

1 k-Dispatch (Dispatch Server Module)

k-Dispatch is a service providing automated task scheduling, execution and monitoring in the process of ultrasound treatment planning. A treatment simulation comprises a workflow to be computed on a remote HPC cluster. Parsing, scheduling and the execution of this workflow is based on the provided input planning file holding simulation setup and patient data. The workflow is defined as a weighted task graph where the nodes represent individual tasks possibly differing in their nature and computational demands. The design of a good workflow execution schedule is the key to minimize the computational cost and meet the time constraints. Since the HPC environments are highly dynamic, the planning itself is time-critical.

2 Workflow Submission Planning

The task graph exploits concurrency and dependencies in the workflow. Selection of suitable execution parameters for particular tasks of the workflow plays the key role in the scheduling process with the primary aim to reduce the computational cost and the overall processing time from the workflow submission to the results delivery, also referred to as makespan. The selection goes beyond the capabilities of common batch schedulers because the selection is aware of all tasks in the workflow as well as the actual cluster utilization. k-Dispatch implements four performance modules to optimize the selection of execution parameters.

2a Optimizer

- Finds execution parameters based on task inputs.
- Generates high quality candidate solutions (a list of execution parameters for all tasks in the workflow).
- Increases performance data diversity by small perturbations of the execution parameters.
- Implements a Genetic algorithm.

Estimator 2b


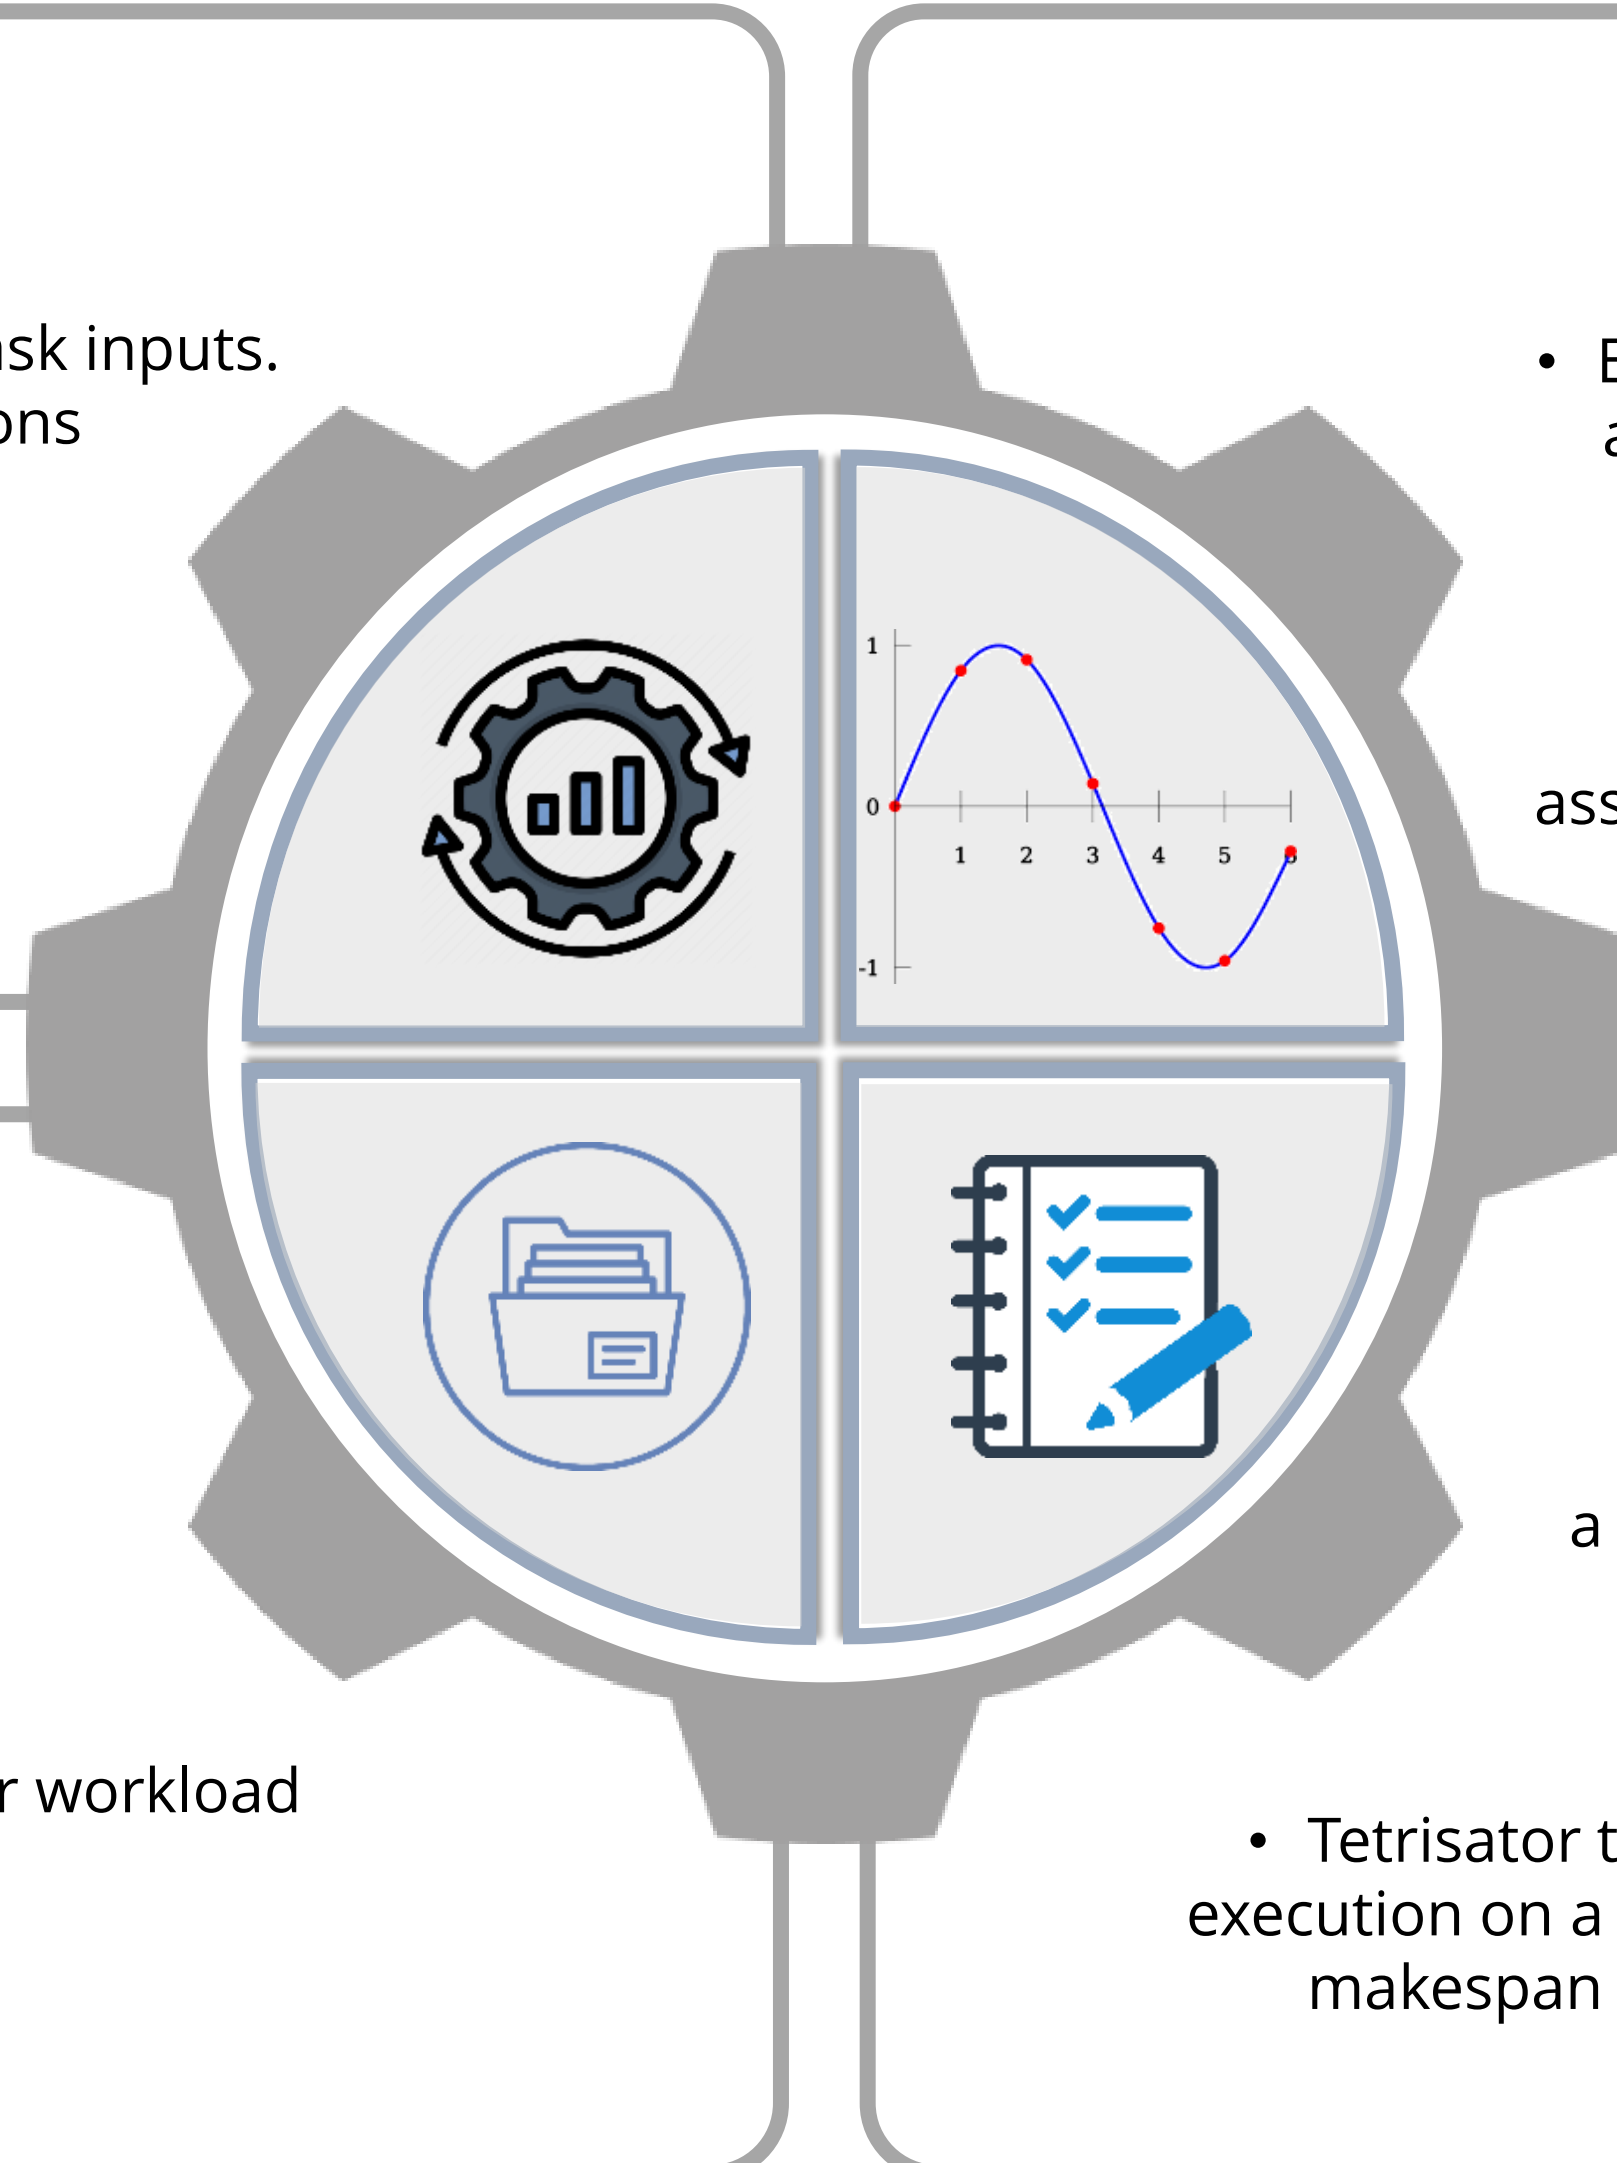
- Estimates the execution time and cost for a given task based on the input data and the amount of required resources.
- Uses linear and quadratic interpolating and extrapolating methods.
- The maximum execution time and associated cost are used if the module fails.

2d Collector

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

Evaluator 2c

- Calculates the makespan, defined as a sum of the execution and queuing time.
- Uses a simplified batch job scheduler simulator called the Tetrinator.
- Tetrinator takes the candidate solution, simulates its execution on a given cluster and calculates the workflow makespan and the cost, i.e., the final quality criteria.



3 Conclusions

k-Dispatch is a service for managing medical applications. Since the execution configuration strongly affects the final tasks mapping, the execution planning is of the highest priority. Currently, k-Dispatch enables users to easily execute predefined workflows on various HPC systems by only providing medical input data.

4 Current and Future Work

- Next steps in the development are to
- (1) include the information about the actual cluster utilization into the cluster simulator,
 - (2) implement a backfilling technique into the cluster simulator,
 - (3) examine more advanced machine learning techniques to improve the interpolation accuracy once the performance database includes tens of thousands of records.