

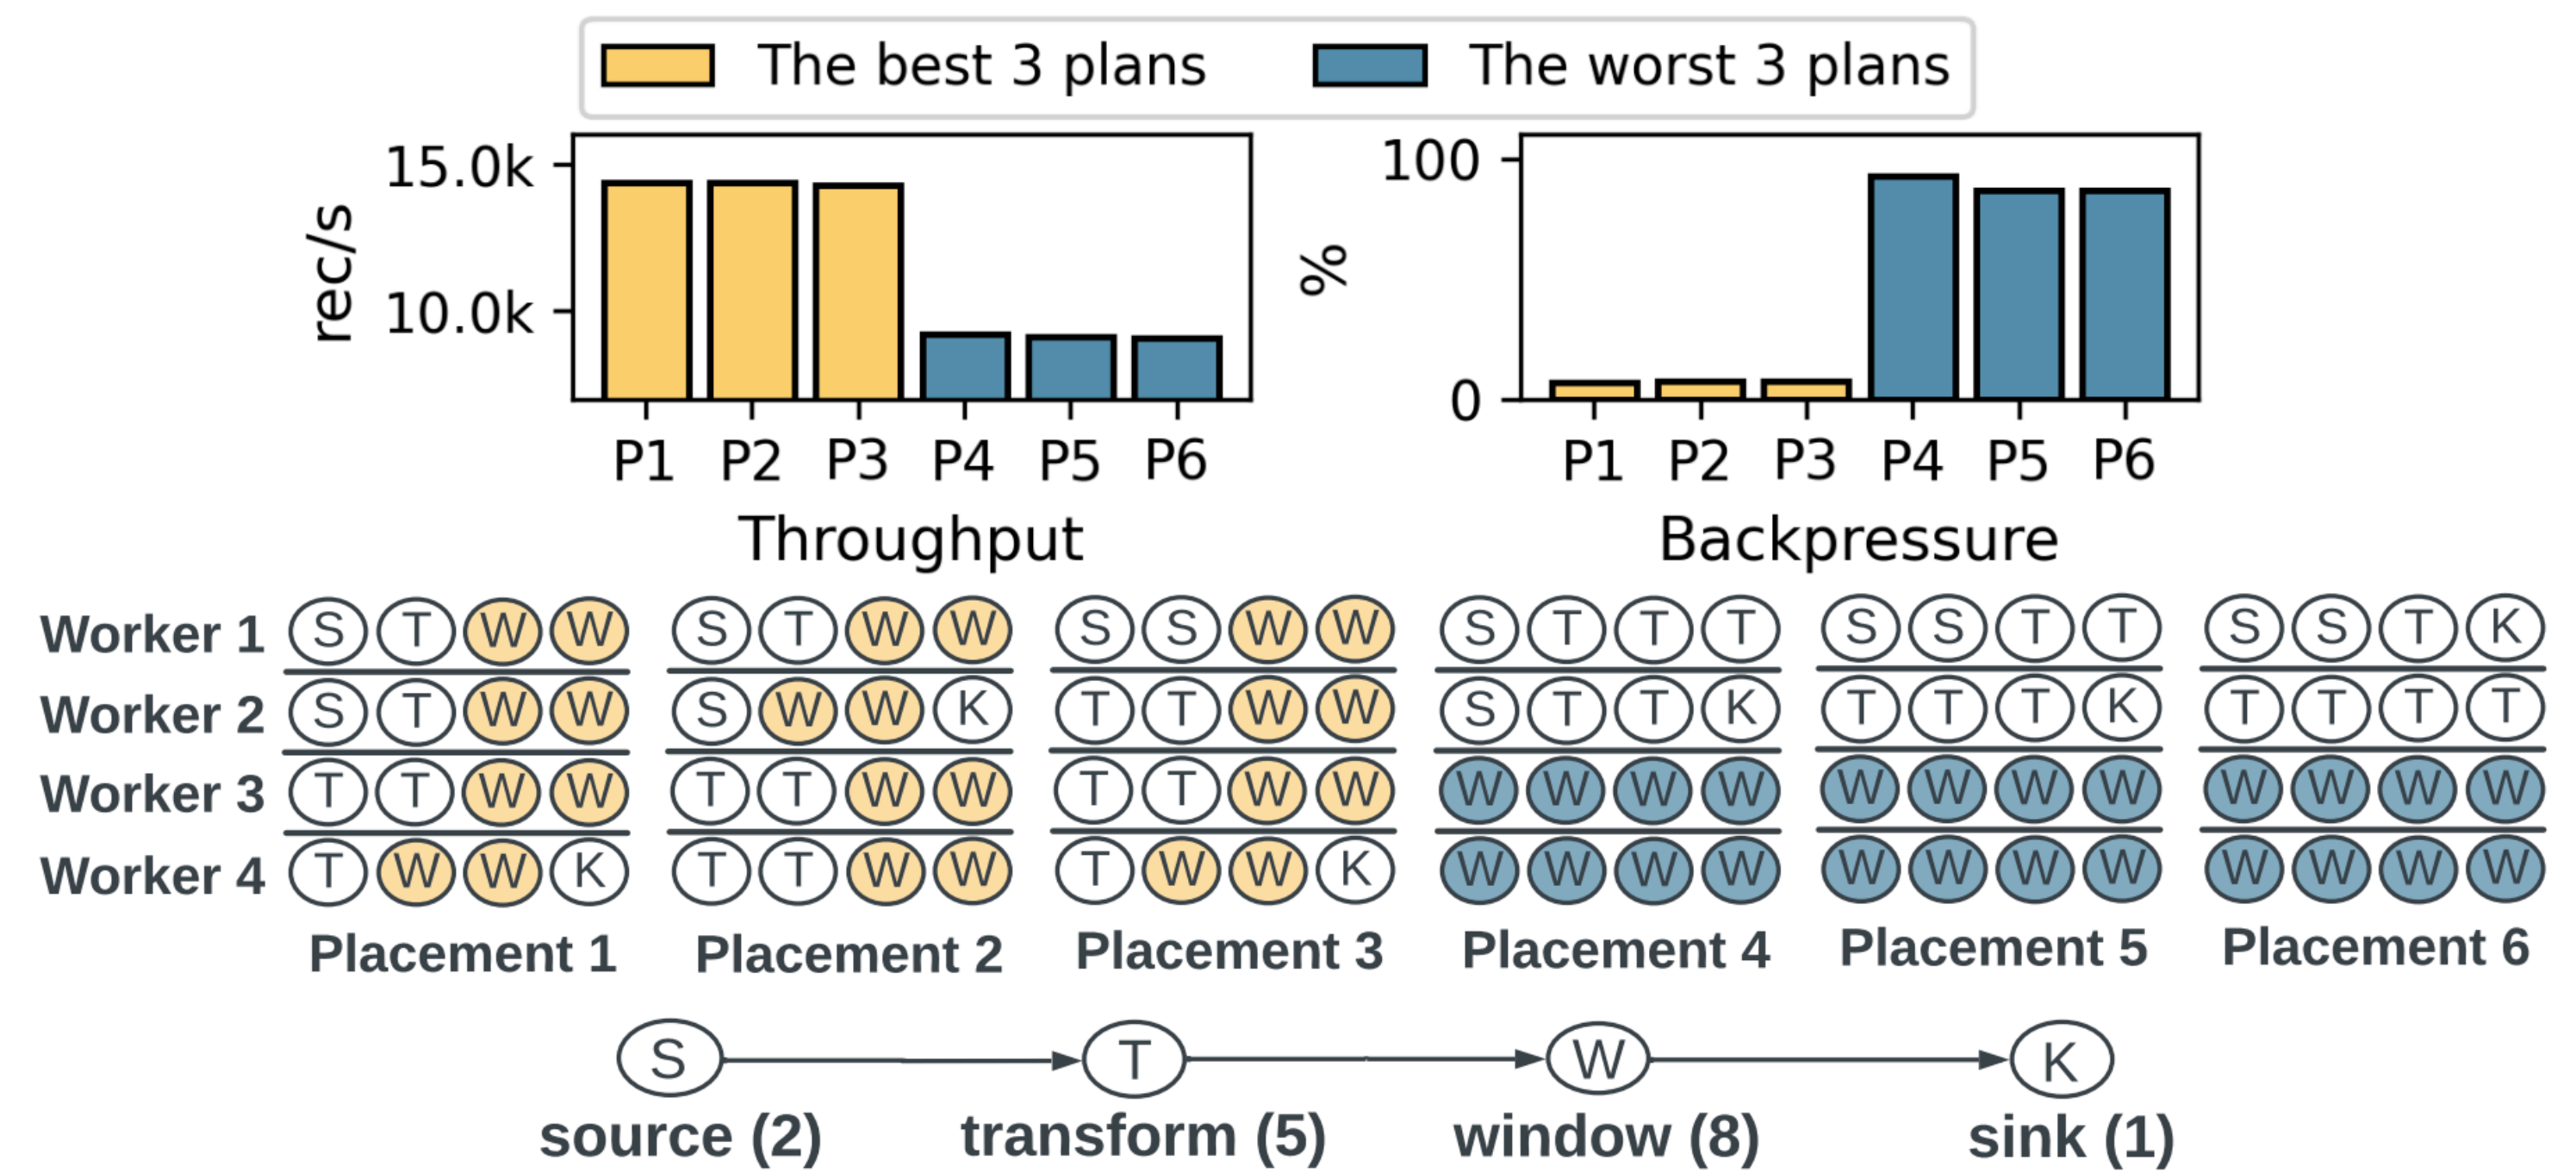
Task placement matters

Task placement can significantly impact the performance of data stream processing applications

- Operators have diverse resource requirements
- Co-locating operators with similar resource requirements will cause resource contention

Limitation of existing approaches:

- Random task placement, e.g. in Flink, can lead to plans with poor performance
- ILP-based approaches are slow and cannot easily incorporate complex objective functions



6 out of 80 alternative placement plans that achieve the best and worst performance for Q1-sliding query

CAPSys: Contention-Aware Placement Search

Goal: Find a good assignment of operator tasks to physical workers, to minimize resource contention and achieve high throughput within resource budget

Massive search space. About 3.25 million alternative placement plans with 5 operators running on 4 workers

- Pruning techniques to effectively reduce the search space and make CAPSys practical for the online setting

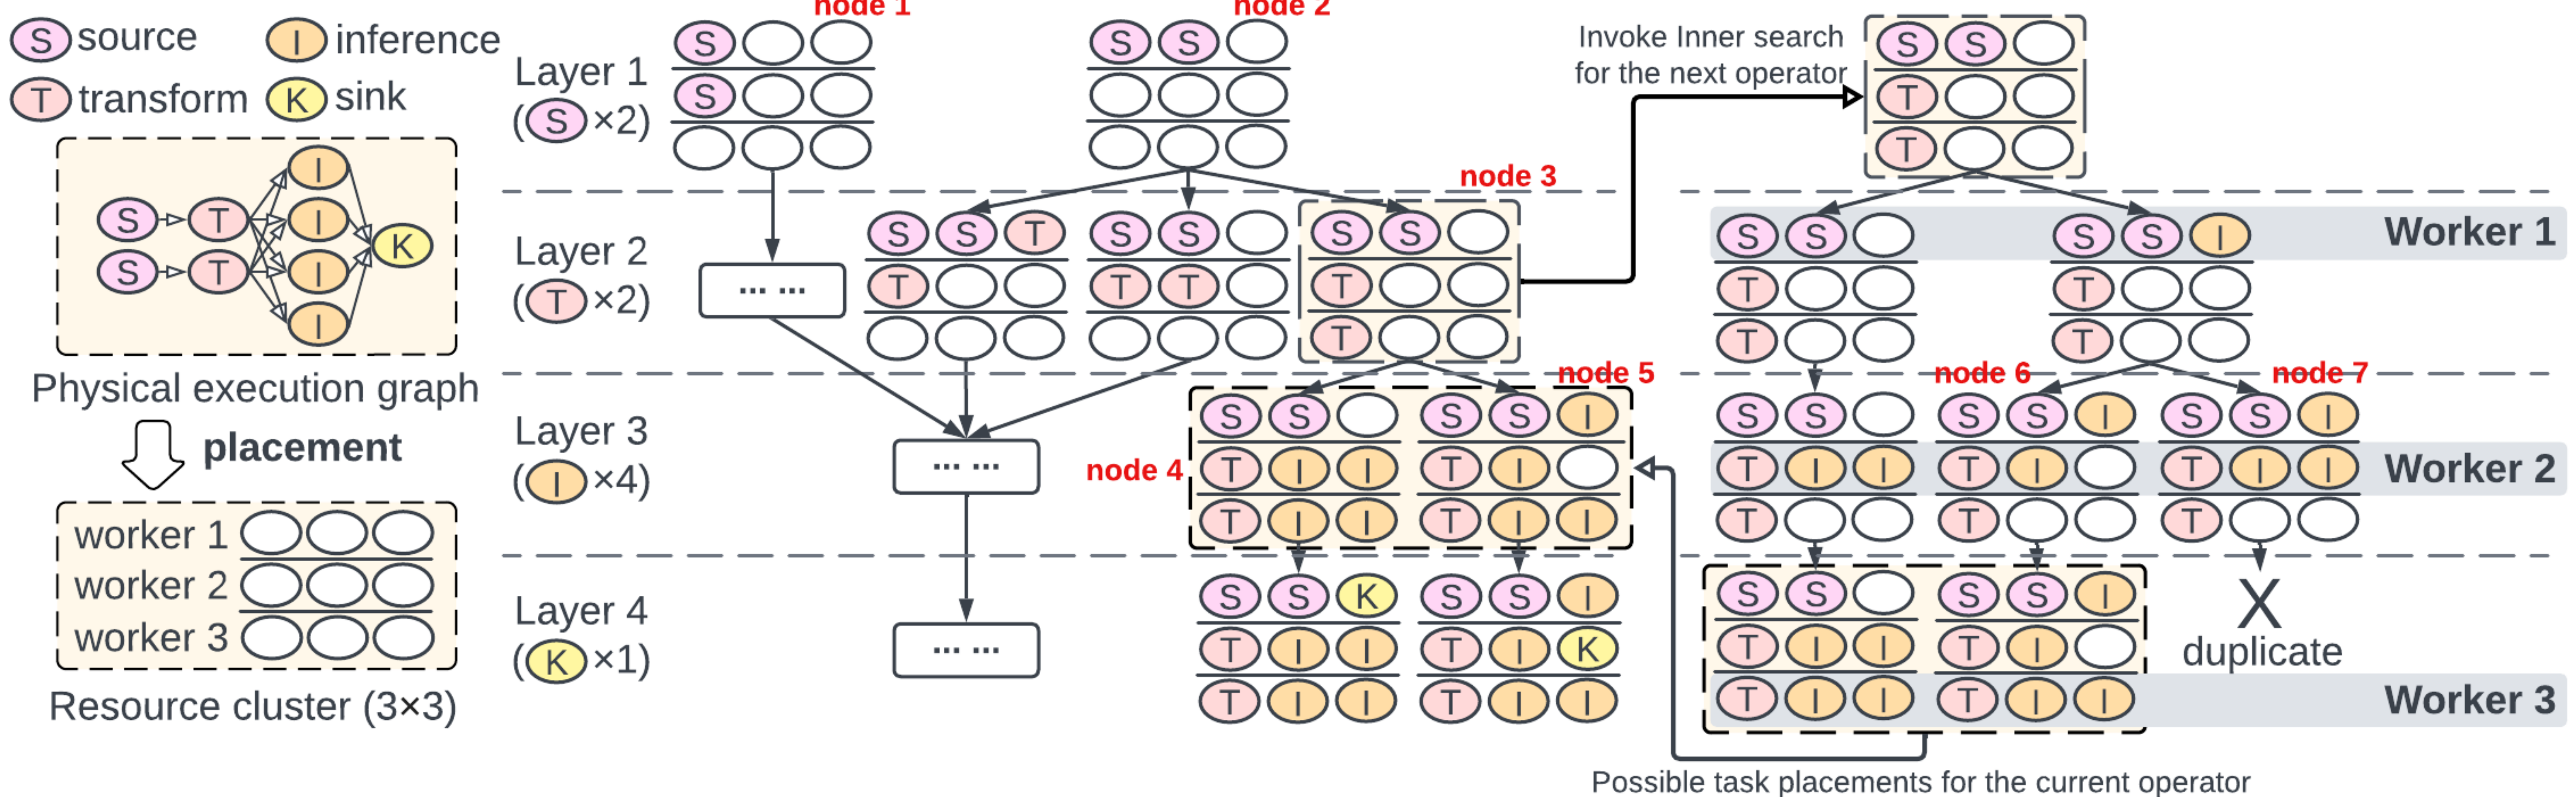
Key idea: Capture the cost of a placement plan by considering the co-location degree of resource-intensive tasks

- Minimize the resource imbalance in the worker cluster
- Consider three dimensions of resources: CPU, Disk IO, Network
- Calculate pareto-optimal solution whose cost is not dominated by any other feasible plan across all dimensions

(a) Task placement problem

(b) Outer search: traverse operators

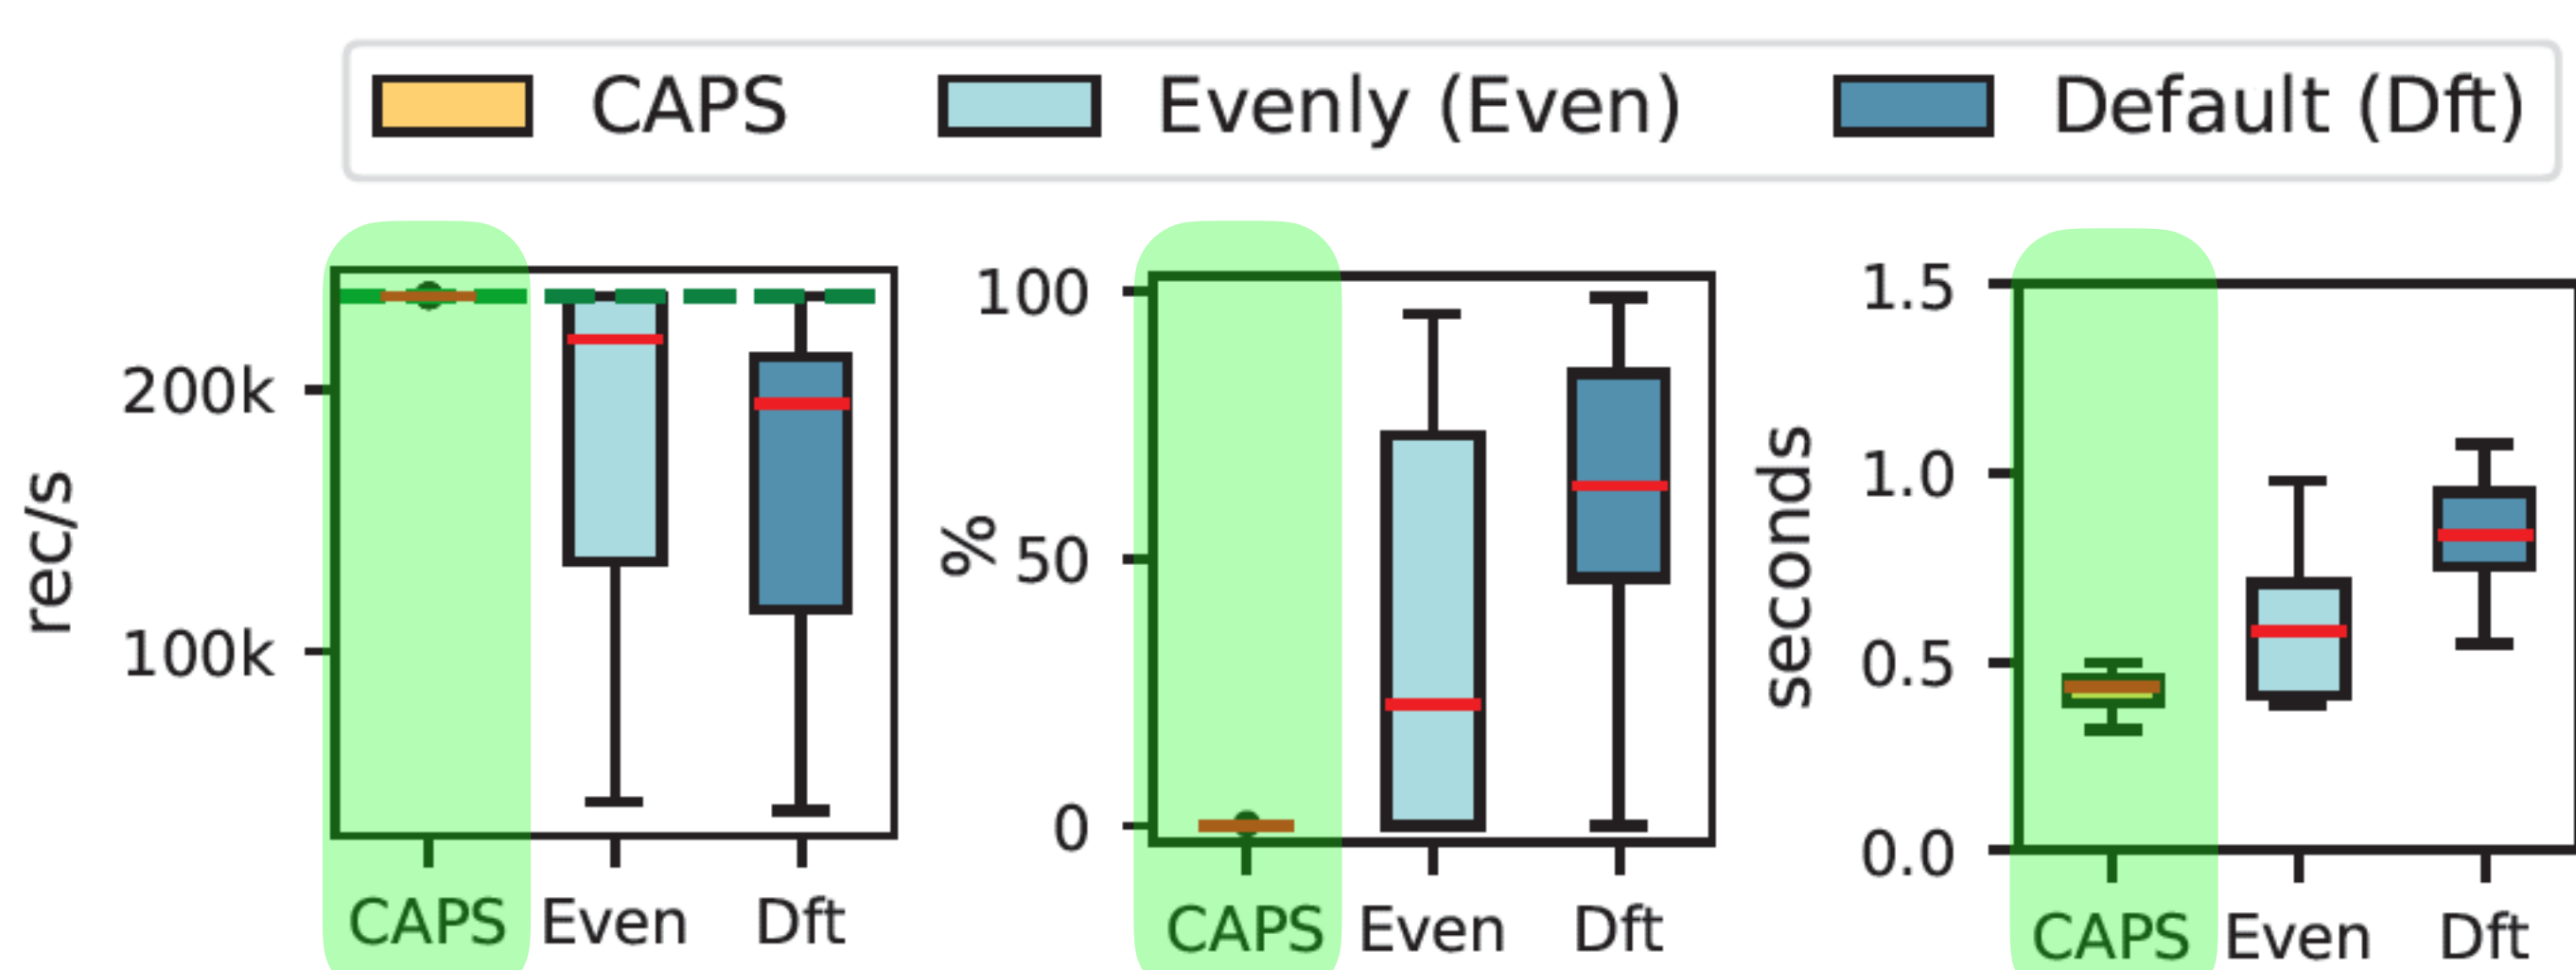
(c) Inner search: traverse workers



Evaluation Results

Individual Query Environment

- Up to 6x higher throughput compared to Flink strategies



Multi-tenant Environment

- CAPSys is the only policy that achieves target throughput across all six queries

Variable workload experiment

- CAPSys avoids over-provisioning, improves convergence behavior and avoids oscillations of auto-scaling controller

Scalability

- CAPSys can find a placement plan within 100ms for a 256 tasks deployment