

Automatic System Tuning for Distributed Stream Processing

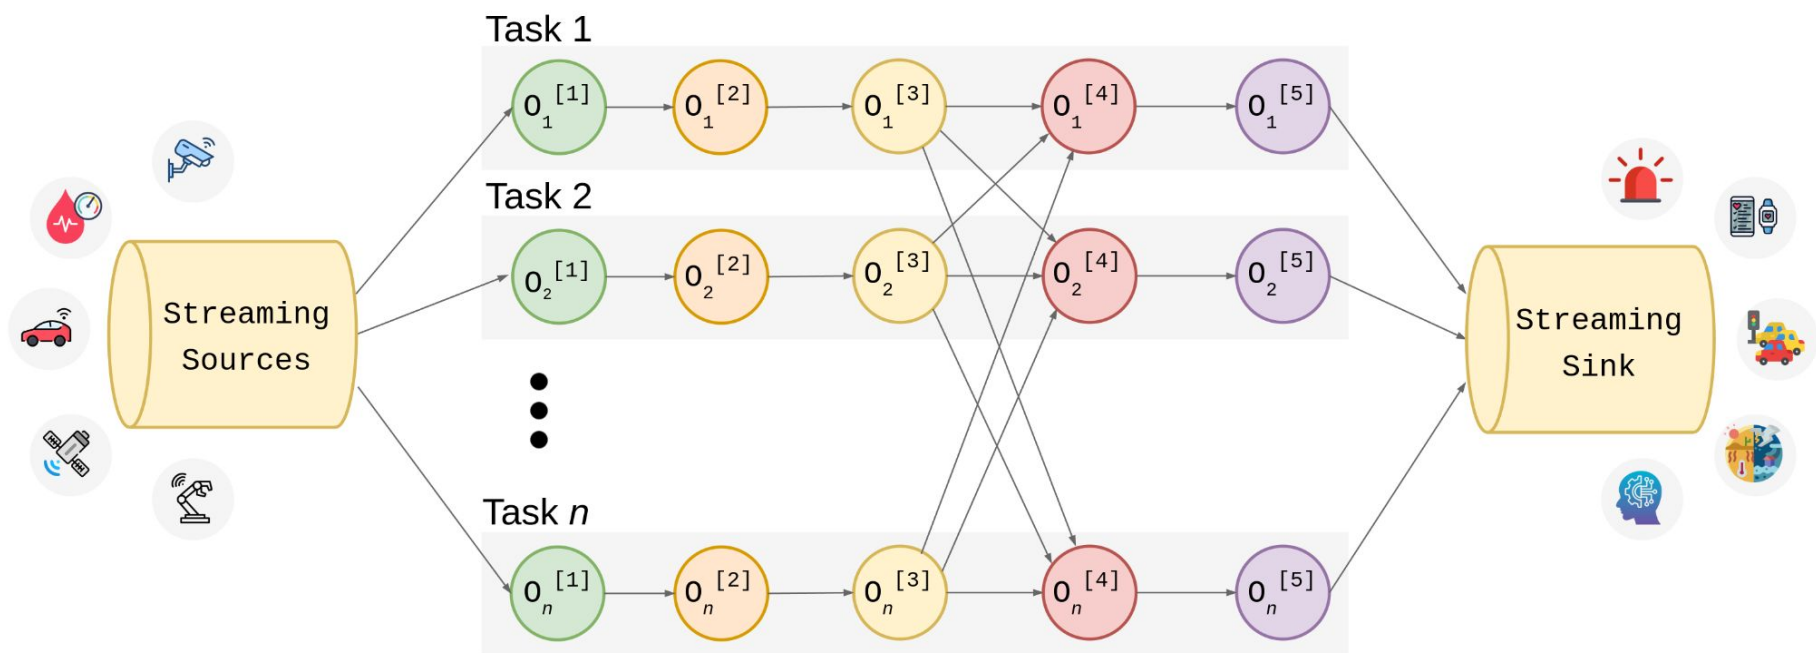
Morgan Geldenhuys, Dominik Scheinert, Odej Kao, Lauritz Thamsen*

Technische Universität Berlin, University of Glasgow*

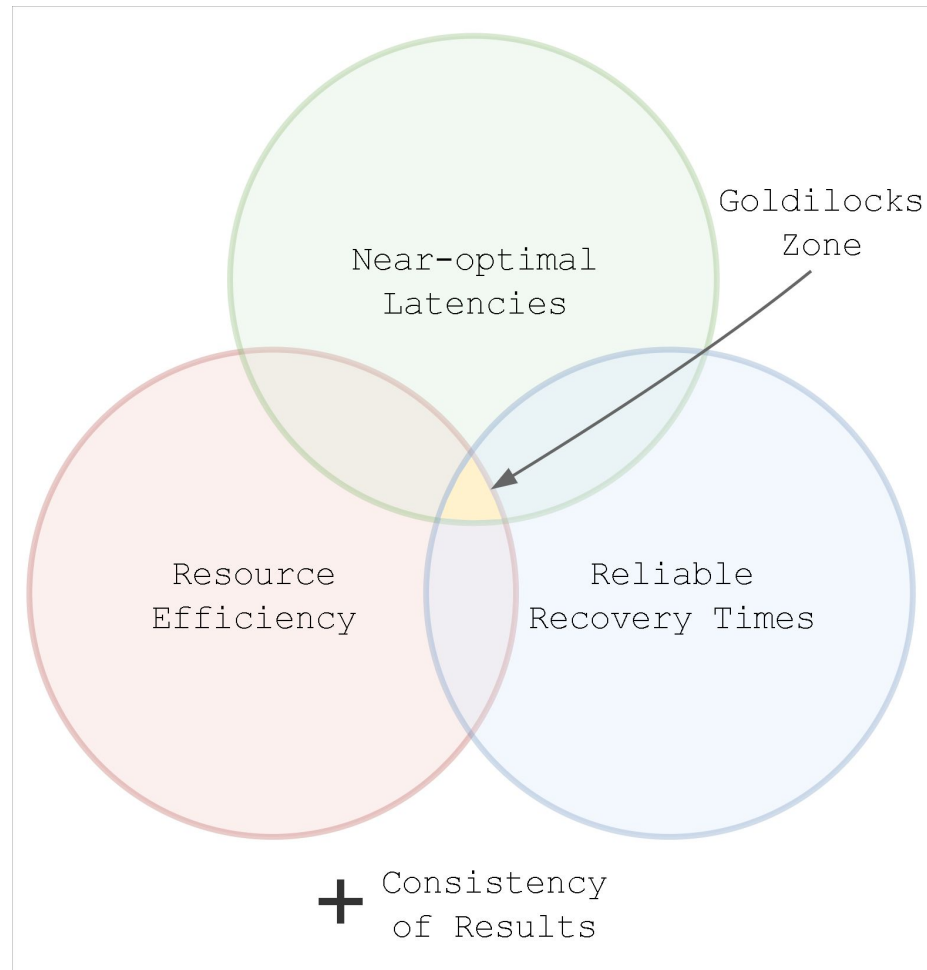
2 May 2024

Background:

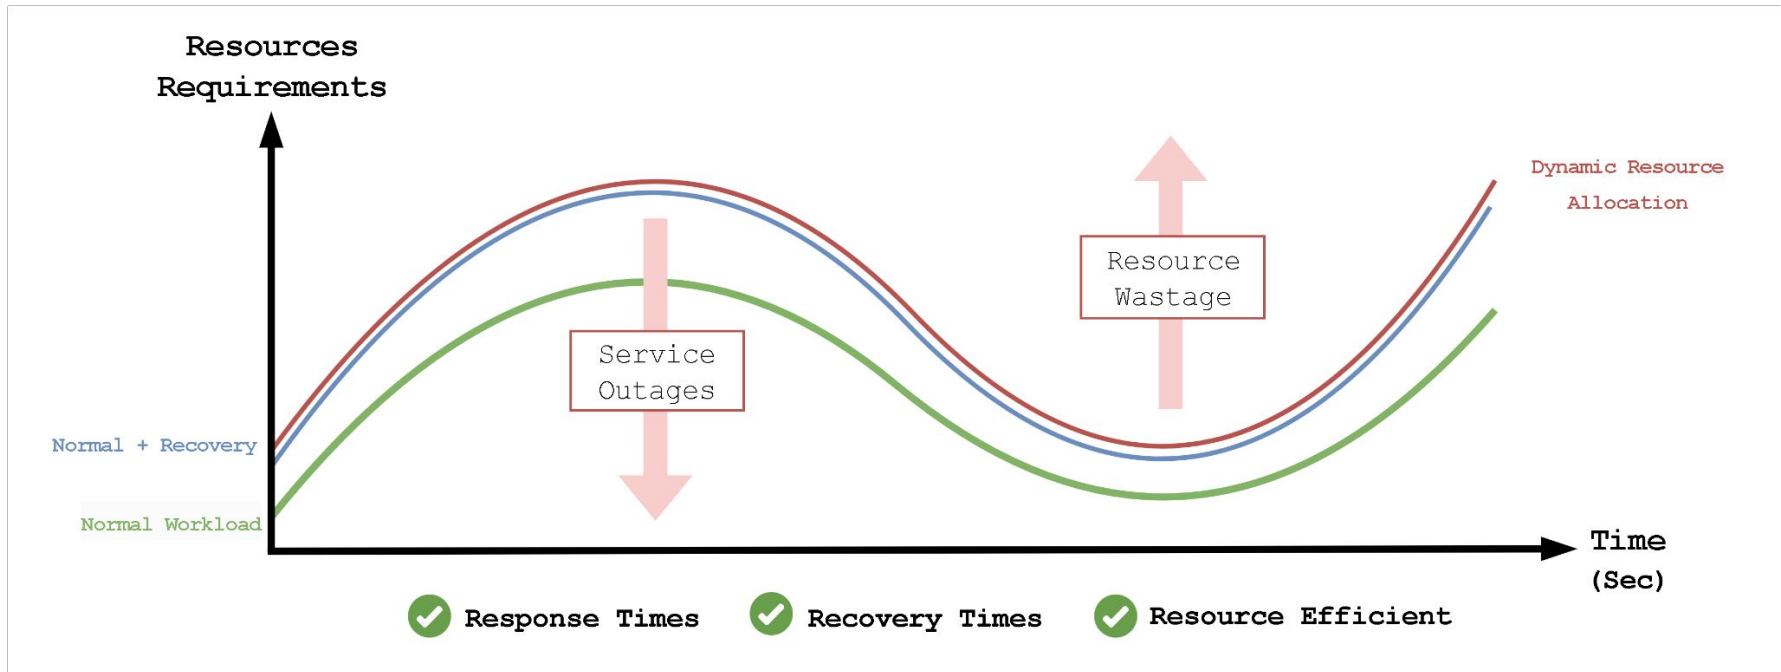
Distributed Stream Processing



Problem: DSP Operational Requirements



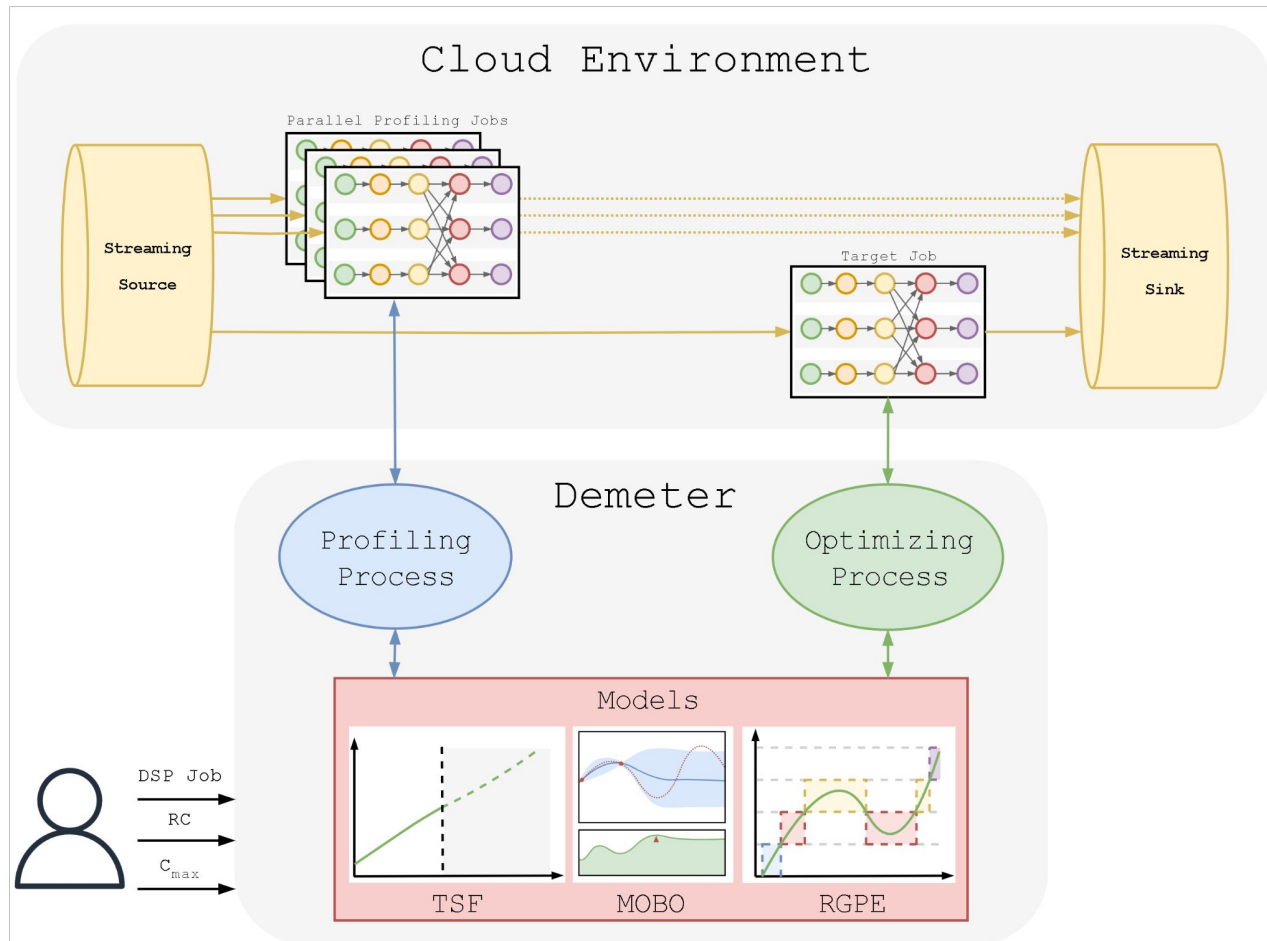
Solution: Runtime Configuration Optimization



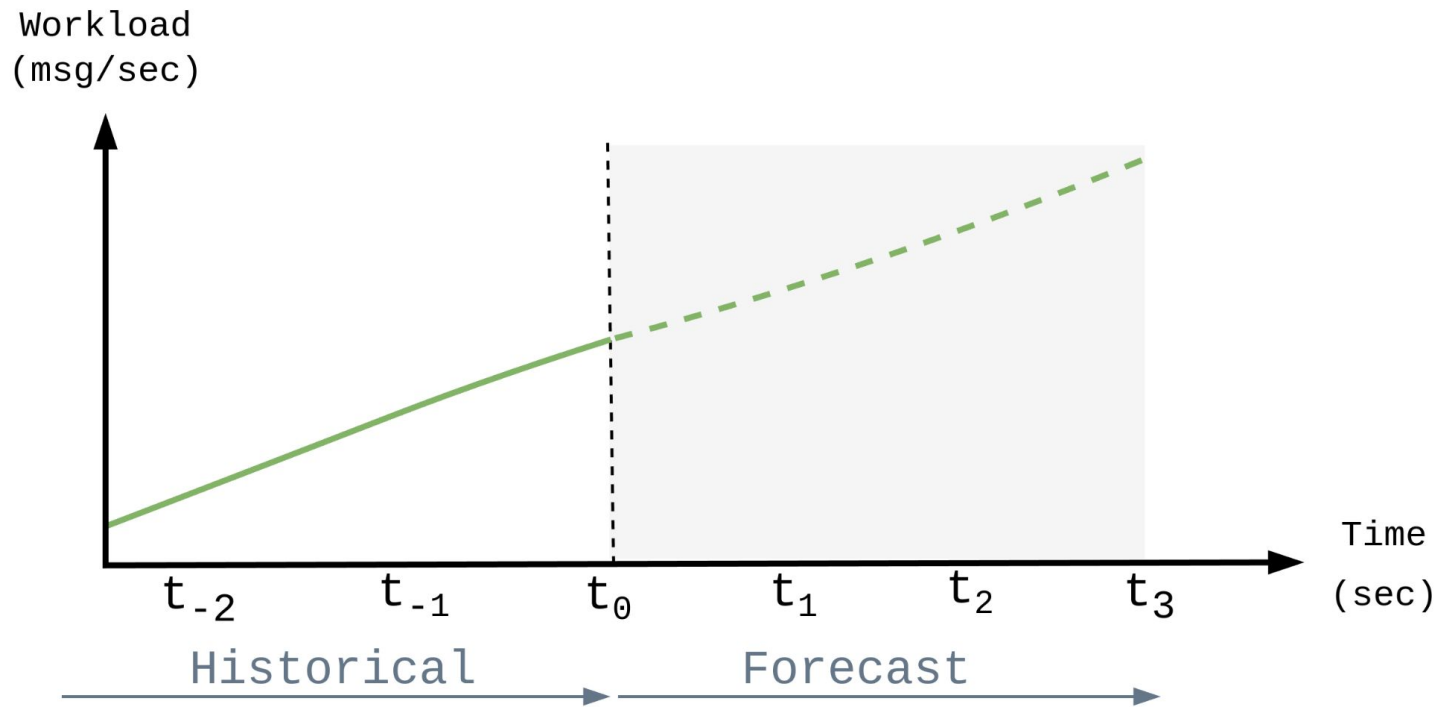
Dynamically match the resource requirements to the changing workload rate through runtime configuration optimization

Approach:

Demeter High-Level Overview

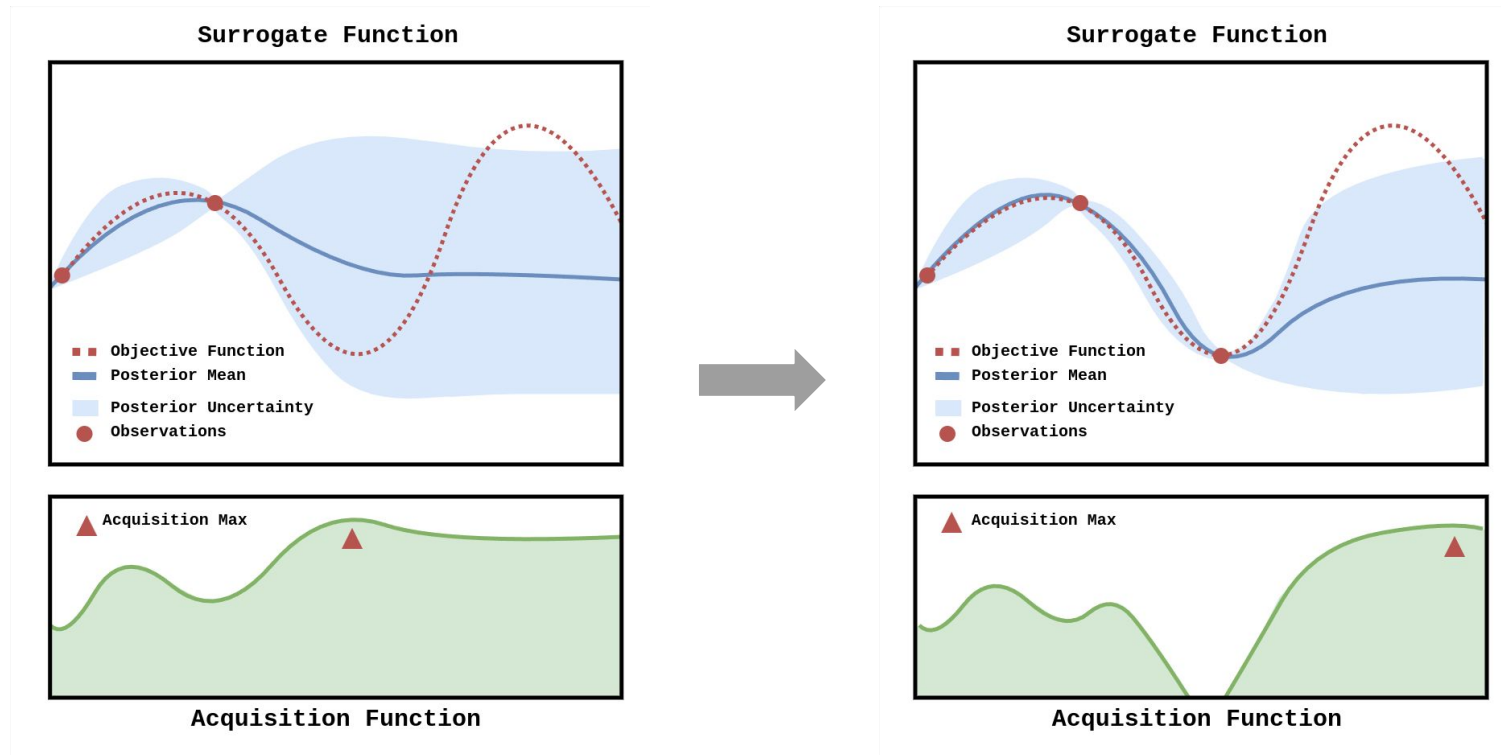


Approach: Modeling Techniques



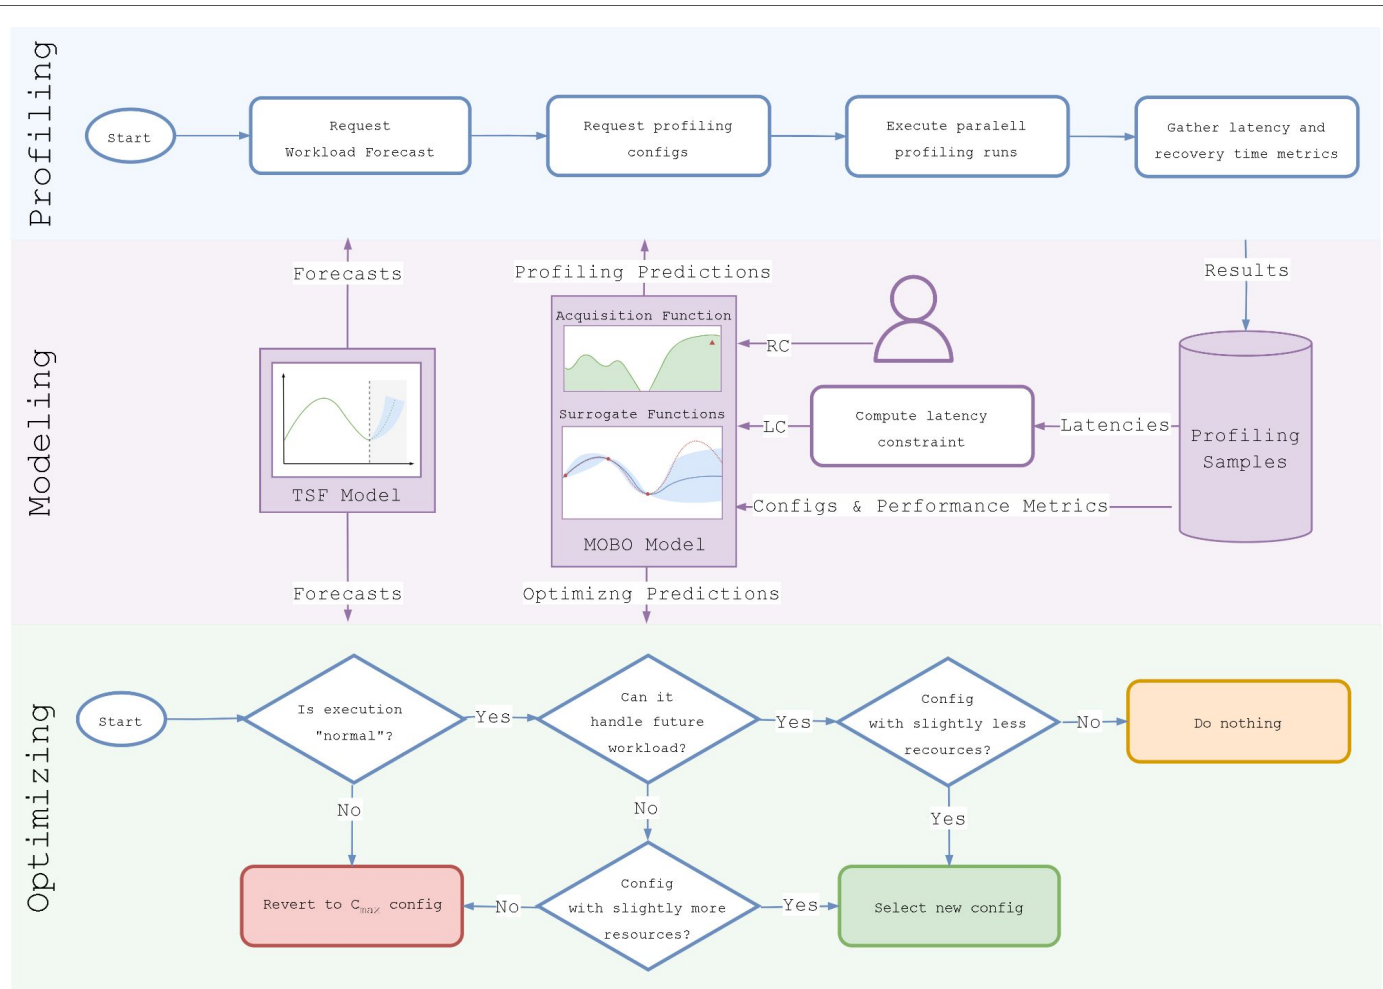
Multi-Step Ahead Time Series Forecasting (TSF)

Approach: Bayesian Optimization




Multi-Objective Bayesian Optimization (MOBO)

Approach: Profiling & Optimizing Processes

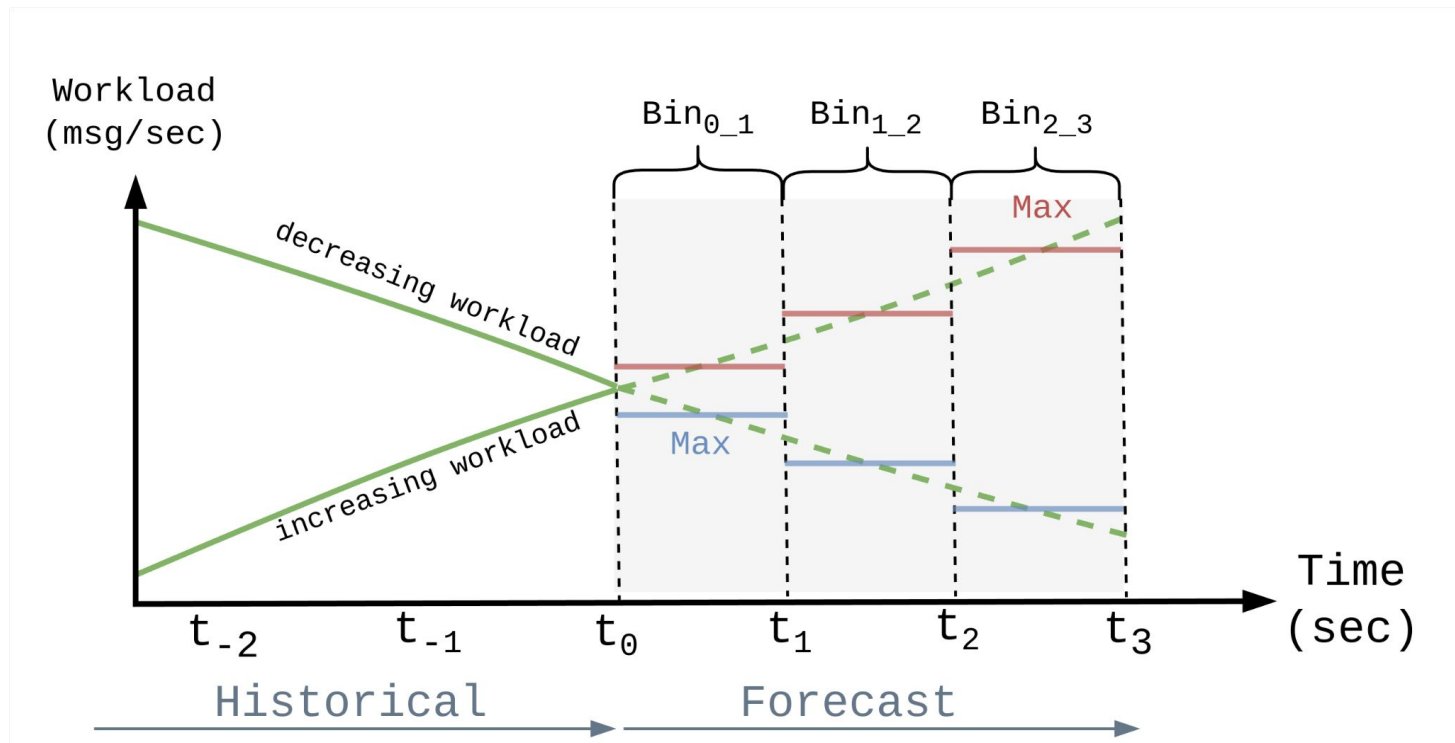


Approach: Modeling Dynamic Workloads




Rank-Weighted Gaussian Process Ensembles (RGPE)

Approach: Proactive Workload Predictions



Approach: Configuration Parameters

- Workers (w)
- CPU Cores (c)
- Checkpoint Interval (t)
- Task Slots (t)
- Memory (m)



Evaluation:


Experimental Setup

- Experiments (Kubernetes + Apache Flink):
 - Yahoo Streaming Benchmark (click-stream analytics; 25K - 100K)
 - Top Speed Windowing (traffic monitoring; 35K - 125K)
- Baselines & Configuration Space:

Method	Workers	Task Slots	CPU Cores	Memory (mb)	Checkpoint Interval (s)
Static (C_{\max})	24	1	1	4096	10
Reactive	1 - 24	1	1	4096	10
DS2	1 - 24	1	1	4096	10
Demeter	1 - 24	1 - 4	1 - 3	1024 - 4096	10 - 90

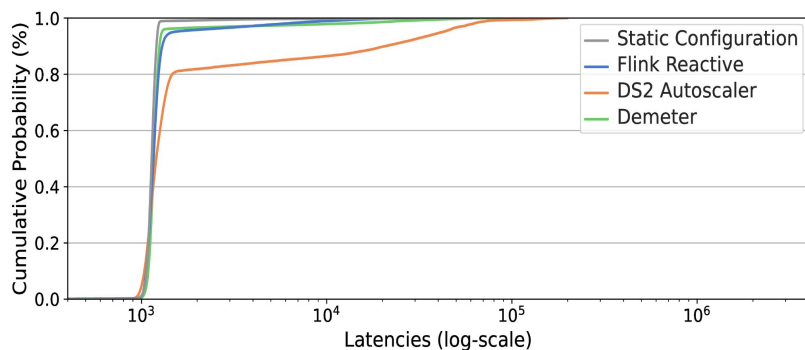
- 18h experiment length; 23 failures; 180s recovery time constraint

Evaluation: Workload, Failures, & Configuration Parameters




YSB Experiment

Evaluation: End-to-end Latencies Comparison



YSB Experiment


- 1st: Static with near-optimal $L_{avg} \approx 1000ms$
- 2nd: Demeter at 95% of near-optimal
- 3rd: Reactive at 94% of near-optimal
- 4th: DS2 at 82% of near-optimal



TSW Experiment

- 1st: Static with near-optimal $L_{avg} \approx 1000ms$
- 2nd: Demeter at 95% of near-optimal
- 3rd: Reactive at 85% of near-optimal
- 4th: DS2 at 70% of near-optimal

Evaluation: Recovery Time Comparison



YSB Experiment

Method	Failures	Changes
Static		-
Demeter		33
Reactive		87
DS2		77

Recovery time averages: 1st Static 96s, 2nd Demeter 99s (0.03), 3rd Reactive 176s (0.83), 4th DS2 171s (0.78)



TSW Experiment

Method	Failures	Changes
Static		-
Demeter		30
Reactive		49
DS2		73


Recovery time averages: 1st Static 107s, 2nd Demeter 113s (0.05), 3rd DS2 162s (0.52), 4th Reactive 174s (0.63)

Recovered within 180s
 Recovered within 360s
 Longer than 360s
 No Result

Evaluation: Resource Usage Comparison



YSB Experiment



TSW Experiment

Summary

- Demeter: Multi-configuration resource optimization for DSP jobs
 - Dynamic workloads
 - Exactly-once processing guarantees
- Experiments showed that Demeter is able to:
 - Deliver a stable near-optimal service
 - Recover reliably from failures
 - Improve resource utilization
- Future Work:
 - Improve the efficiency of the bayesian optimization component
 - Implement bayesian optimization component using GPU resources
 - Perform extended execution runs to observe efficiency over time

Thank you

Have a great conference!

Morgan K. Geldenhuys

Researcher @ DOS

<https://www.dos.tu-berlin.de>

morgan.geldenhuys@tu-berlin.de

