

Optimizing Biomedical Ultrasound Workflow Scheduling Using Cluster Simulation


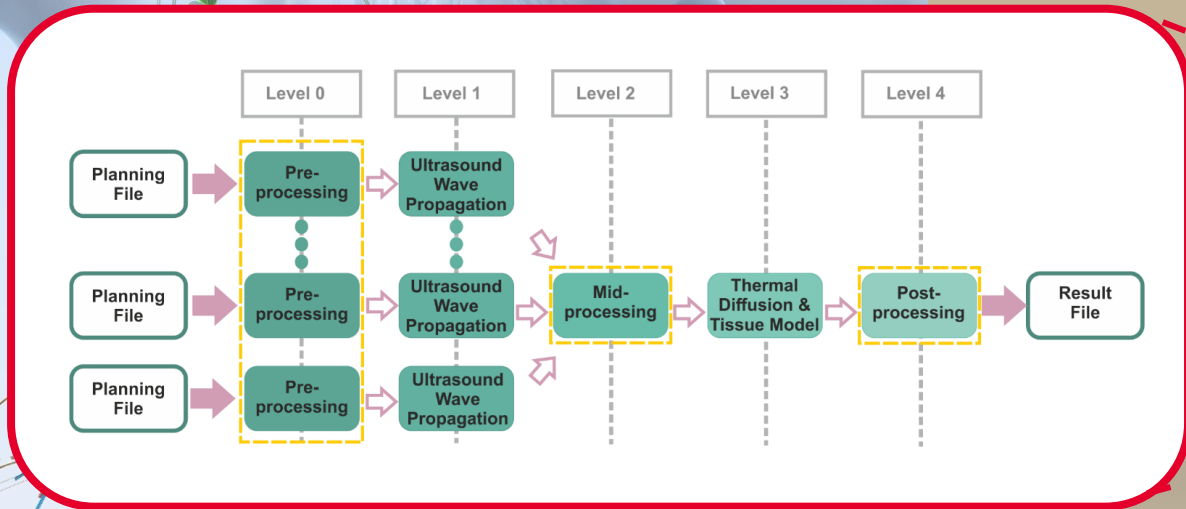

Marta Jaros¹, Dalibor Klusacek², Jiri Jaros¹

¹ Brno University of Technology, Faculty of Information Technology,
Centre of Excellence IT4Innovations, Brno, Czech Republic
{martajaros, jarosjir}@fit.vutbr.cz

² CESNET a.l.e., Brno, Czech Republic
klusacek@cesnet.cz



- Ultrasound treatment planning
- Examples of medical applications:
 - Surgery planning
 - Targeted drug delivery
 - Neurostimulation

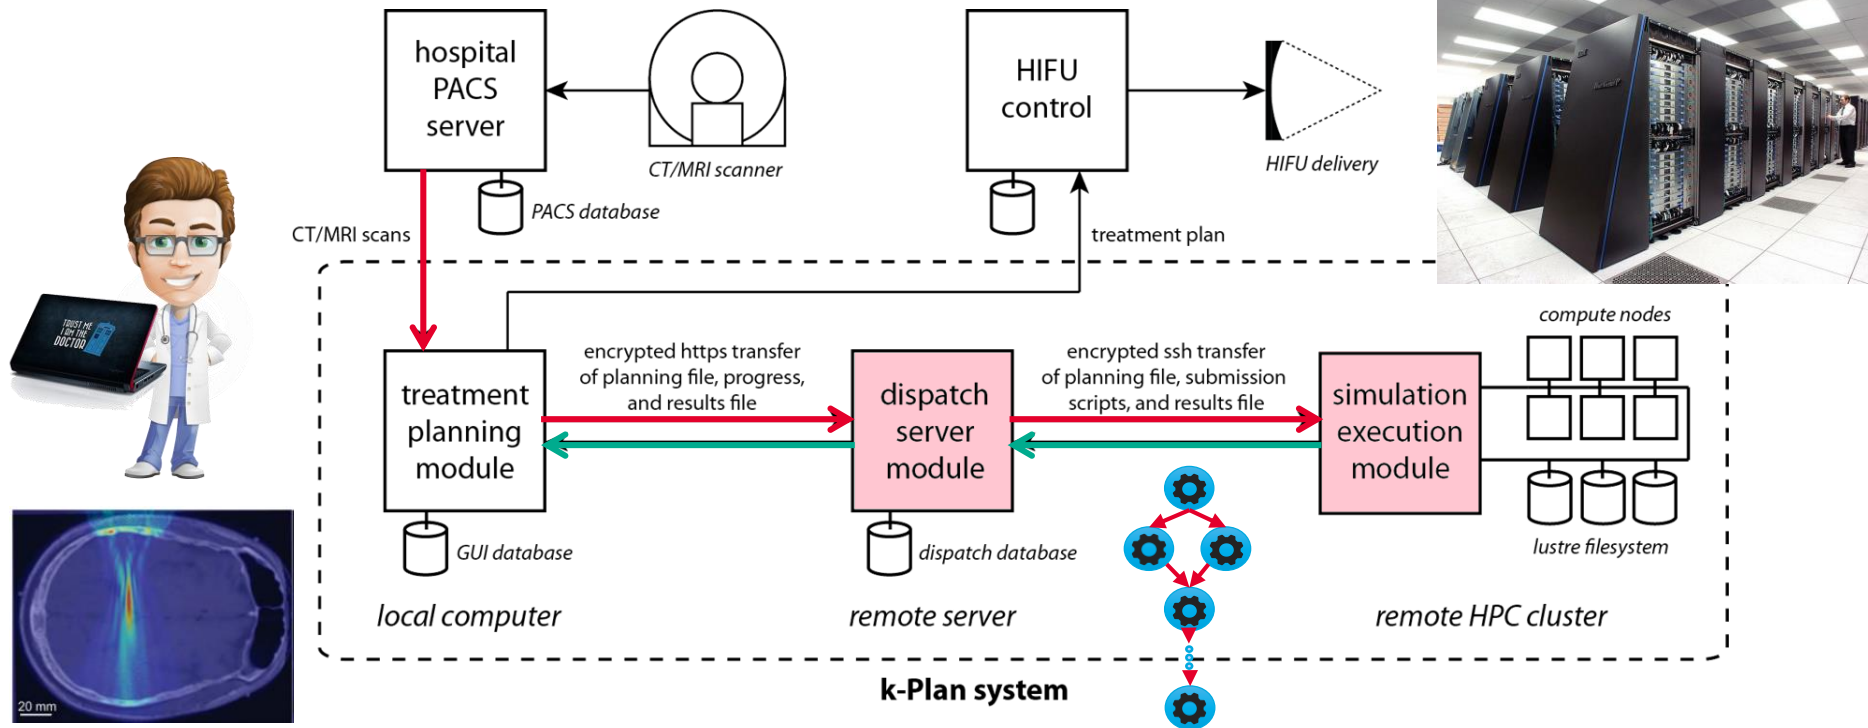


Execution Planning

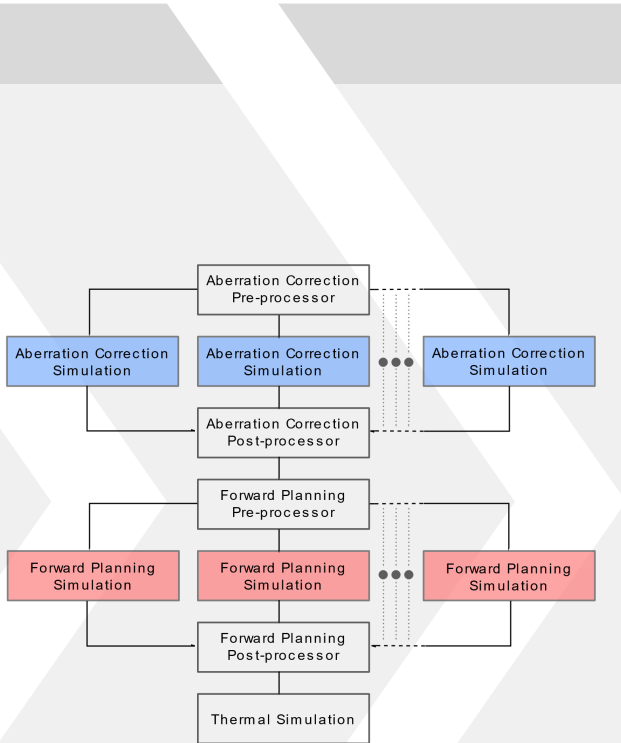
Workflow Execution

Data Exchange Between Tasks




Cluster and Tasks Monitoring



Provide an automated, effective and failure-free workflow execution.

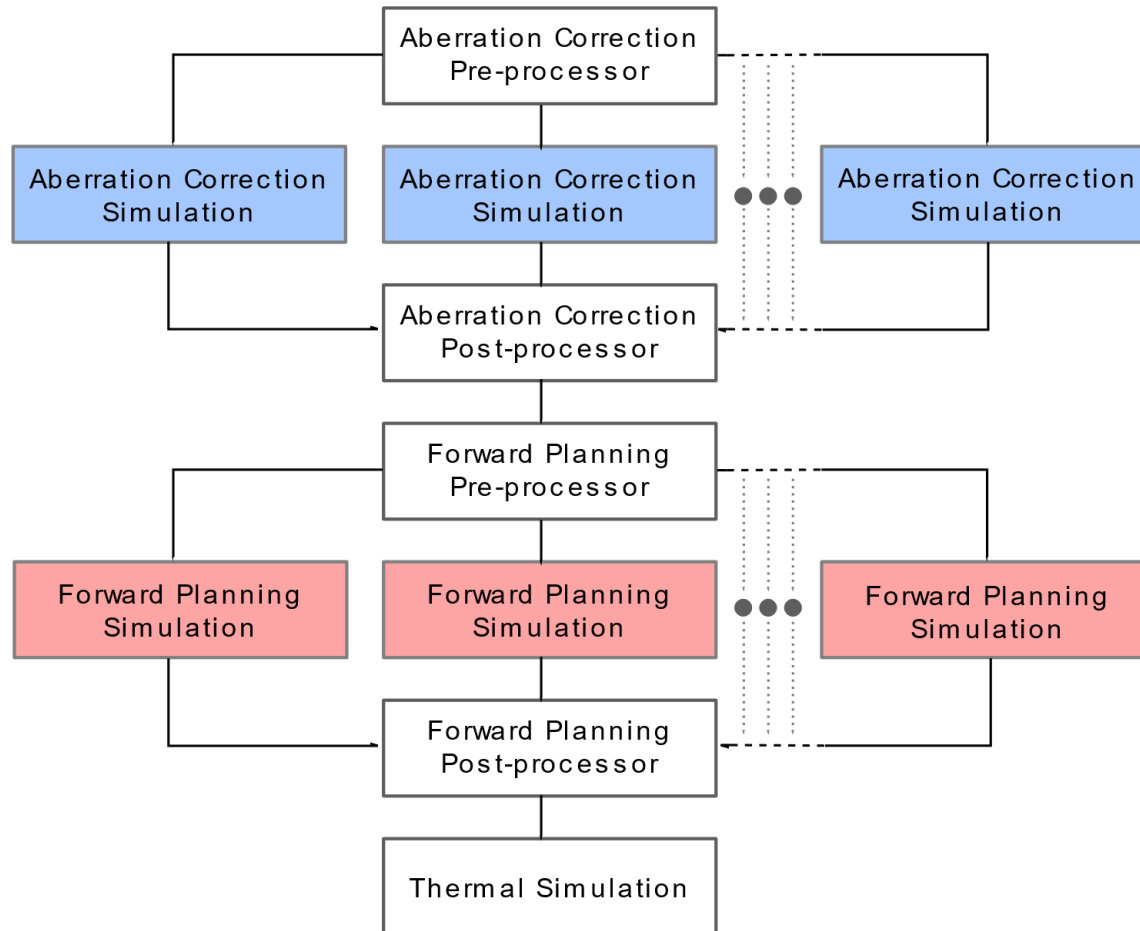


Performance Data



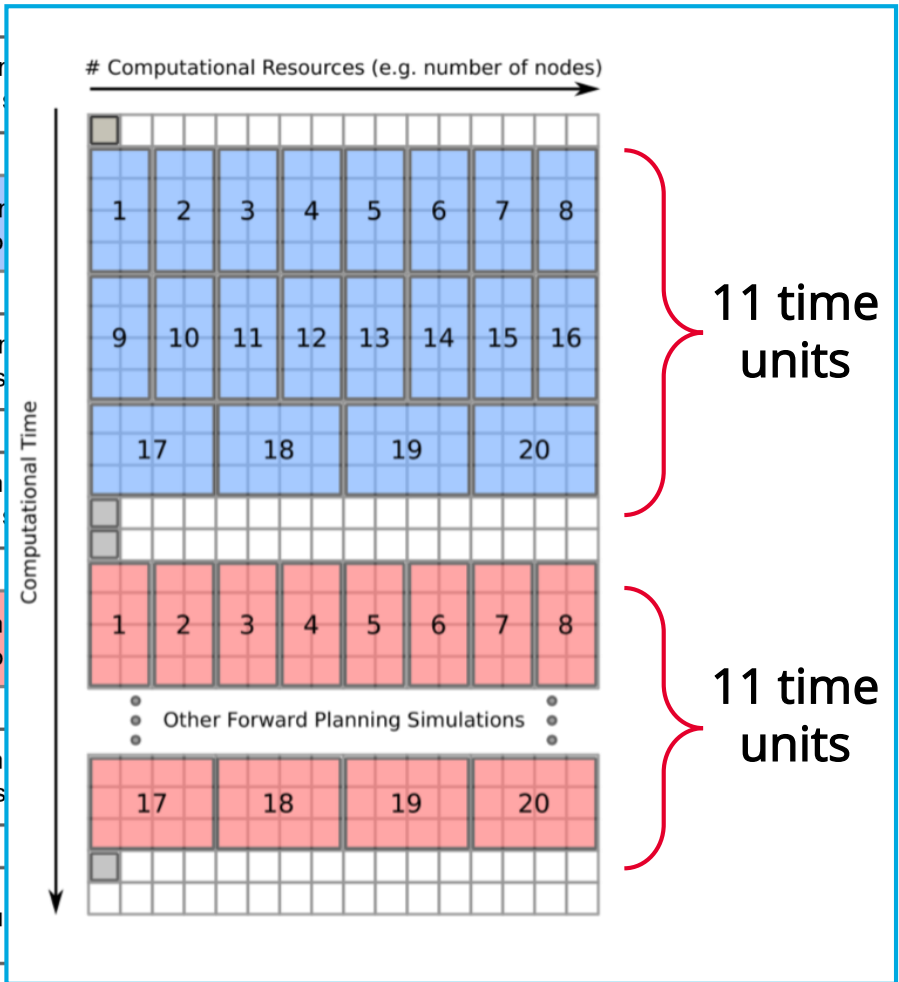
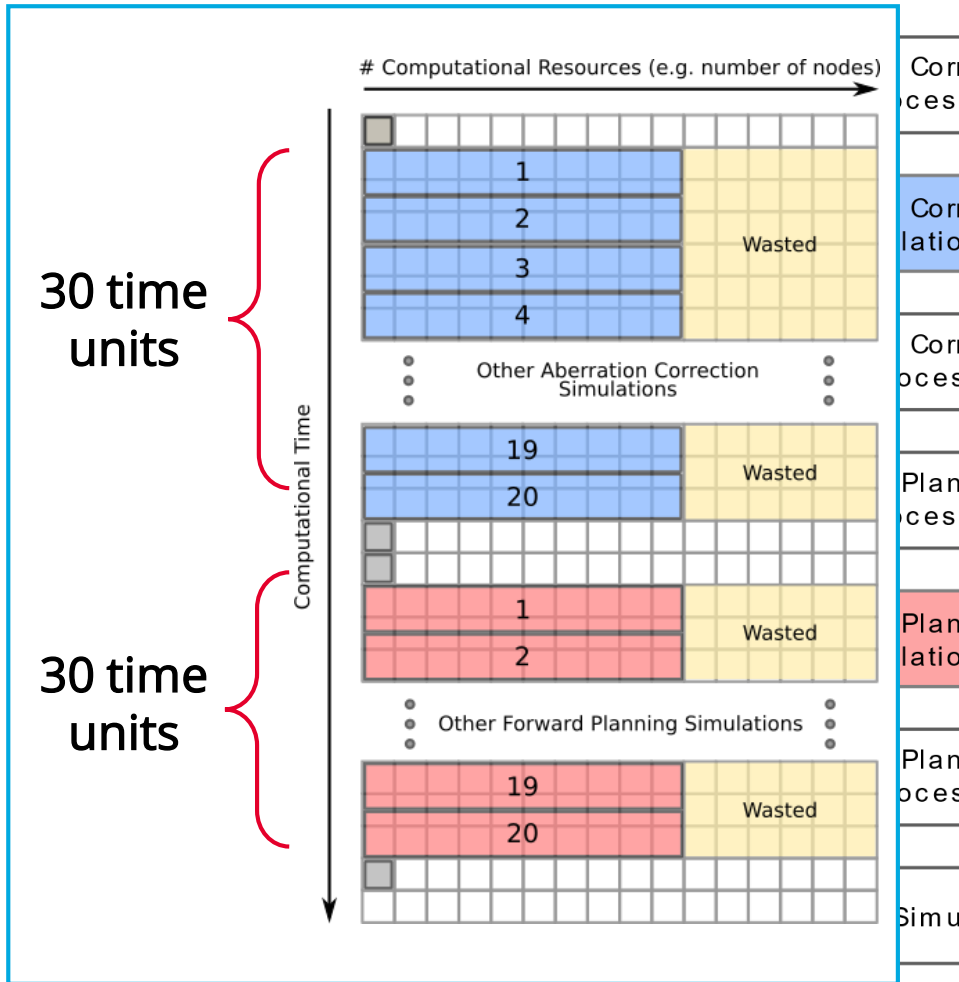
- Select the executable
- Select the execution parameters (HPC, # CPUs, HPC queue, ...)





Example #1

Example #2



Task-level optimization

Workflow-level optimization

Optimizer

- Finds execution parameters based on task inputs
- Increases performance data diversity by small perturbations of the execution parameters
- Can be based on Hill Climbing or Simulated annealing

Interpolator

- Estimates the execution time and cost for a given amount of resources
 - Uses linear and cubic-spline interpolation
- If fails, maximum execution time and associated cost are used

Collector


- Updates performance data in the database after each successful run
- Provides feedback to the optimizer after successful run
- Adapts optimizing process to the cluster workload variations

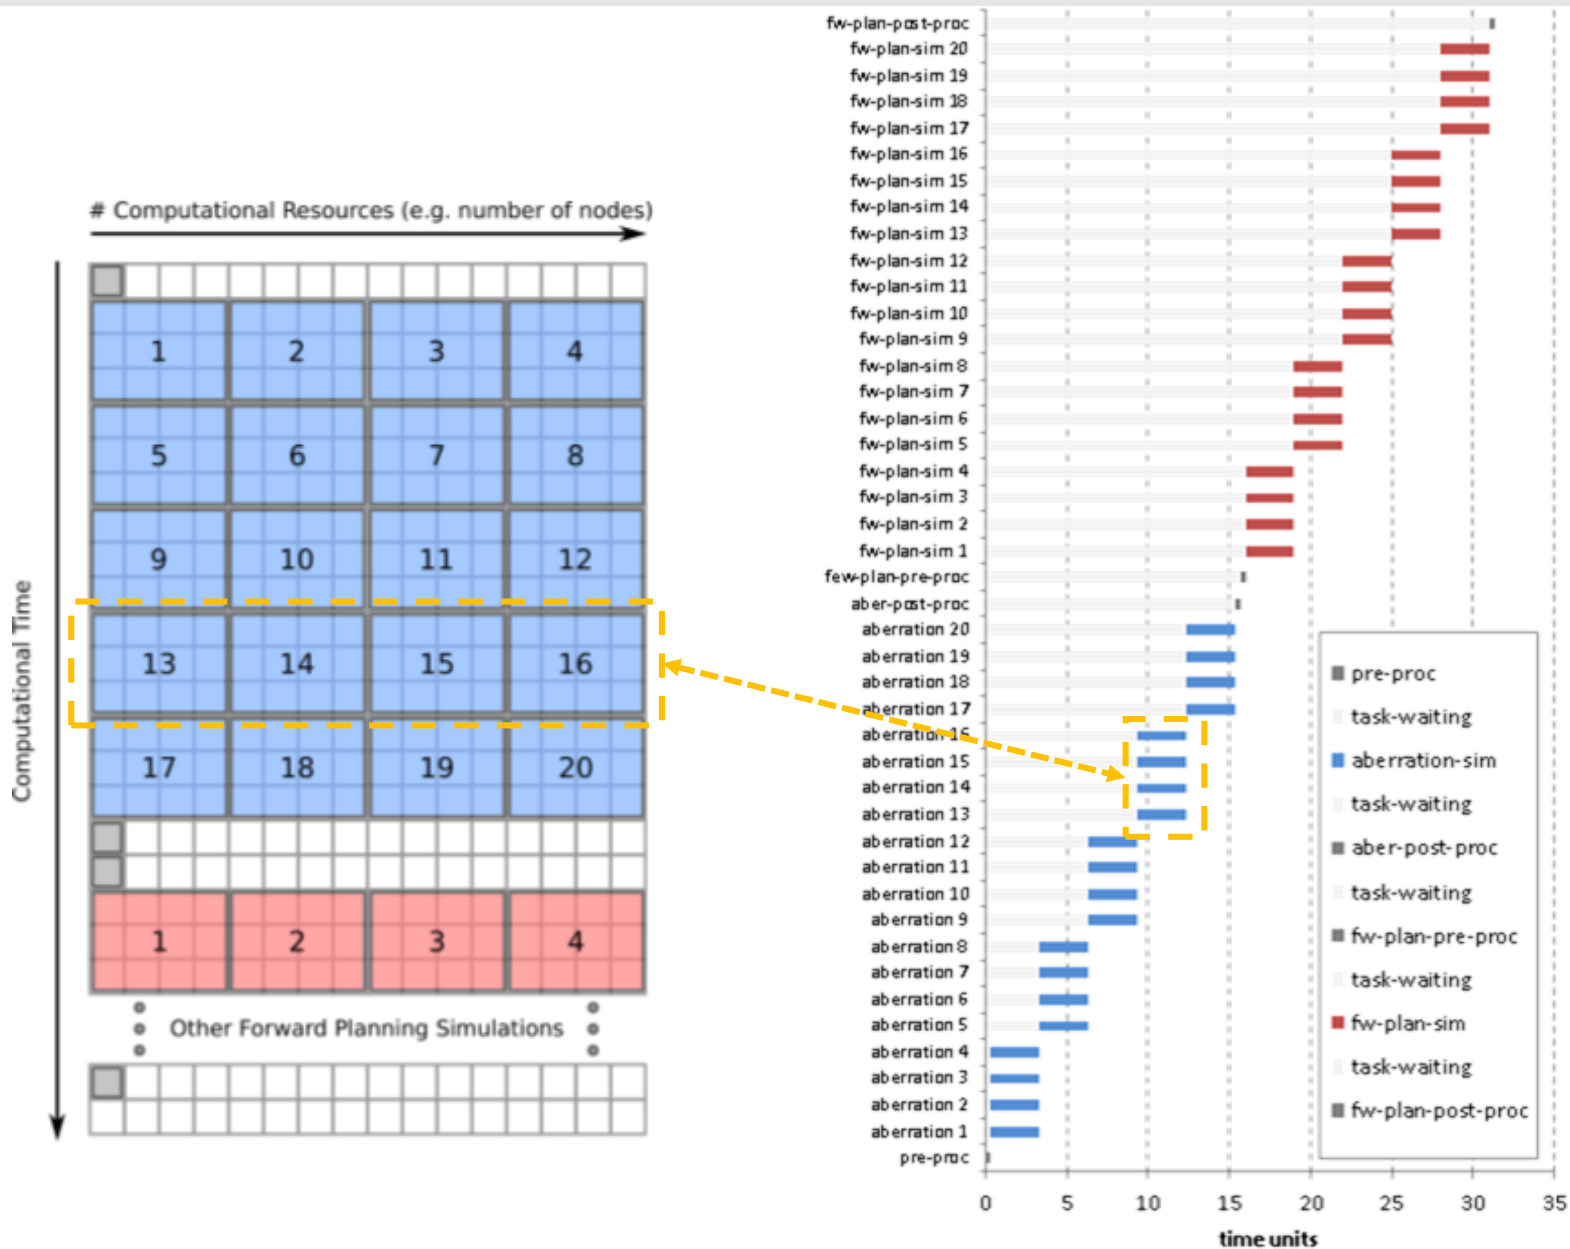
Evaluator

- Calculates the makespan (execution + queueing time)
- Runs the simulator (ALEA)

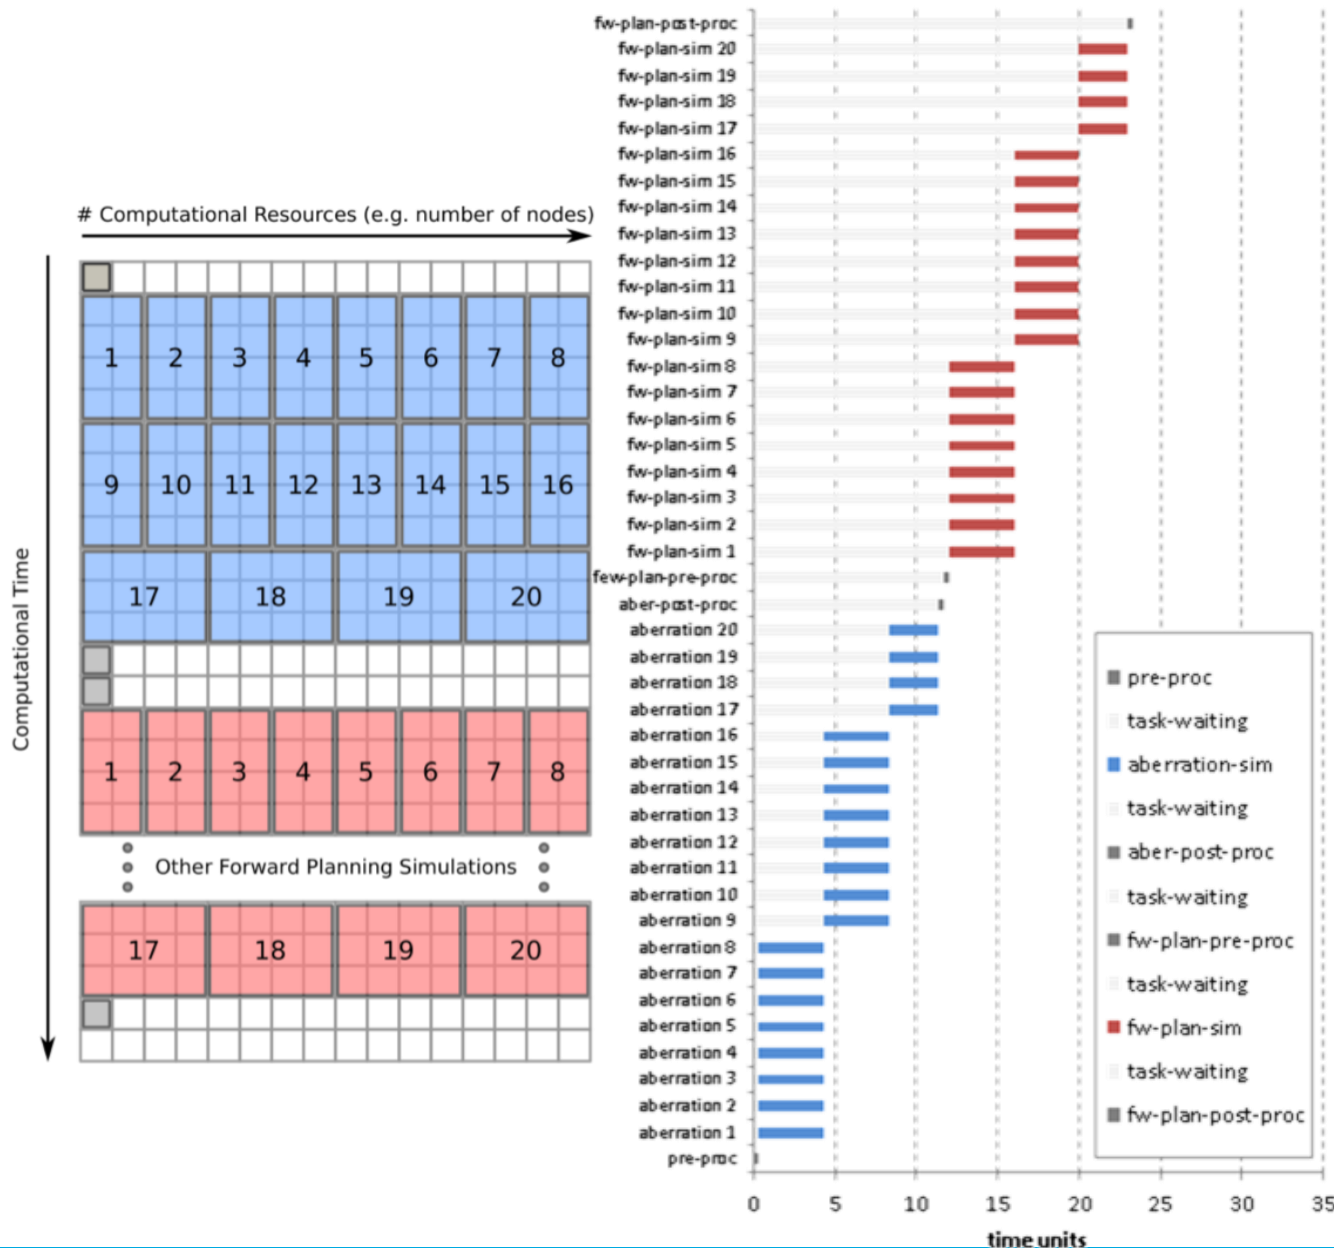


- Calculates the final quality criteria:
$$f = w * (t + q) + (1 - w) * c$$







Experiment #2 + Results



Total execution time (makespan)



Used time slots (nodes x makespan)



Experiment #2 is of 26 % better (cheaper/faster).

- **Optimizer:**
 - Selection of a fast algorithm (1 minute execution)
 - Adaptation on a changing workload
 - Multi-workflow optimization
- **Interpolator:**
 - Data filtering (anomalies, age, missing data)
 - Reliable interpolation in noisy data (oscillations)
- **Evaluator:**
 - Cluster workload snapshots
 - User priorities and backfilling parameters
- **Implementation challenges:**
 - Coupling Java based ALEA with Python based k-Dispatch

- Workflow planning based on performance data has proven the ability to save significant amount of computation time and cost (up to 26 % for a simple benchmark).
- New architecture composed of four modules has been introduced:
 - Optimizer - in a prove of concept stage.
 - Interpolator - partially implemented and being tested on large data sets.
 - Estimator - first attempt to integrate the ALEA simulator.
 - Collector - fully implemented and routinely used for a data collection.
- Future work:
 - Evaluation of different optimizer strategies.
 - Deployment on real systems with real workflows.

Questions?
Thank You for Your Attention