

A novel heuristic for data distribution in massively parallel phylogenetic inference using site repeats

BENOIT MOREL, TOMAS FLOURI AND ALEXANDROS STAMATAKIS SCIENTIFIC COMPUTING GROUP,
HEIDELBERG INSTITUTE FOR THEORETICAL STUDIES

MOTIVATION: PHYLOGENETIC INFERENCE

- Find the tree that maximizes the phylogenetic likelihood for a given multiple sequence alignment (MSA).
- Users typically assign MSA columns/sites to partitions (e.g., individual genes). For each partition/gene a distinct set of evolutionary model parameters are estimated. For analyzing whole-genome datasets, we need to initially split and subsequently distribute these partitions to hundreds of cores for a tightly coupled parallelization of compute- and memory-intensive likelihood calculations on massively parallel systems.
- We split the likelihood evaluation for each tree into three steps:
 - An initialization step, with constant time β for each partition assigned to a specific core.
 - A computation step, with an execution time linear to the total number of sites assigned to a specific core.
 - Execute MPI_Allreduce call to calculate the overall likelihood of the tree for deciding which tree to evaluate next.

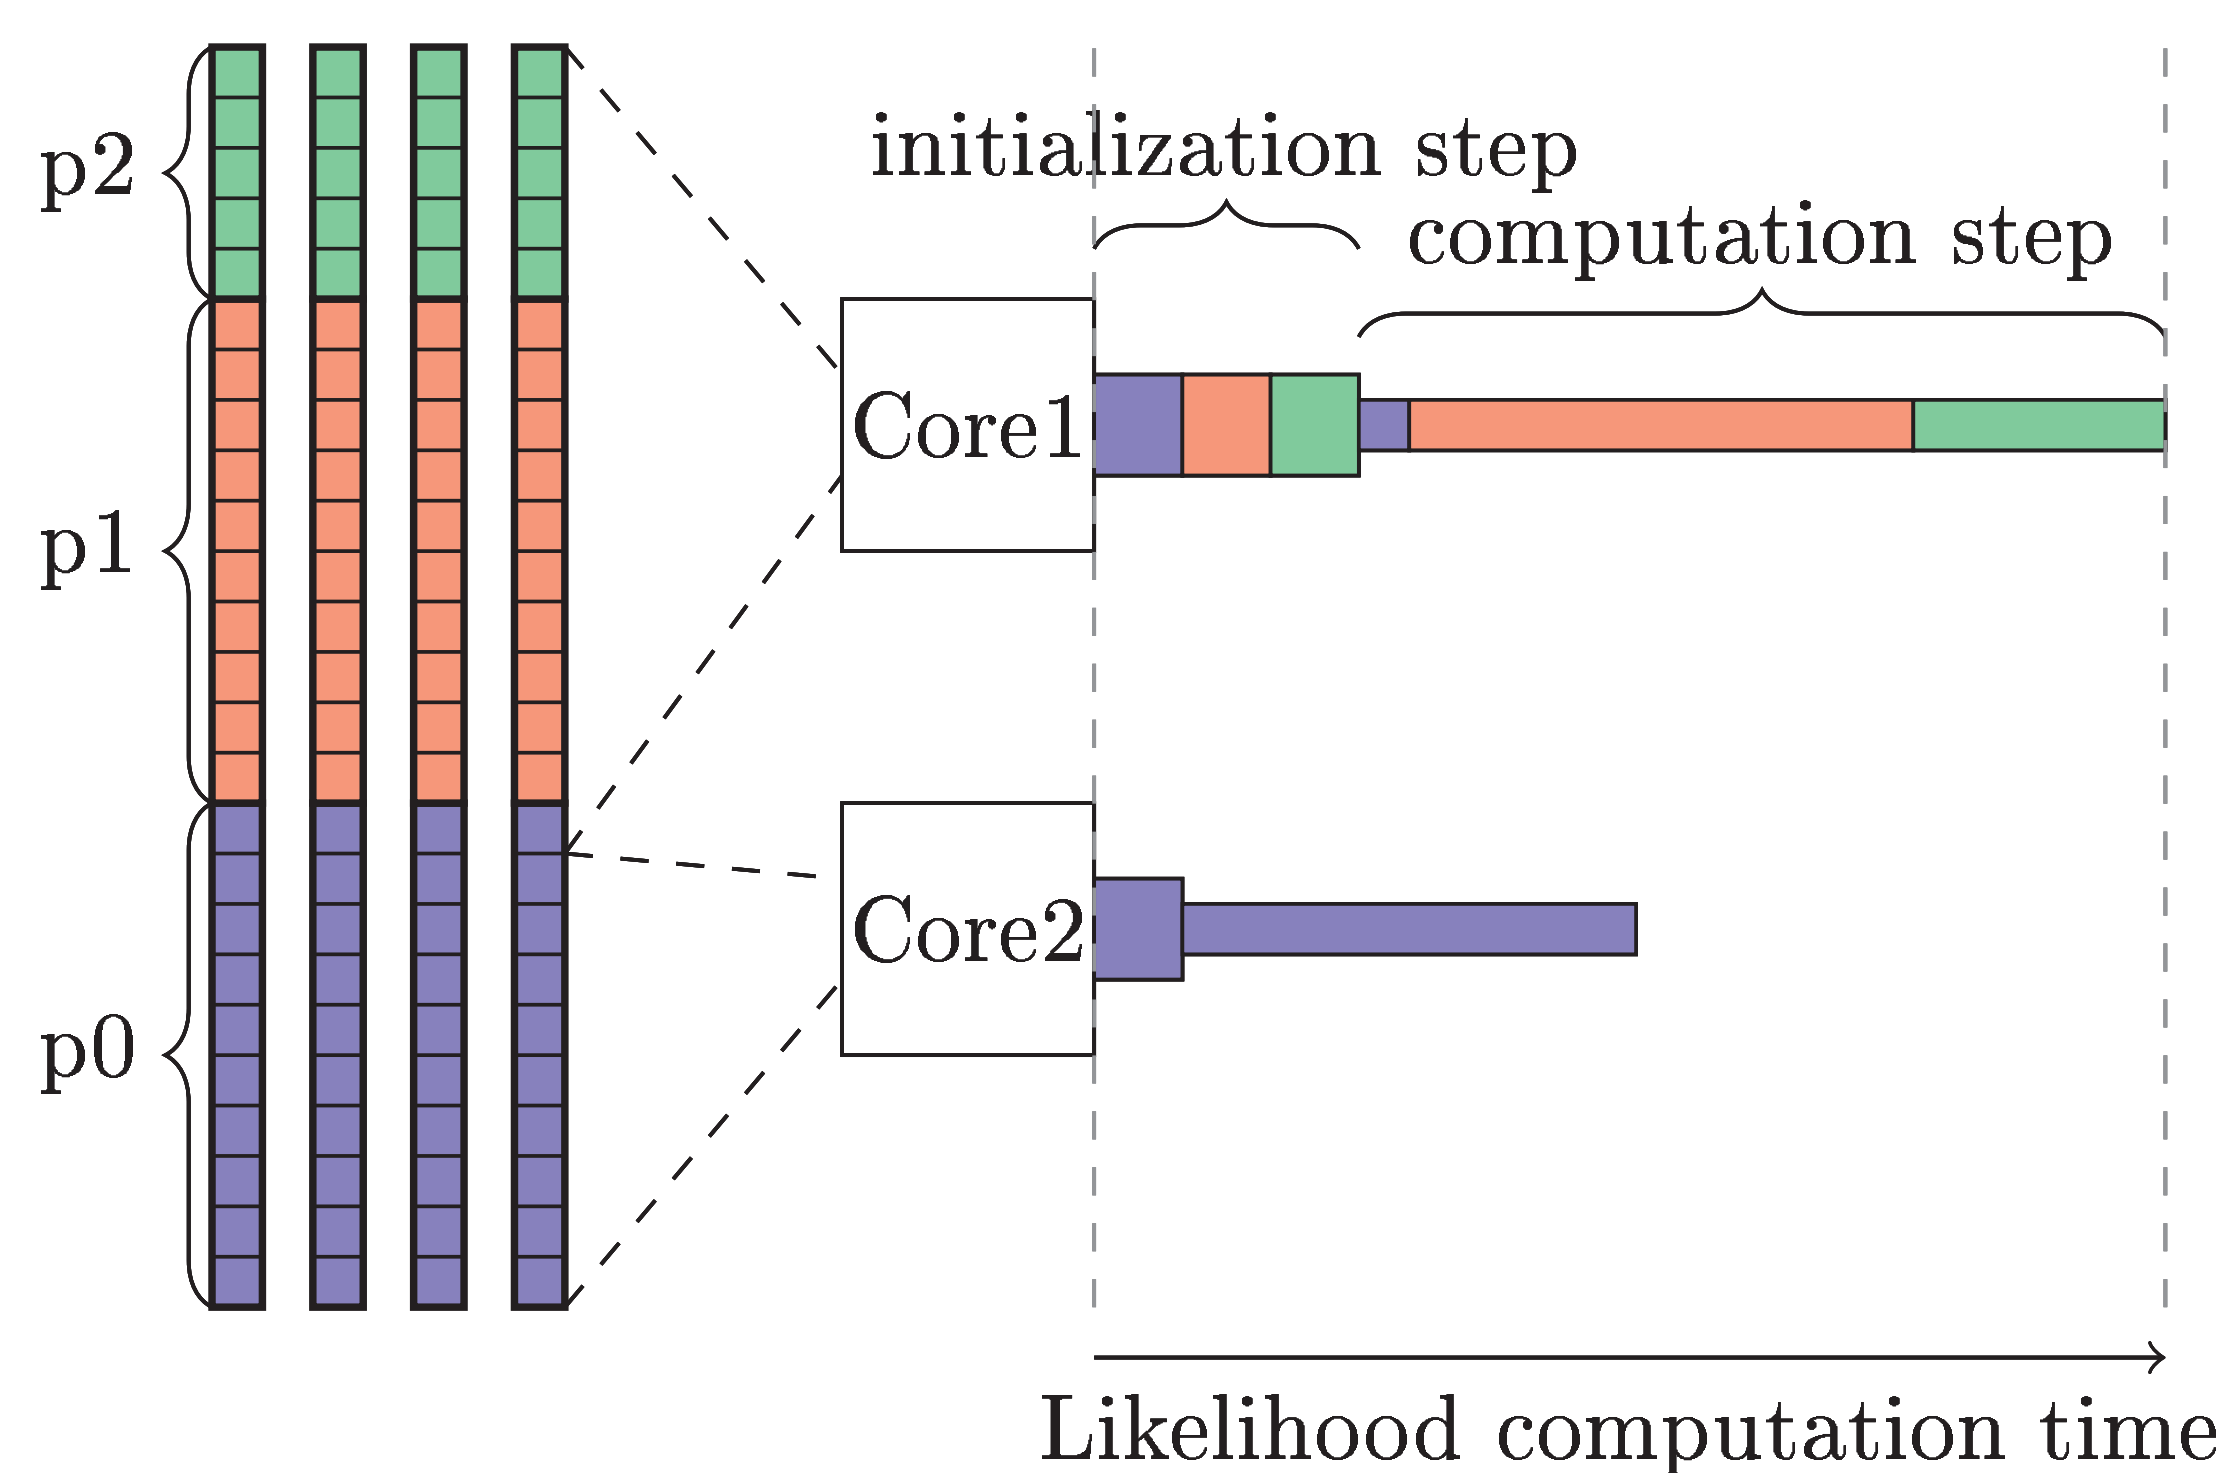


Figure 1: Unbalanced distribution of partitions and timeline for one tree likelihood computation

A NEW CHALLENGE

We recently introduced an algorithmic optimization of the likelihood function [1] which identifies repeating site patterns in subtrees to omit redundant computations. However, it violates the two main assumptions of our previous data distribution algorithm:

- All sites have the same computational cost.
- The computational cost of a site does not depend on the tree.

ABSTRACT PROBLEM

Determine a site-to-core assignment that satisfies:

- The average computational load per core and over all trees

$$\overline{load}_c = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \left(\sum_{s \in S_c} \phi_t(s) \right)$$

is balanced among the cores, where c is a core, \mathcal{T} the set of evaluated trees, S_c the set of sites assigned to c , and $\phi_t(s)$ the predicted computational cost of site s for tree t .

- Minimize the maximum number of distinct partitions assigned to a core.

OUR SOLUTION

We already have an approximation algorithm for data distribution that, given the number of sites per partition, computes a provably almost optimal solution in linearithmic time to the number of partitions [2]. The algorithm assumes that the per-site computational cost is one ($\phi_t(s) = 1$). It initially sorts partitions by size and assigns them to the cores as long as it does not need to split partitions. Thereafter, it splits the remaining partitions and distributes them among cores.

To accommodate sites repeats, we developed the following heuristic:

- Generate a sample of n random trees $\tilde{\mathcal{T}} = \{t_1, t_2, \dots, t_n\}$.
- For each site s calculate its average computational cost over these n random trees:

$$\phi(s) = \frac{1}{|\tilde{\mathcal{T}}|} \sum_{t \in \tilde{\mathcal{T}}} \phi_t(s)$$

- Apply the original algorithm from [2] with $\phi(s)$ as input instead of one.

RESULTS

The heuristic yields an overall speedup of up to 50% on the worst-case, generated datasets, and up to 20% on real datasets. Large-scale empirical data analyses require 10-70 million CPU hours on supercomputers. Hence this improvement is substantial.

The following histograms depict the per-core time for computing the phylogenetic likelihood of one random tree using site repeats with the new and the old solutions.

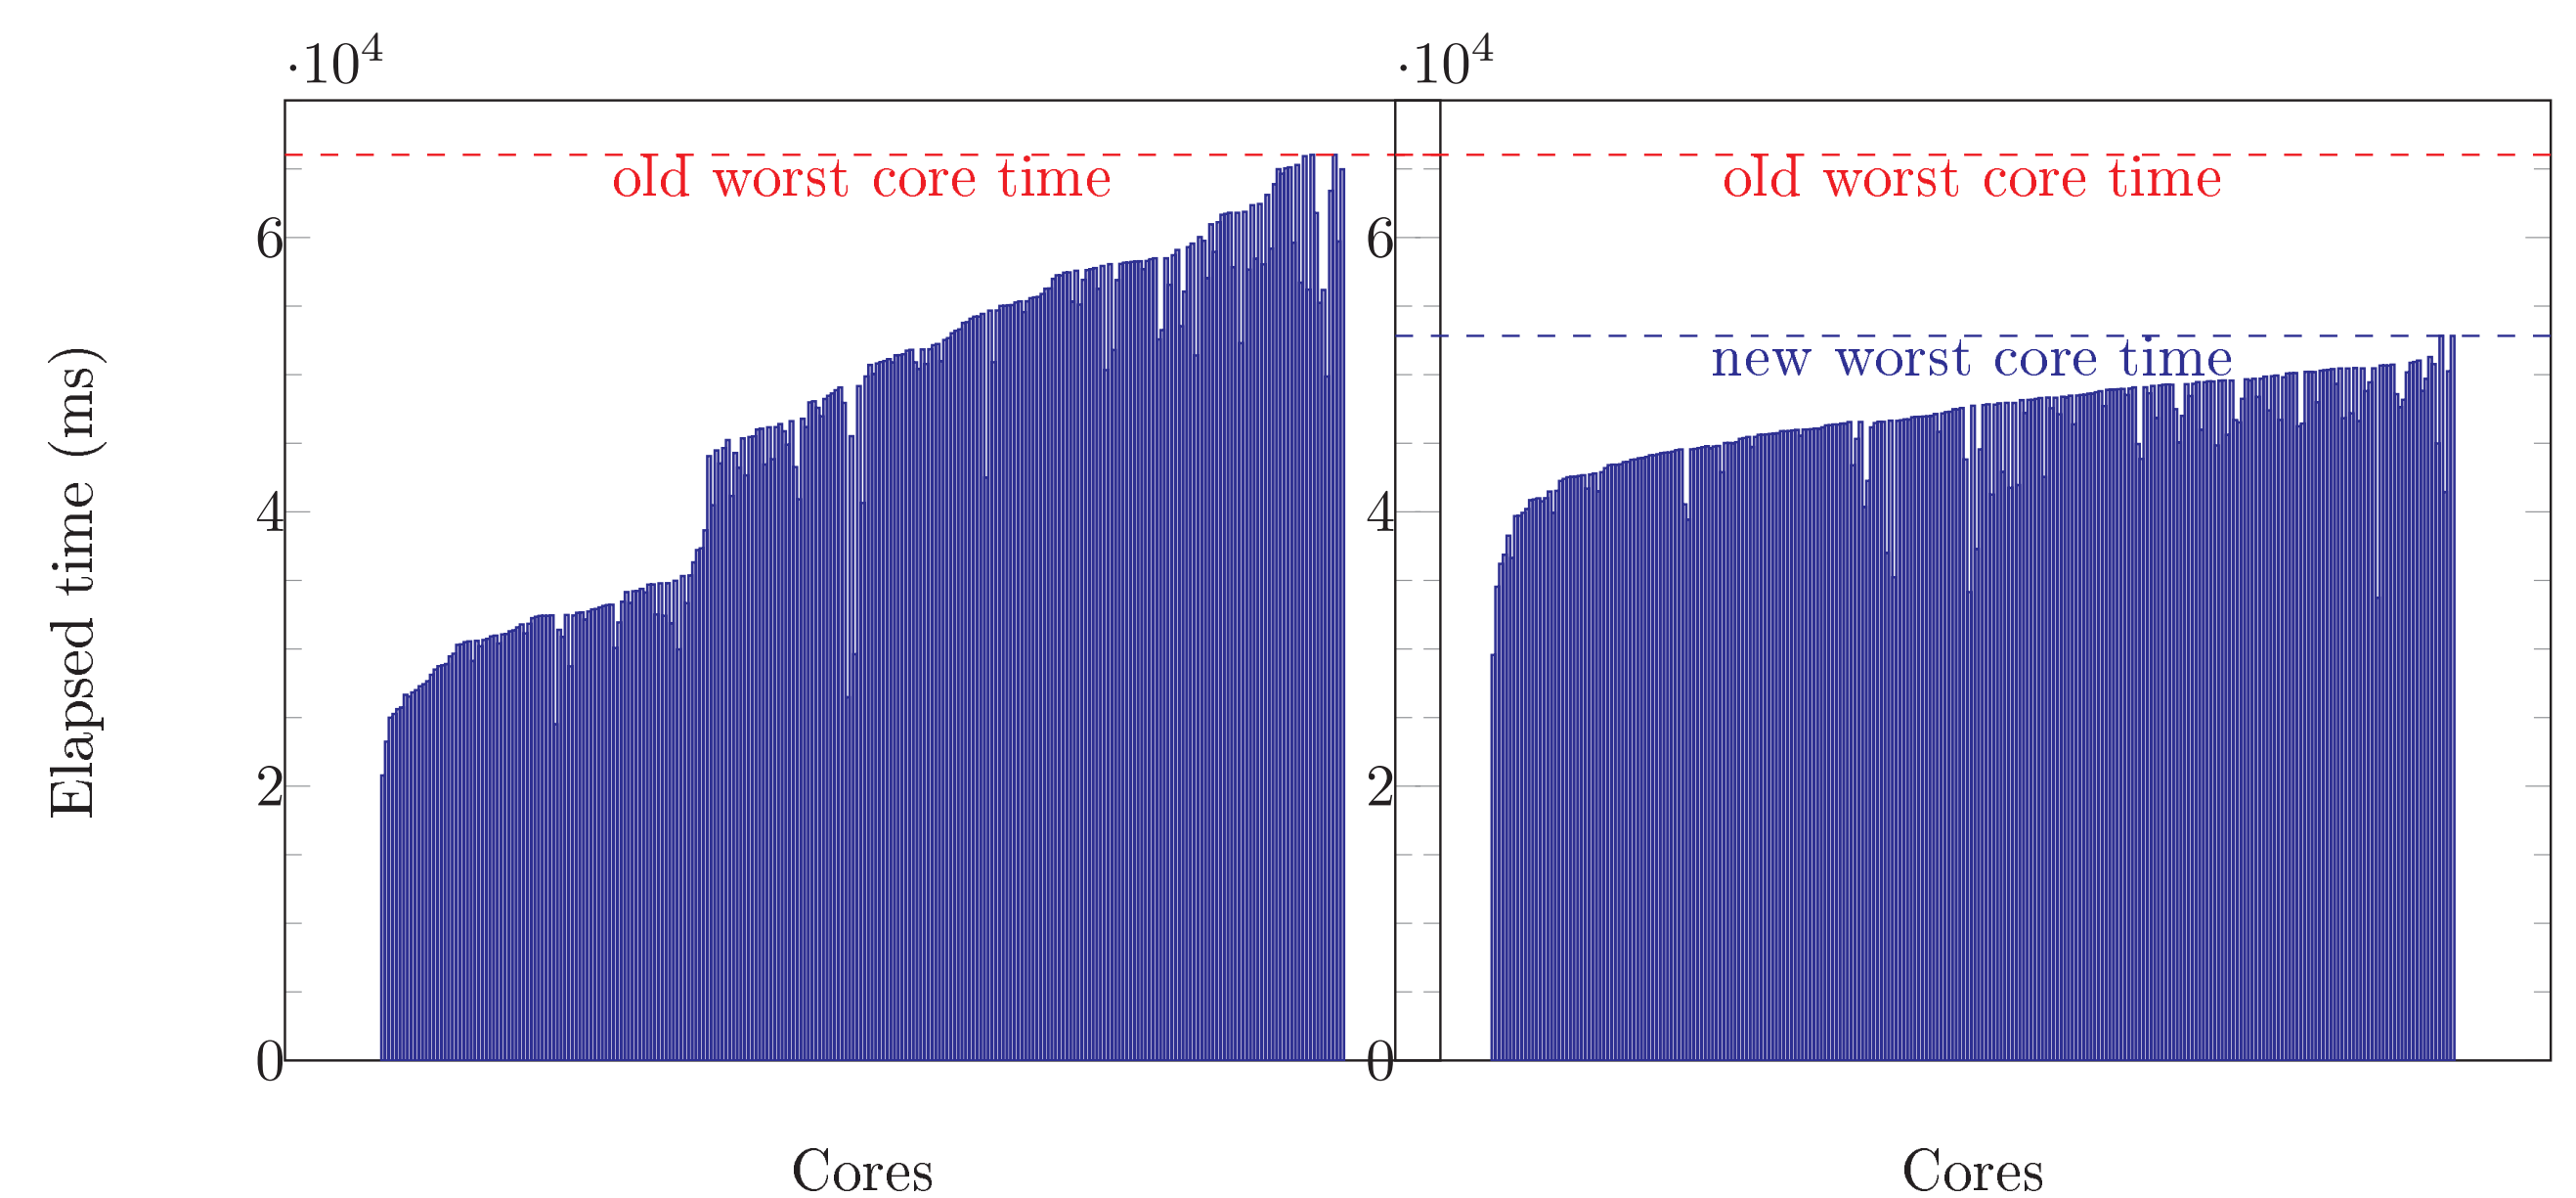


Figure 2: 256 cores, 16 nodes, 144 species, 50 partitions, 370,000 sites (real dataset)

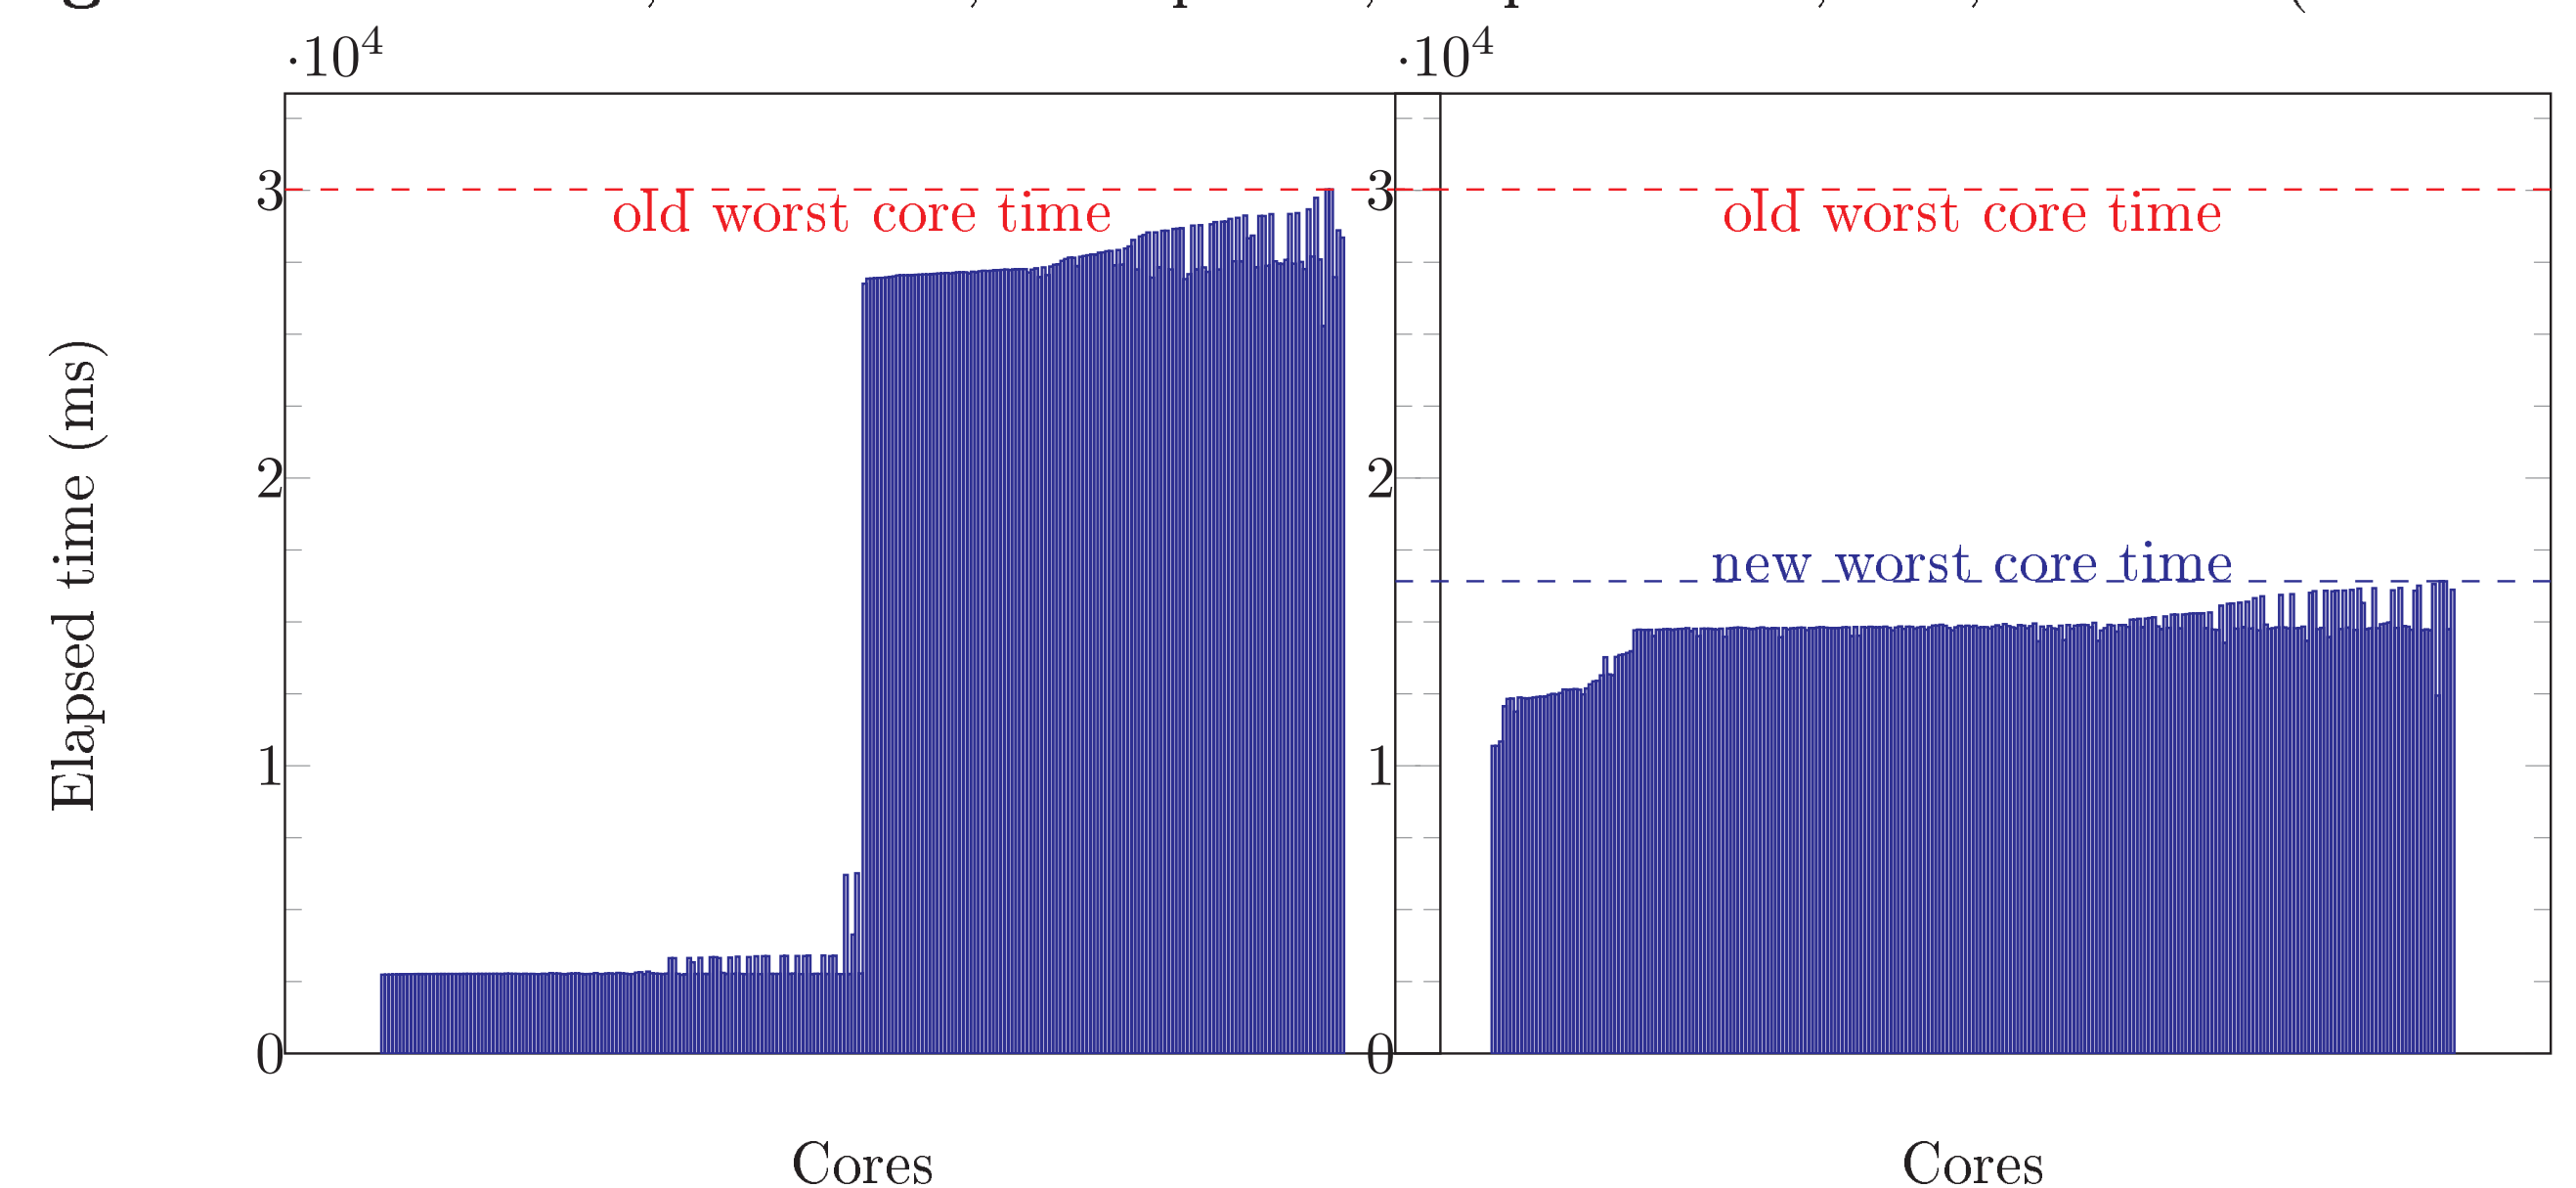


Figure 3: 256 cores, 16 nodes, 50 species, 50 partitions, 500,000 sites (generated dataset)

FUTURE WORK

Find a method to minimize the number of repeats lost during partition splits. A possible approach is to use hypergraphs to model the shared repeats among all sites of a partition and to apply graph partitioning algorithms.

REFERENCES

- [1] K. Kobert et al.: "Efficient detection of repeating sites to accelerate phylogenetic likelihood calculations". In Systematic Biology, 2016.
- [2] T. Flouri et al.: "The divisible load balance problem and its application to phylogenetic inference". In Proceedings of WABI, 2014.