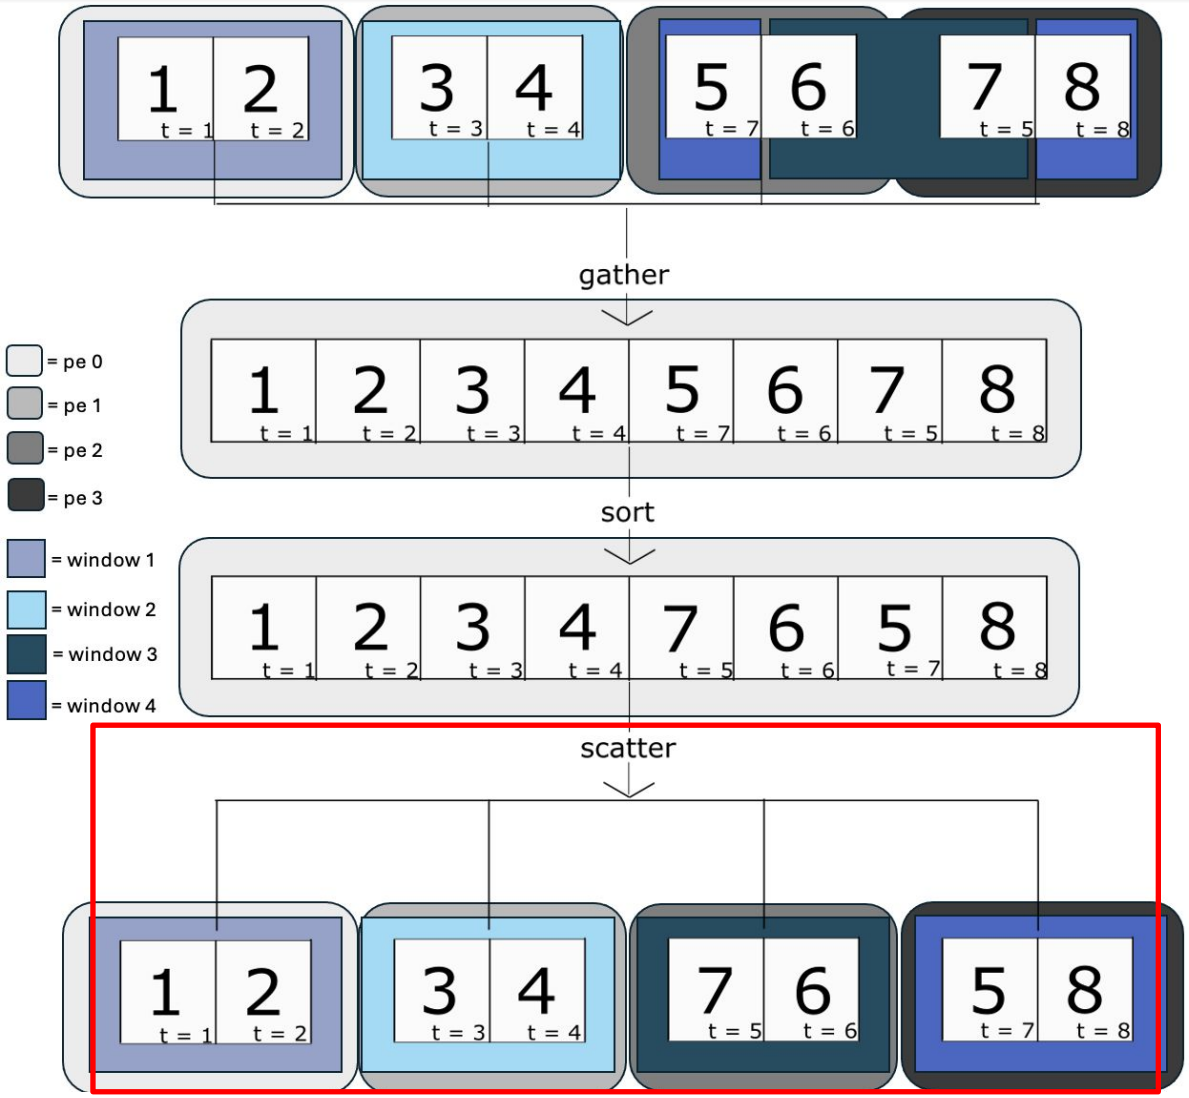

One-sided communication

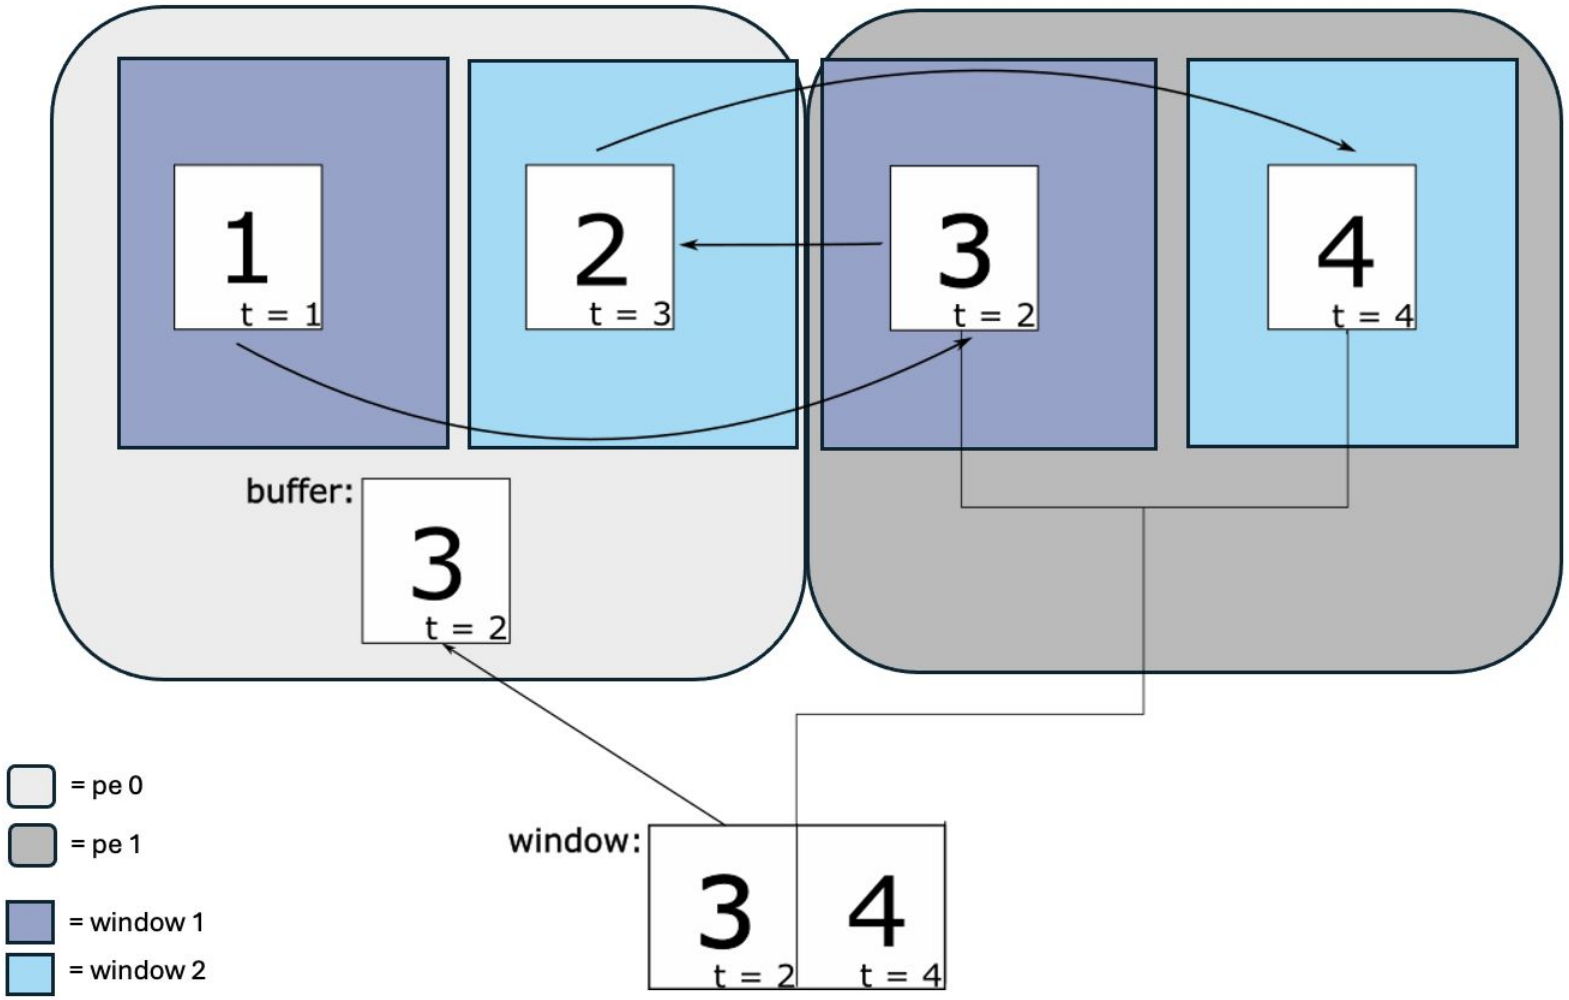

One-sided communication

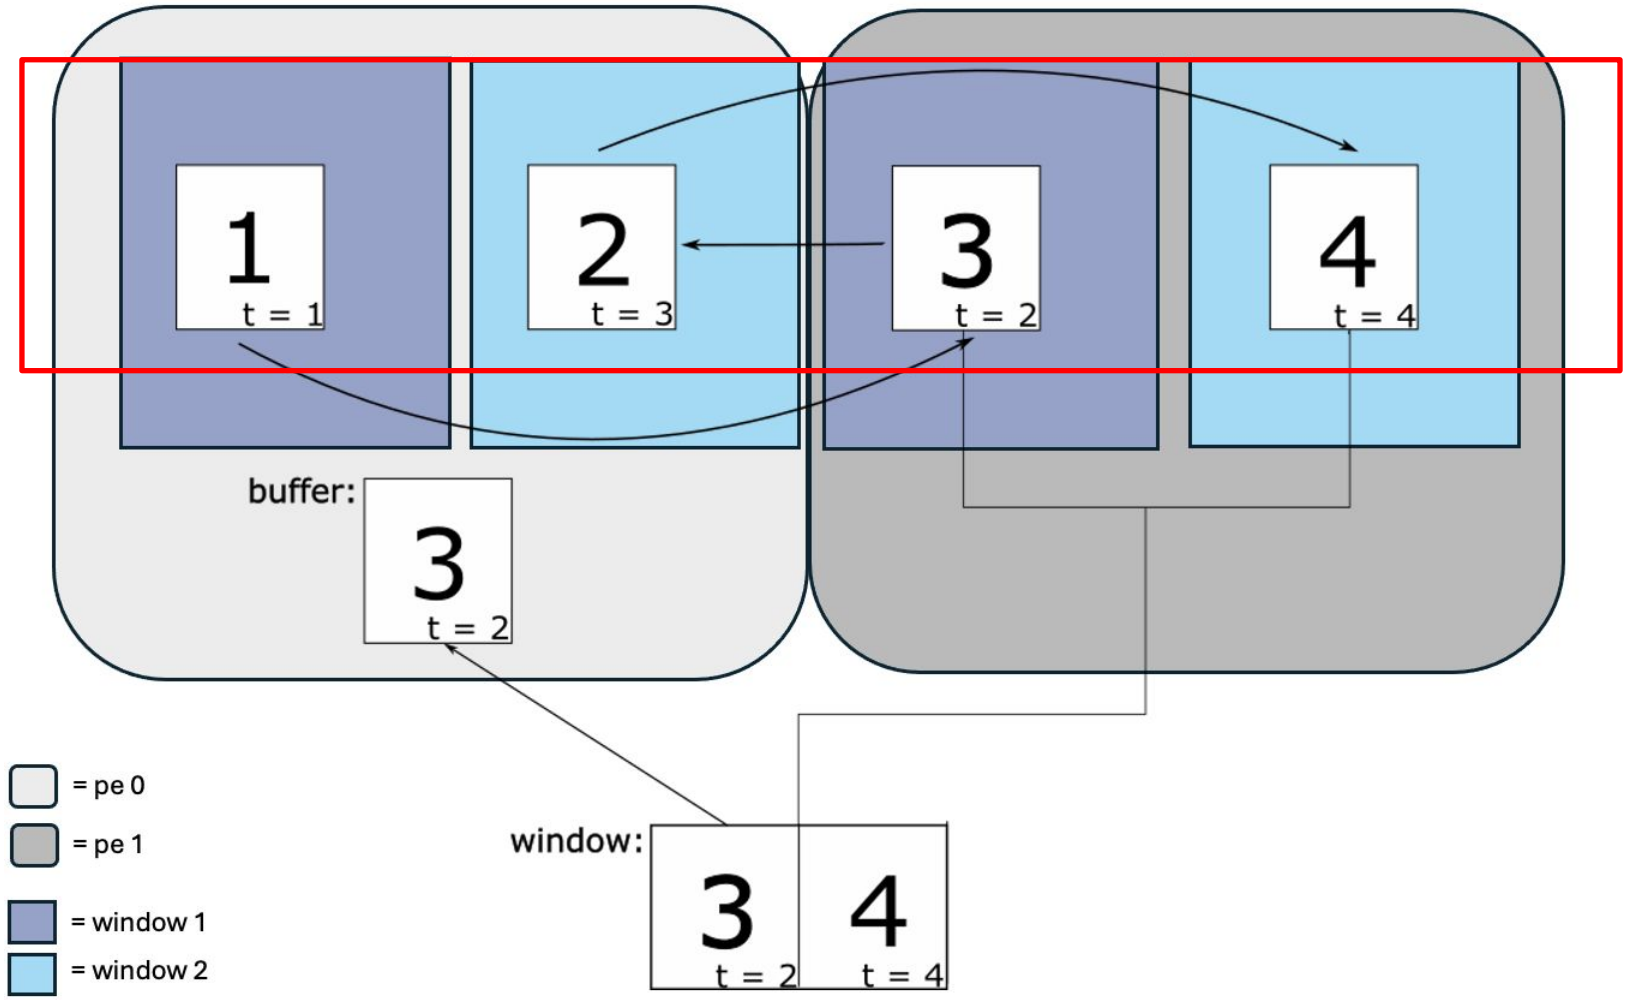

One-sided communication

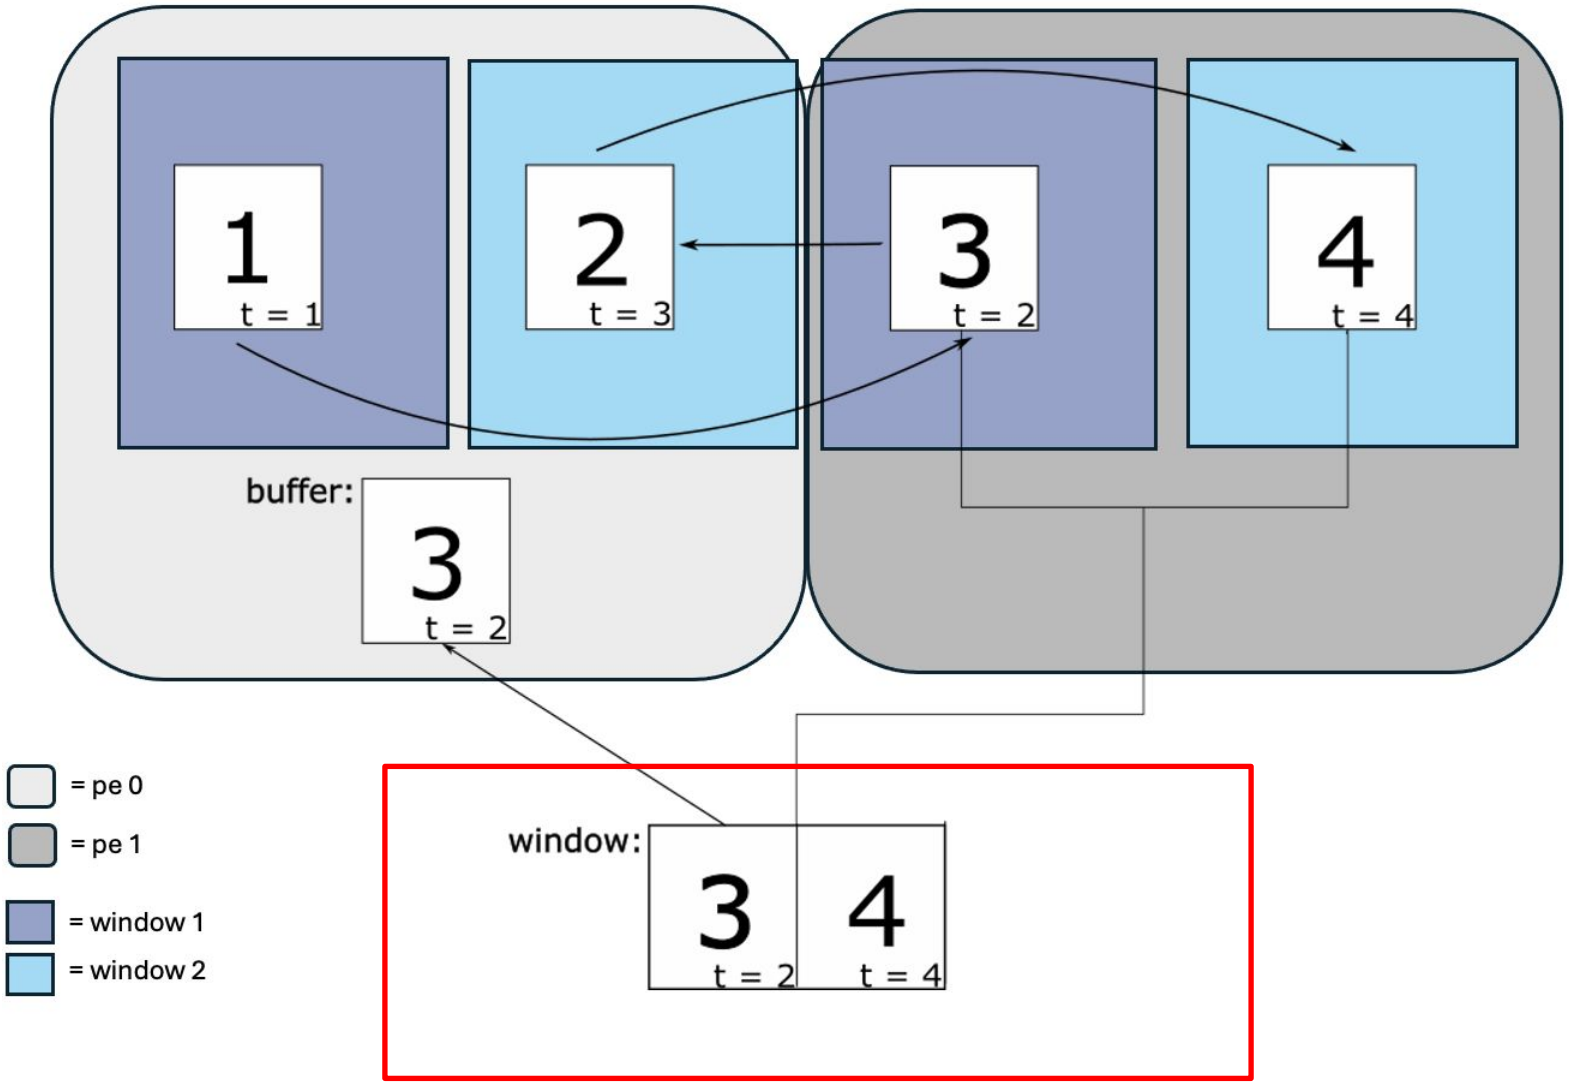

One-sided communication

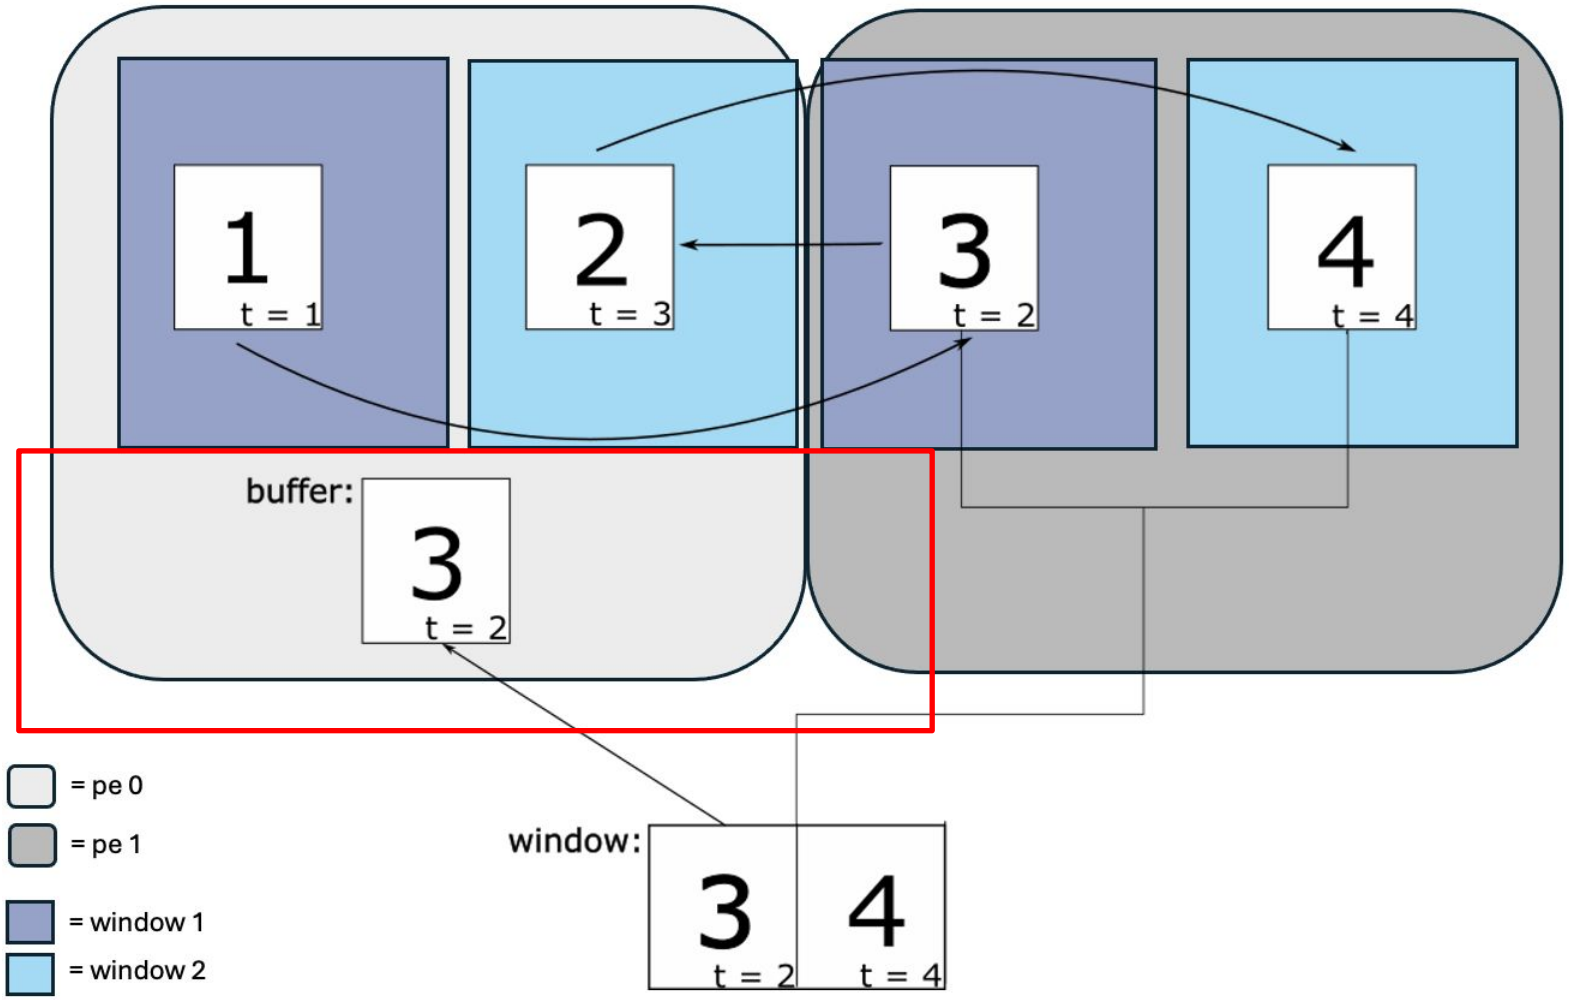

One-sided communication: direct linked list traversal vs. key caching

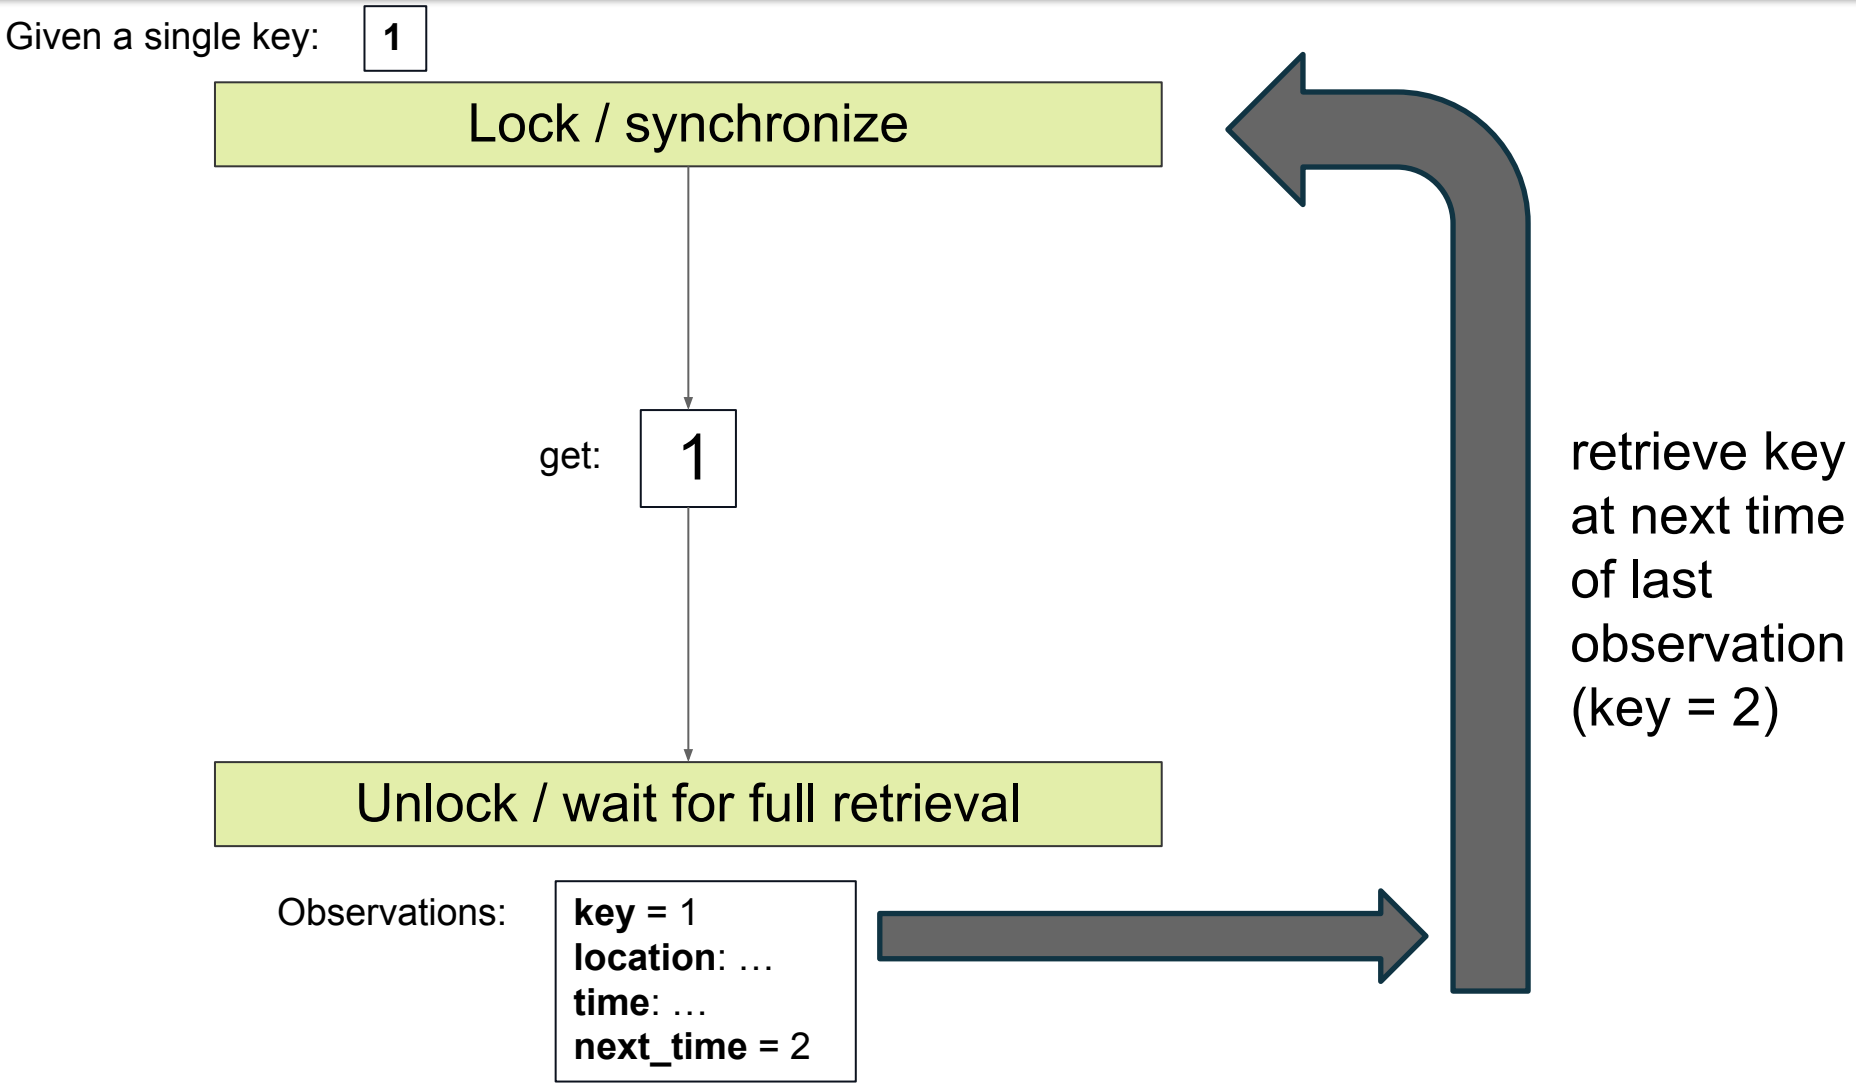

One-sided communication: direct linked list traversal vs. key caching

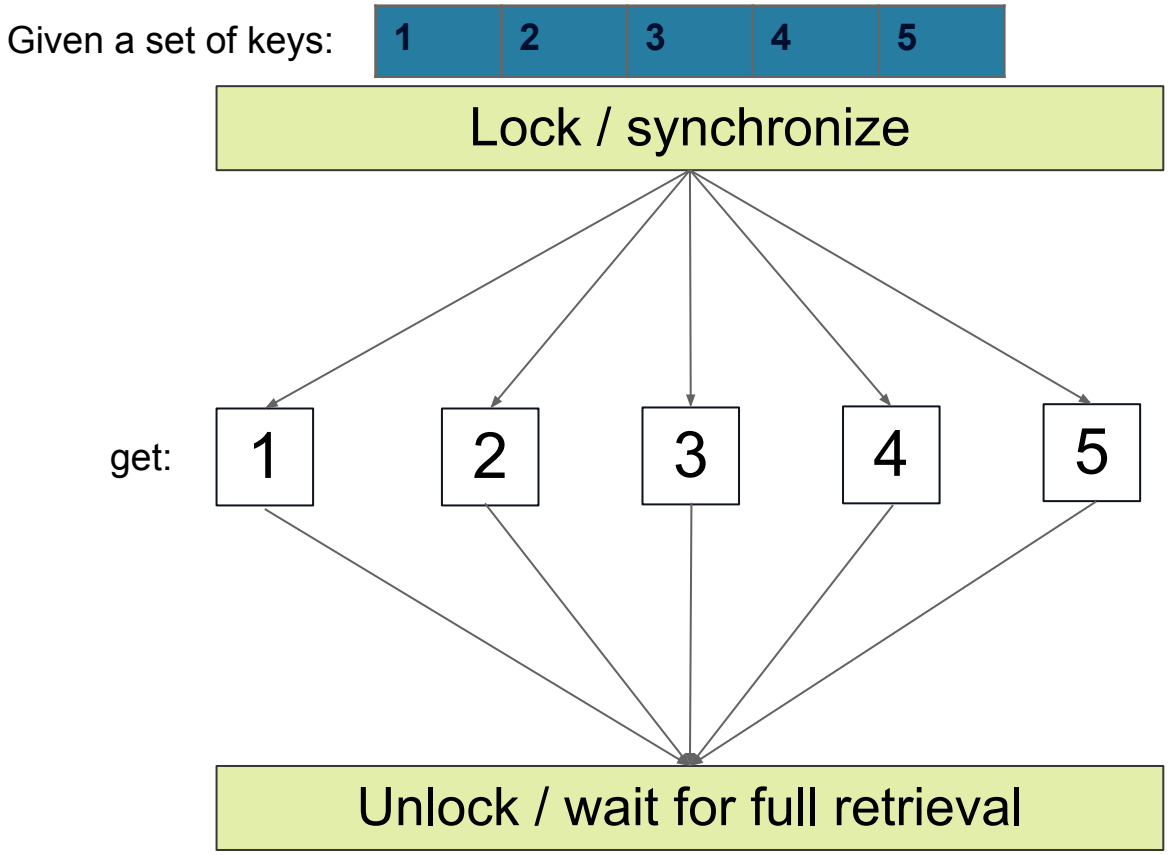

Observations:

key = 1 **location**: … **time**: … **next_time** = 2	**key** = 2 **location**: … **time**: … **next_time** = 3	**key** = 3 **location**: … **time**: … **next_time** = 4	**key** = 4 **location**: … **time**: … **next_time** = 5	**key** = 5 **location**: … **time**: … **next_time** = 6

Gather-Sort-Scatter: scaling / performance

N = 64,950,921

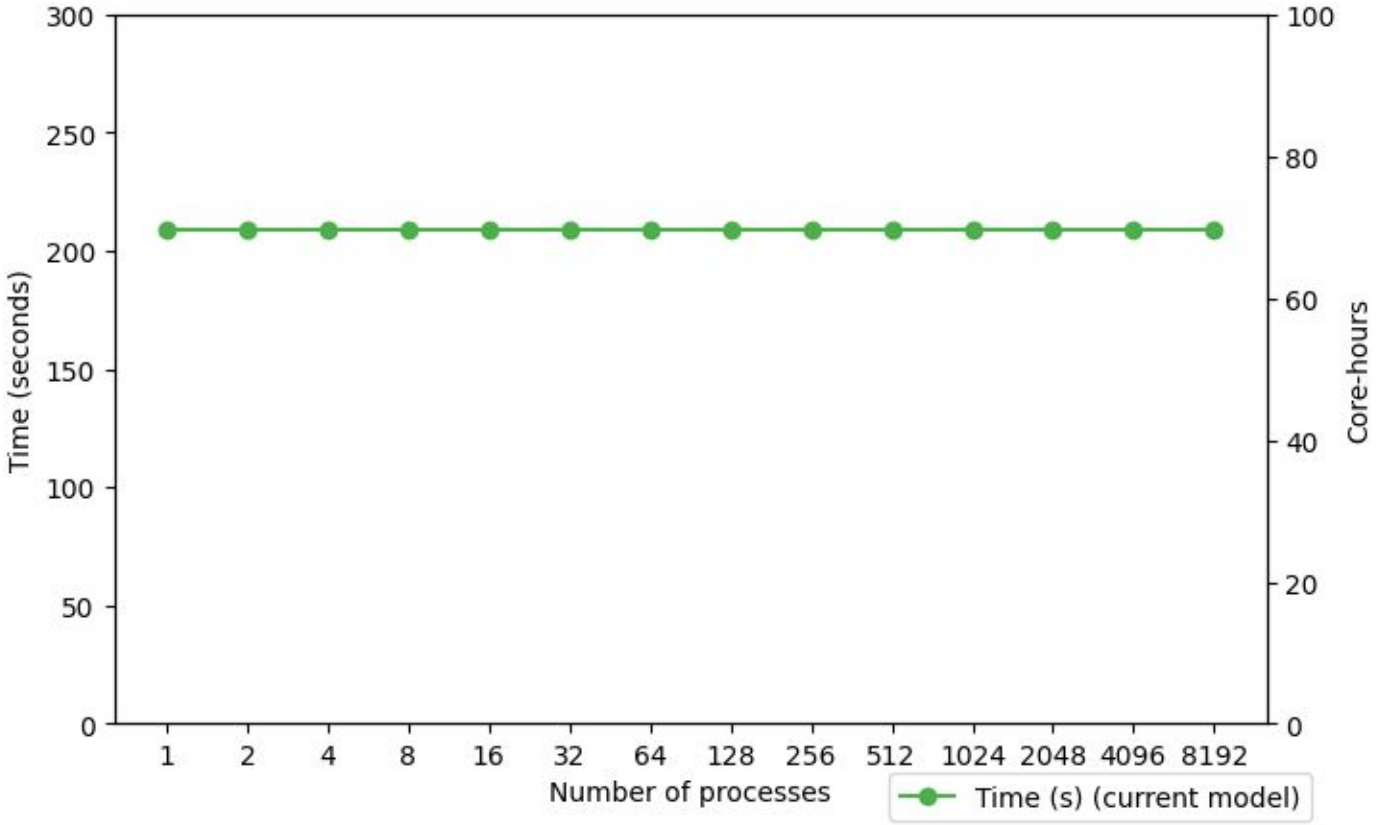

Gather-Sort-Scatter: scaling / performance

N = 64,950,921

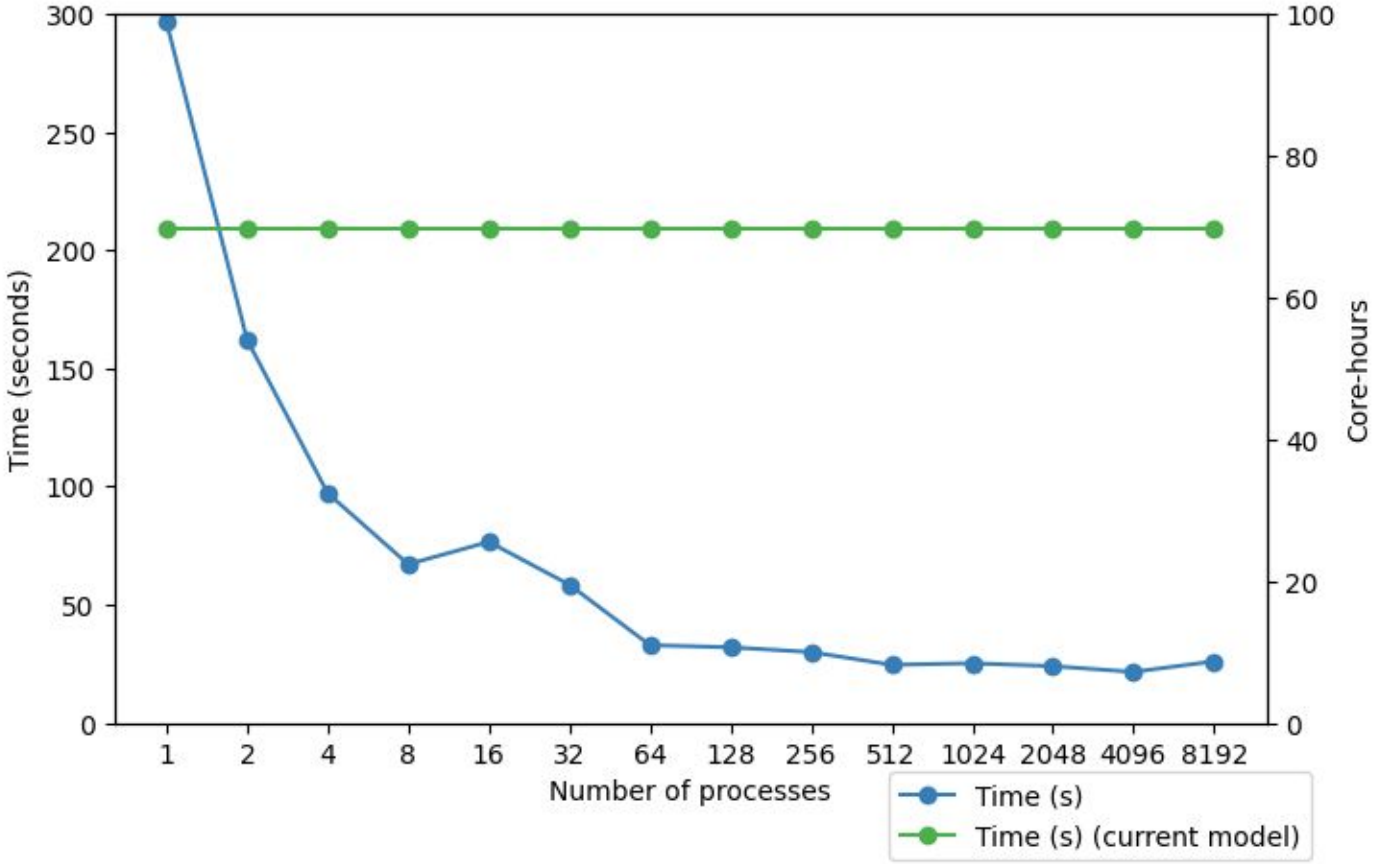

Gather-Sort-Scatter: scaling / performance

N = 64,950,921

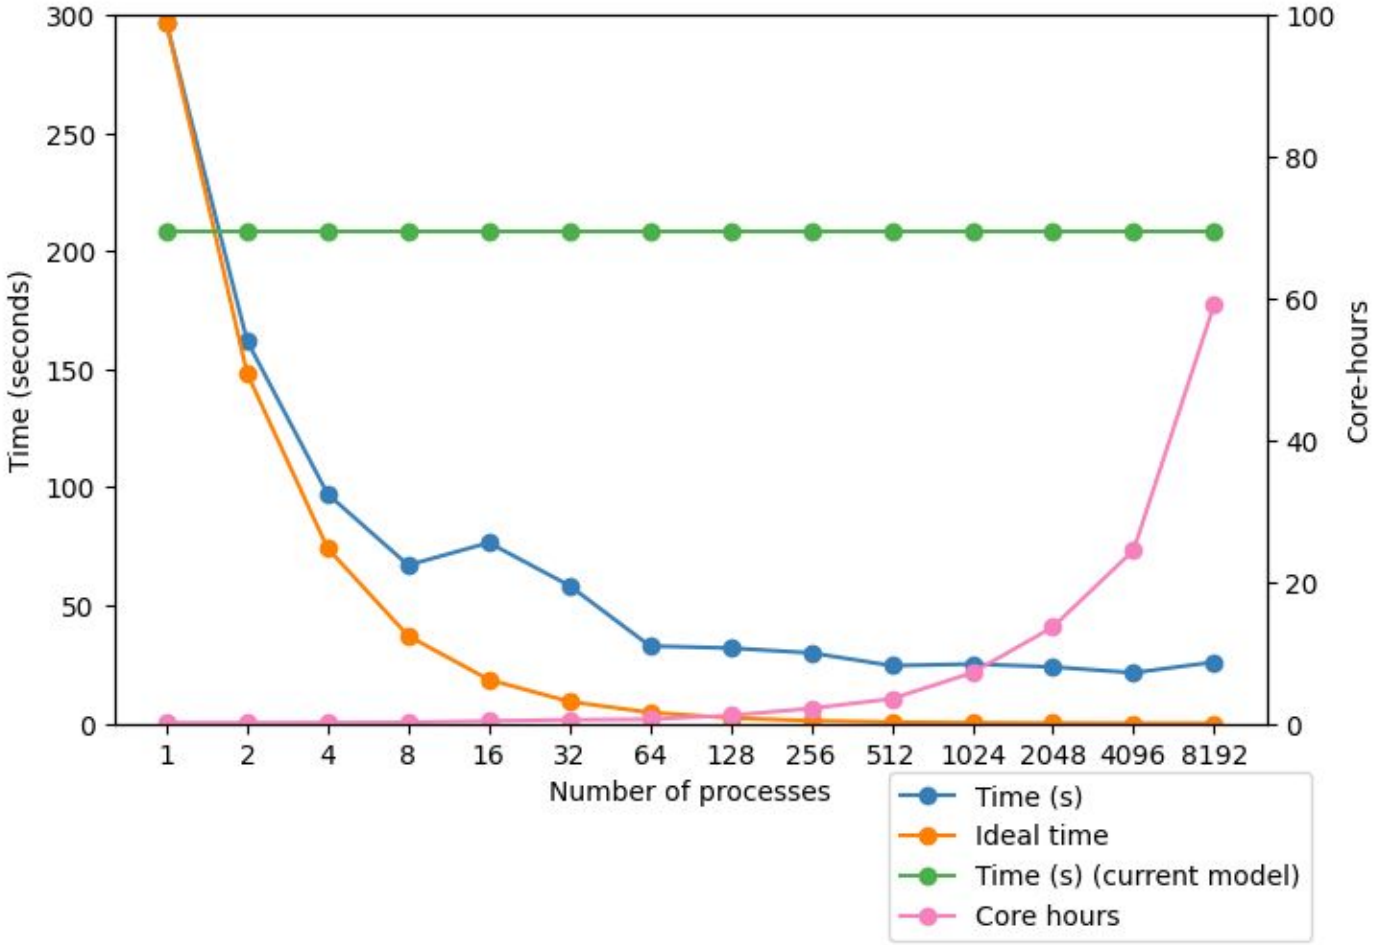

Gather-Sort-Scatter: scaling / performance

N = 64,950,921

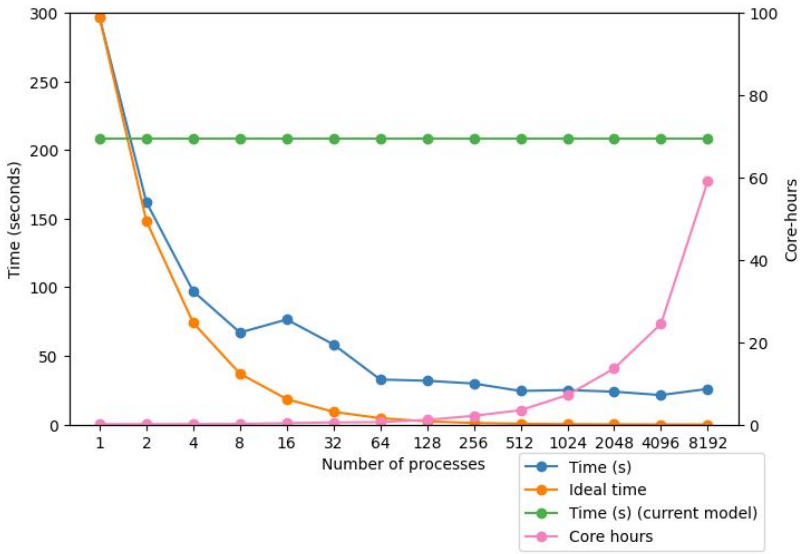

N = 208,901,231

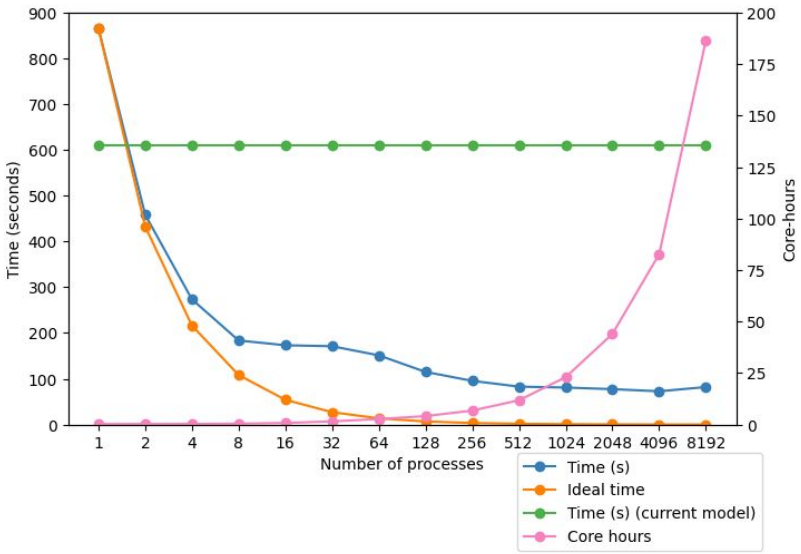

N = 626,703,695

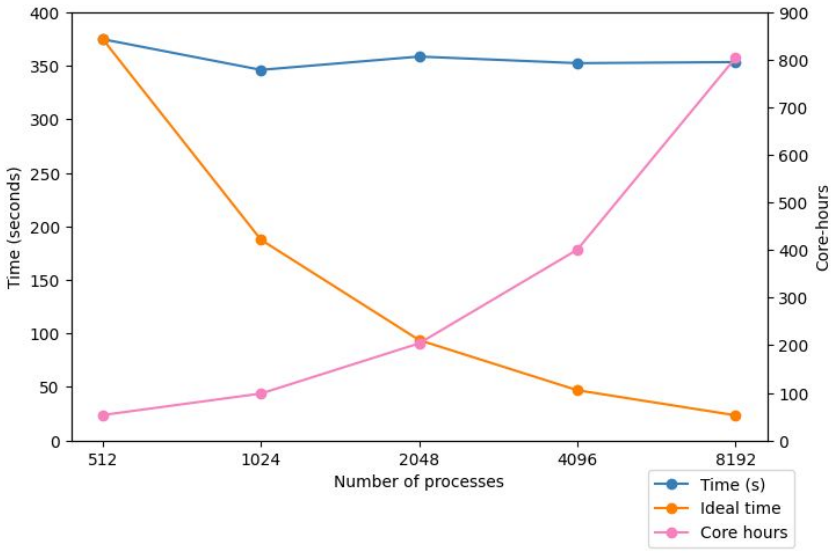

Peak per-process memory

N = 64,950,921	N = 208,901,231	N = 626,703,695
7.46 GB	23.6 GB	70.5 GB

One-sided communication: read times & mem usage

N = 64,950,921

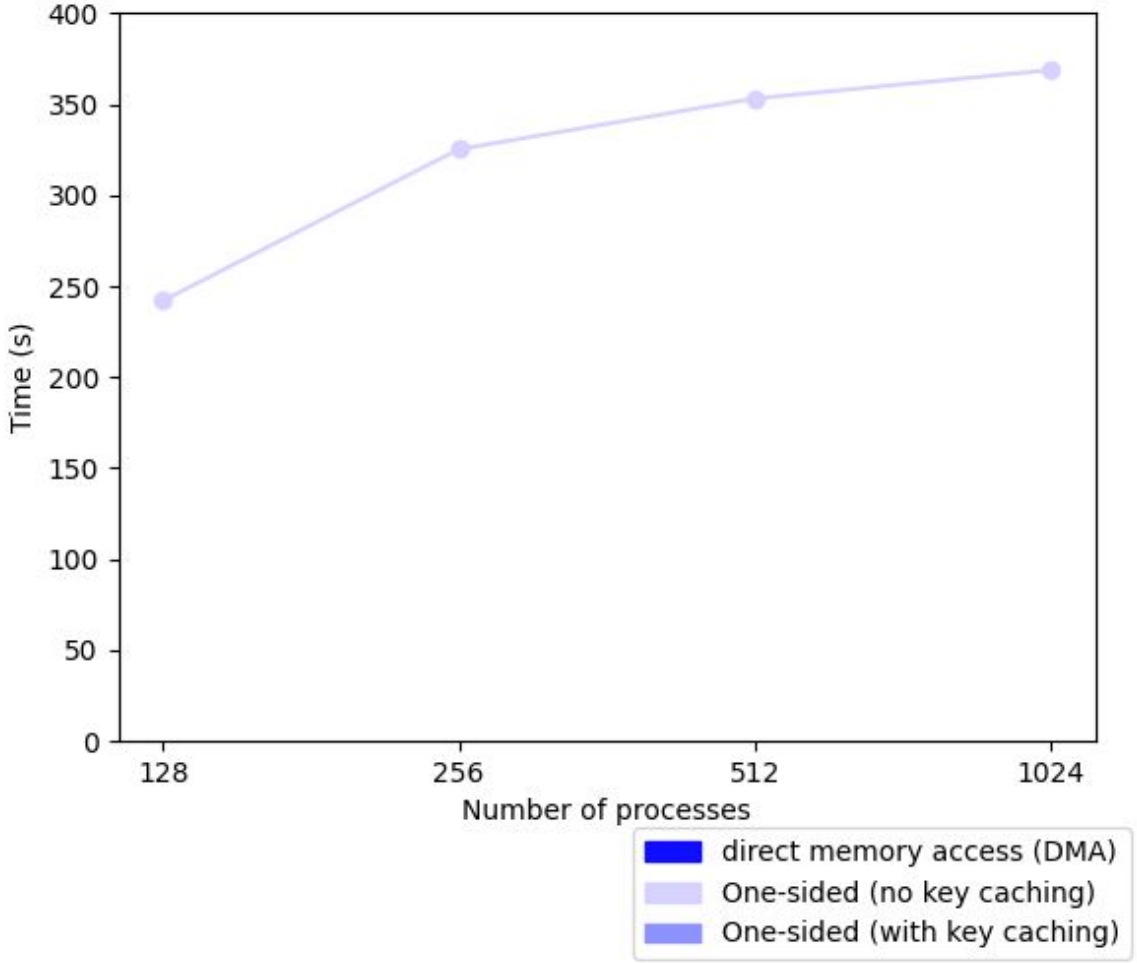

One-sided communication: read times & mem usage

N = 64,950,921

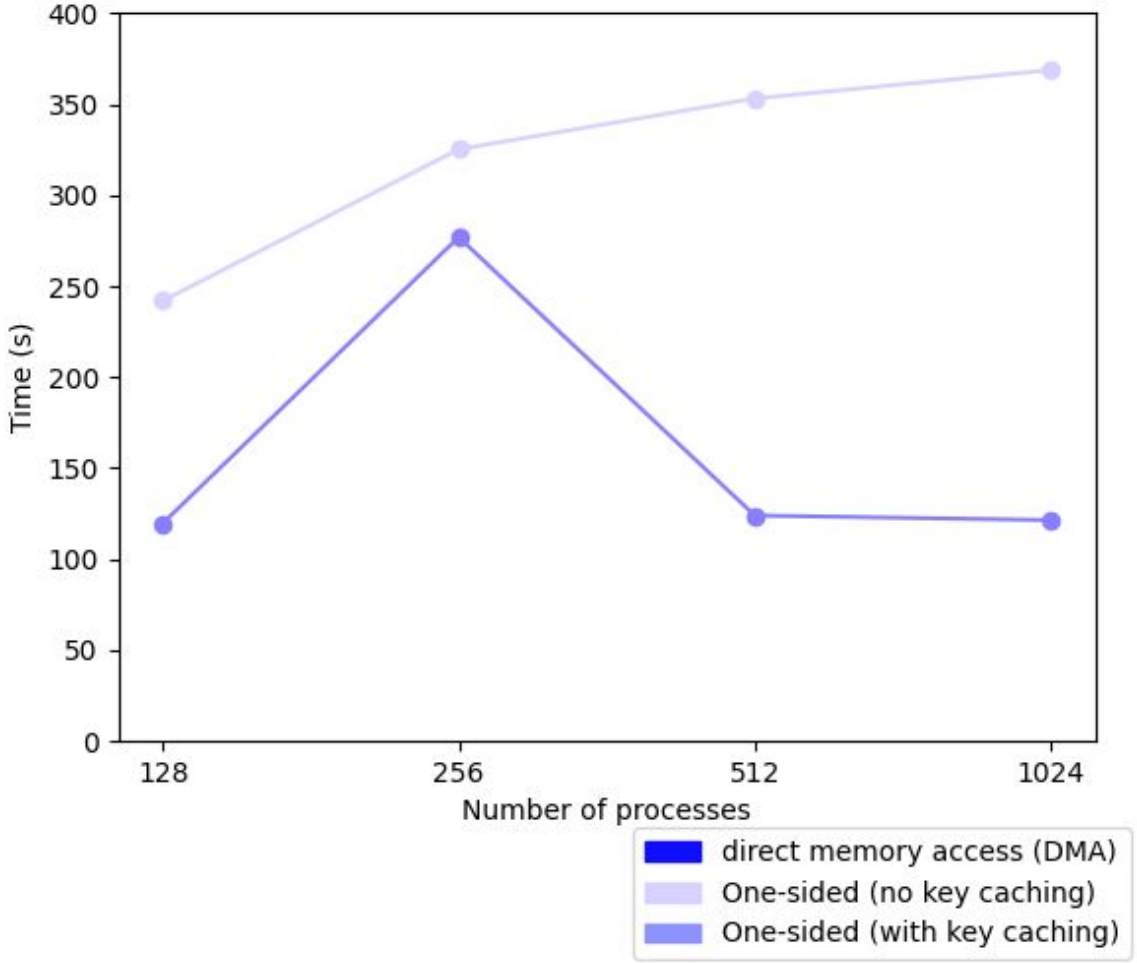

One-sided communication: read times & mem usage

N = 64,950,921

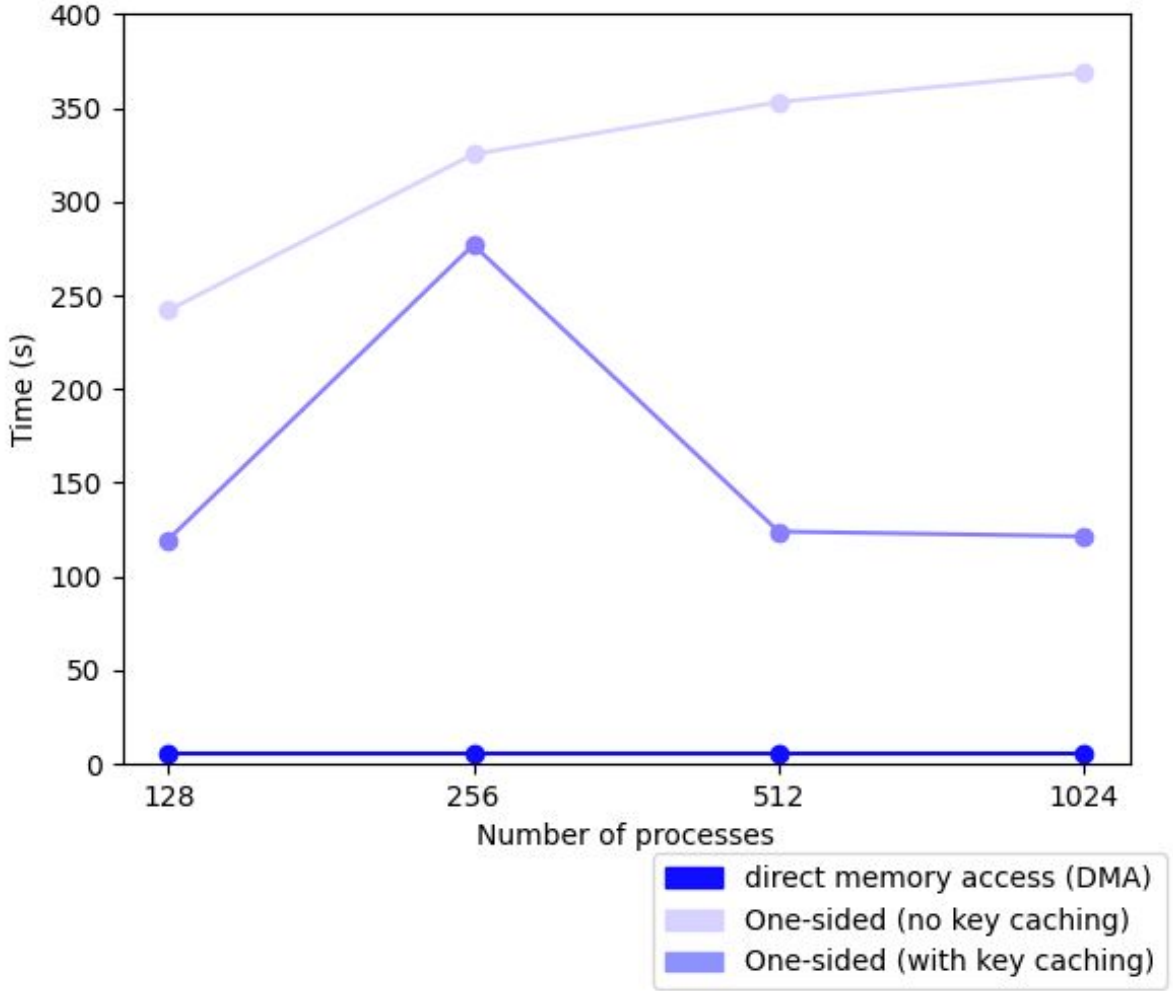

One-sided communication: read times & mem usage

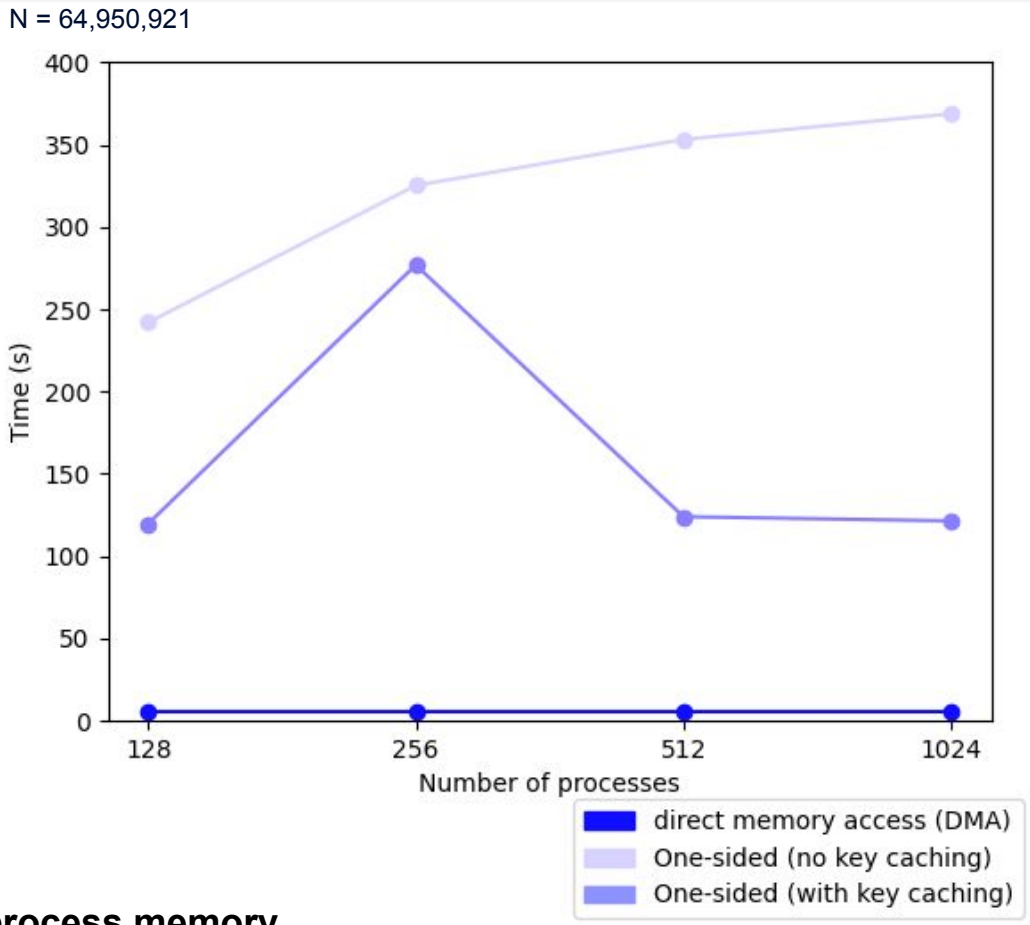

Peak per-process memory

N = 64,950,921	N = 208,901,231	N = 626,703,695
205 MB	279 MB	462 MB

Drawbacks and Improvements

- Both methods have clear drawbacks
  - Memory usage
  - Scalability
  - Time to retrieve

- Ideas for ideal method?
  - One-sided + key-caching?
    - Retrieve key values using collectives?

Discoveries and Enhancements

- Reading subsequence of observations across multiple processes *faster*
- Distributing observations reduced per-process memory used
  - *More* observations could be read **(> 600 mil!)**
- Stepping towards futureproofing DART!

Acknowledgements

- **Helen Kershaw** – profiler and compiler assistance; sound HPC advice; poster & presentation advice; code review
- **Marlee Smith –** profiler assistance, code review, poster & presentation advice
- **Jeff Anderson, Moha Gharamti –** in-depth explanations of DART algorithms
- **Dan Amrhein –** presentation advice
- **Eva Sosoo, Ben Fellman, Virginia Do, Jessica Wang, Jerry Cyccone, all SIParCS interns –** making this summer amazing!

Thank you for attending!