What to do when HPC-FaaS Problems Stare at Your Face?

Rohan Basu Roy

Ph.D. Candidate

Advisor:
Prof. Devesh Tiwari

LVX VERITAS VIRTVS

Northeastern University

HPC Applications are Being Designed as DAGs

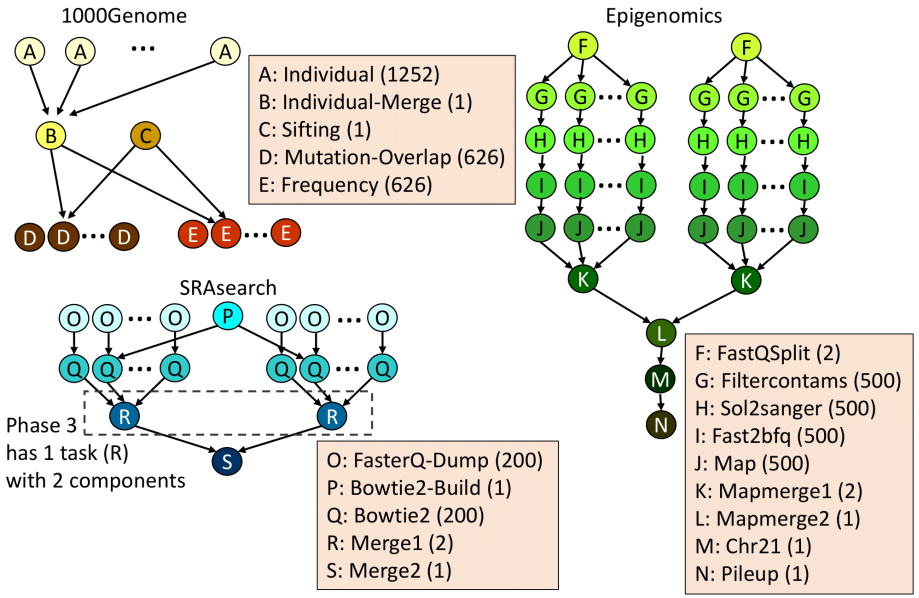

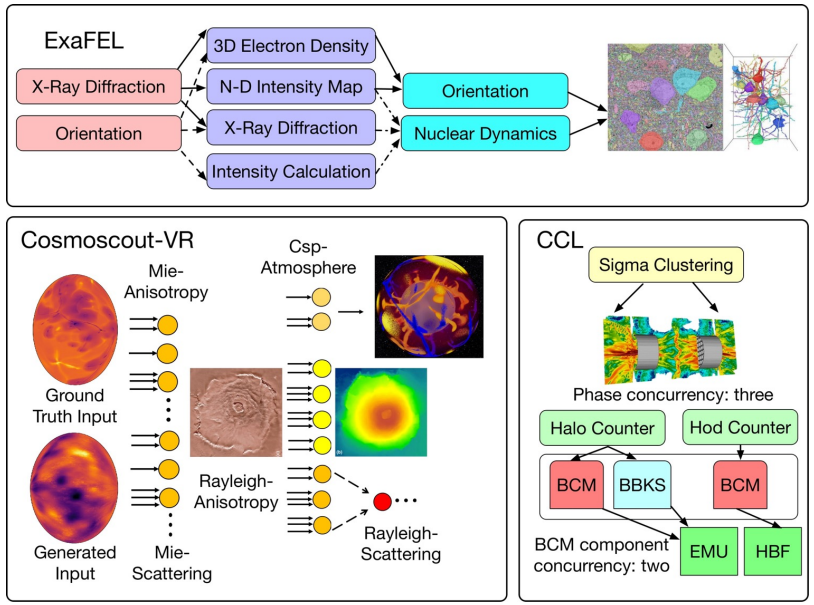

DAG execution on a serverless platform has both advantages and disadvantages

Challenge 1: Cold Starts are Always Problematic in Serverless Computing

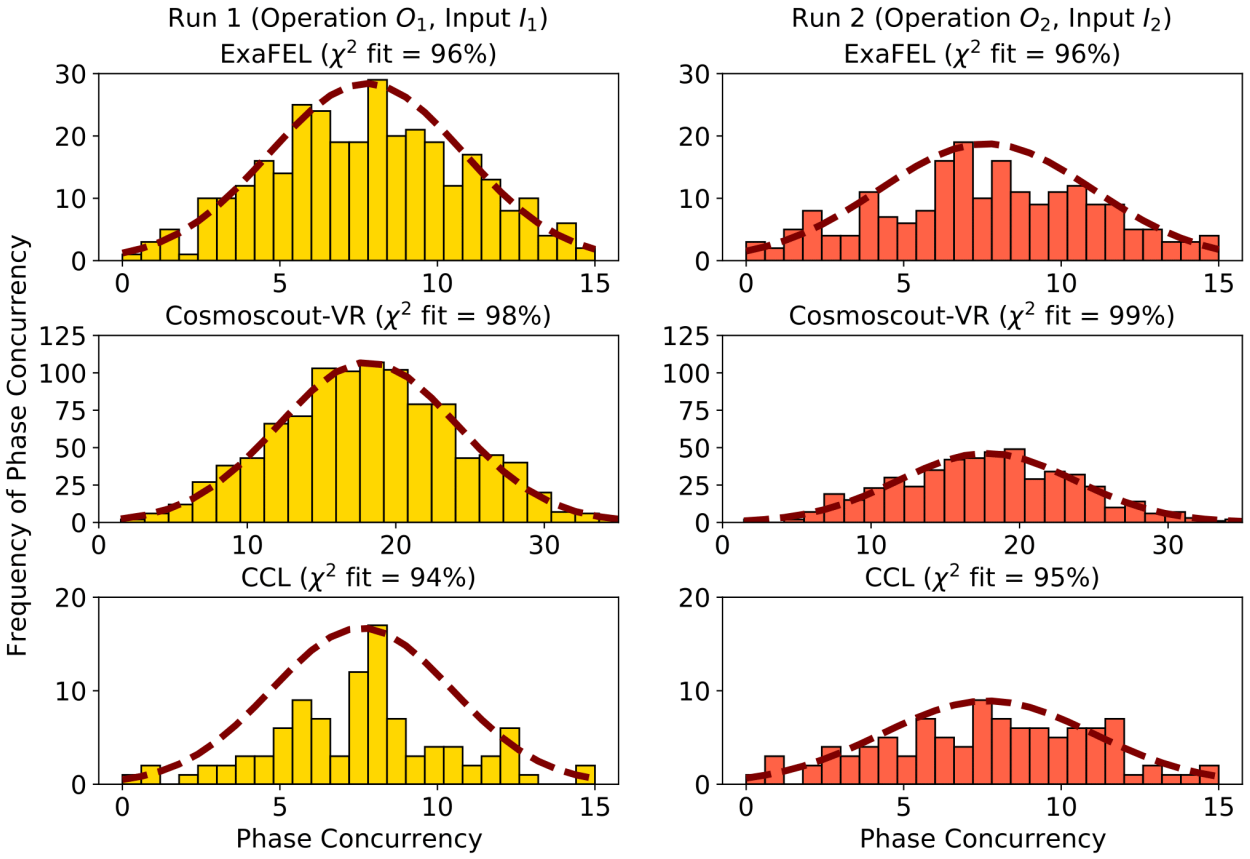

The invocation of components follow a Weibull distribution in HPC DAGs

This characteristics enables to decouple warming up the function code from warming up the microVM (hot start).

DayDream [SC 22]

Employ hot starting of components to reduce cost and service time

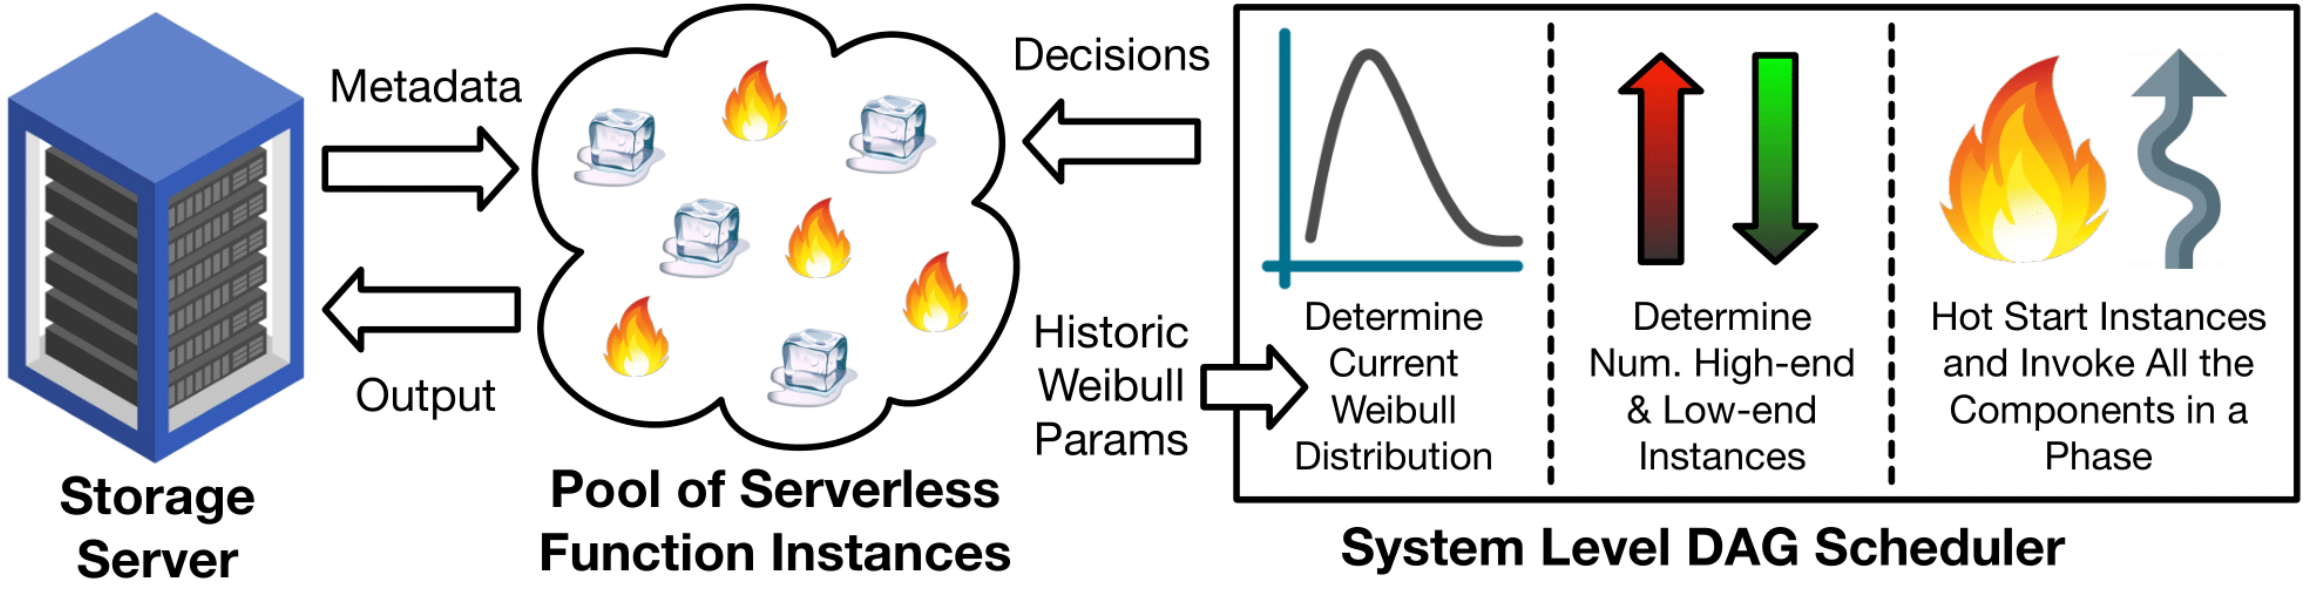

Challenge II: Scaling Overhead Can be a Significant Portion of Serverless Service Time

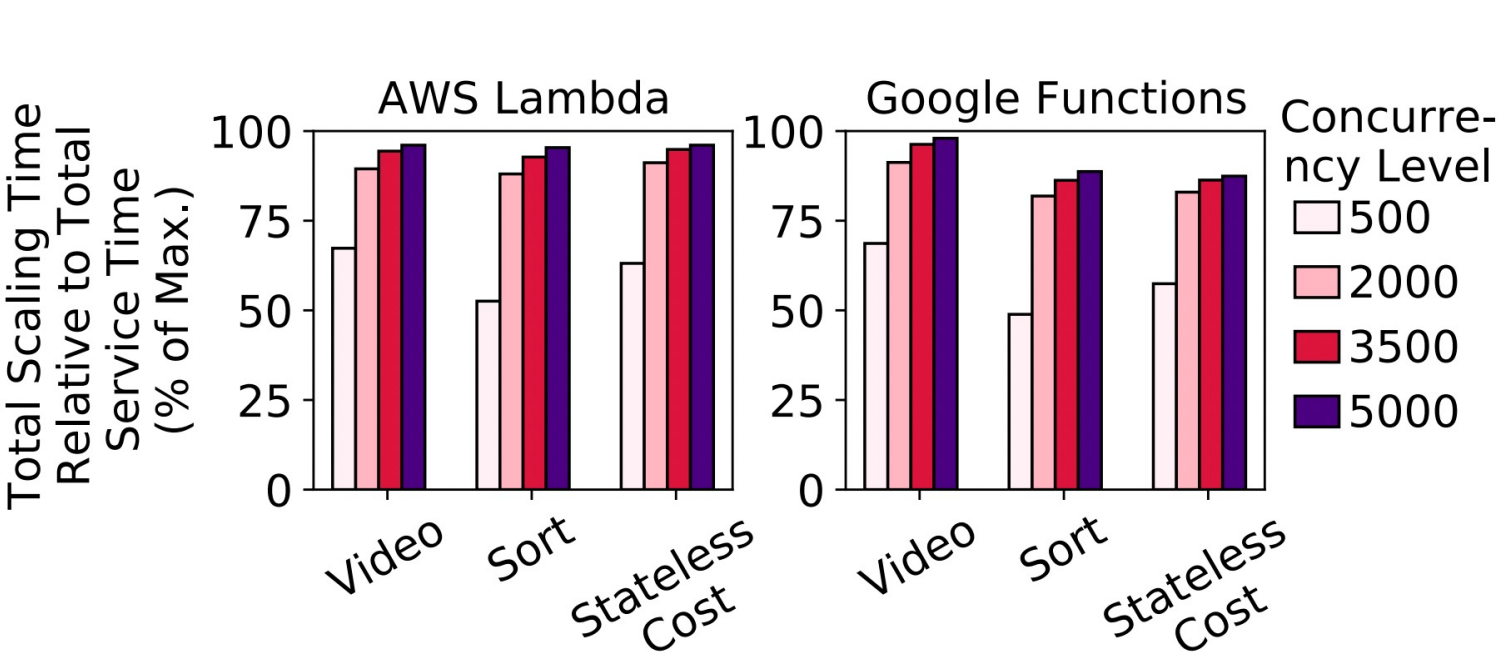

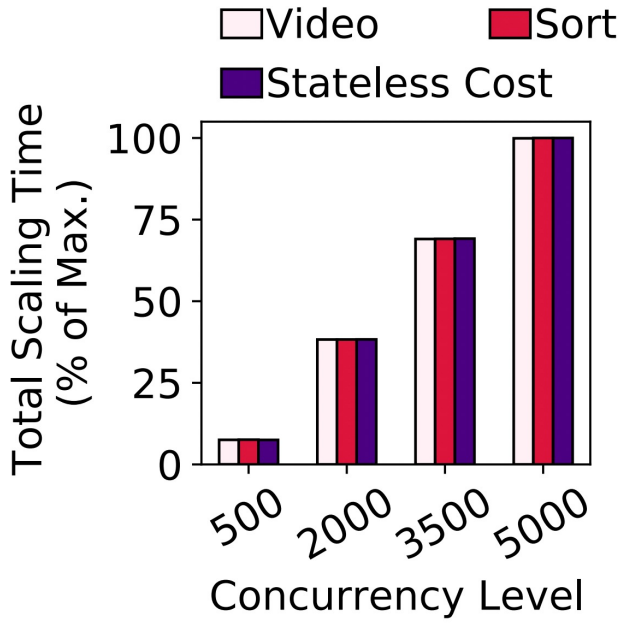

Service Time (Execution Time + Scaling Time) becomes worse with concurrency due to increased scaling time

Reduce Scaling Time by Packing Multiple Functions in a Function Instance

Invoked Functions

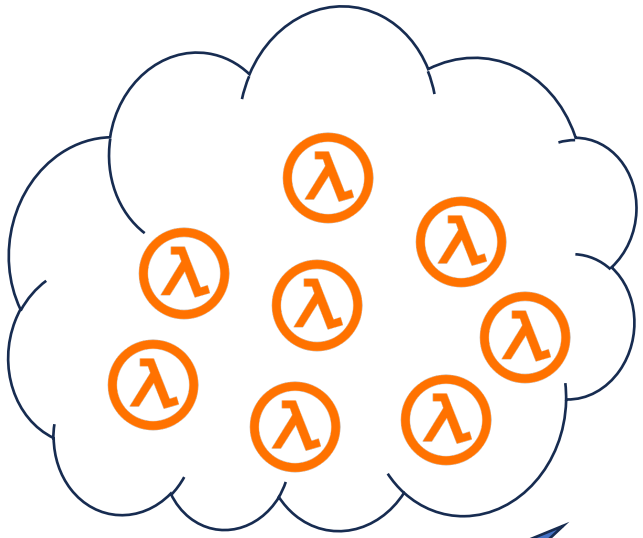

Functions Packed Inside Instances

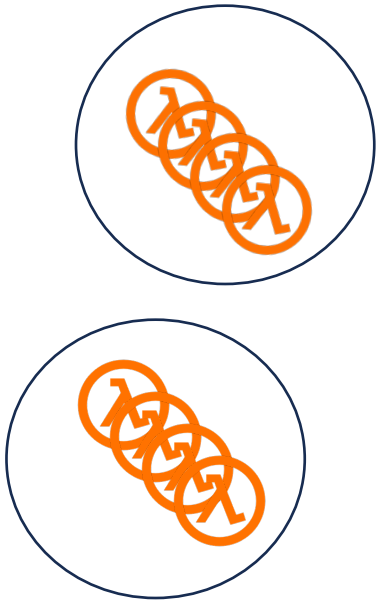

Packing effectively reduces the number of instances to be spawned

High Scaling Time

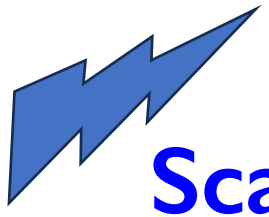

Low Scaling Time

ProPack [HPDC 23]

Determine the Optimal Packing Degree of Concurrent Function Invocations

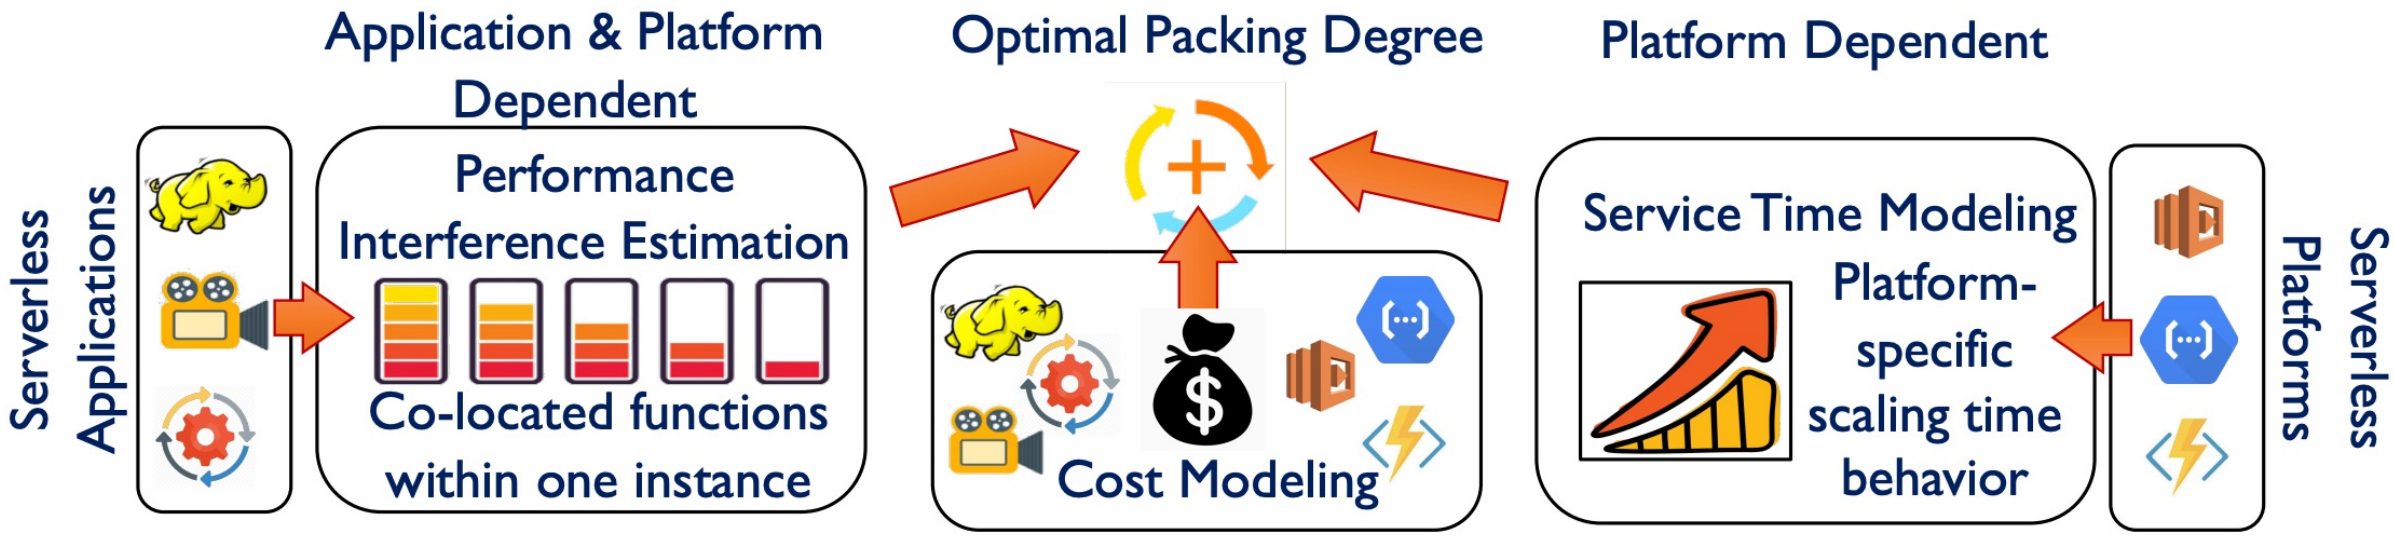

With applying ProPack on FuncX, functions scale 22% faster than AWS Lambda for a concurrency of 1000

Challenge III: Serverless Functions Cannot Directly Communicate with Each Other

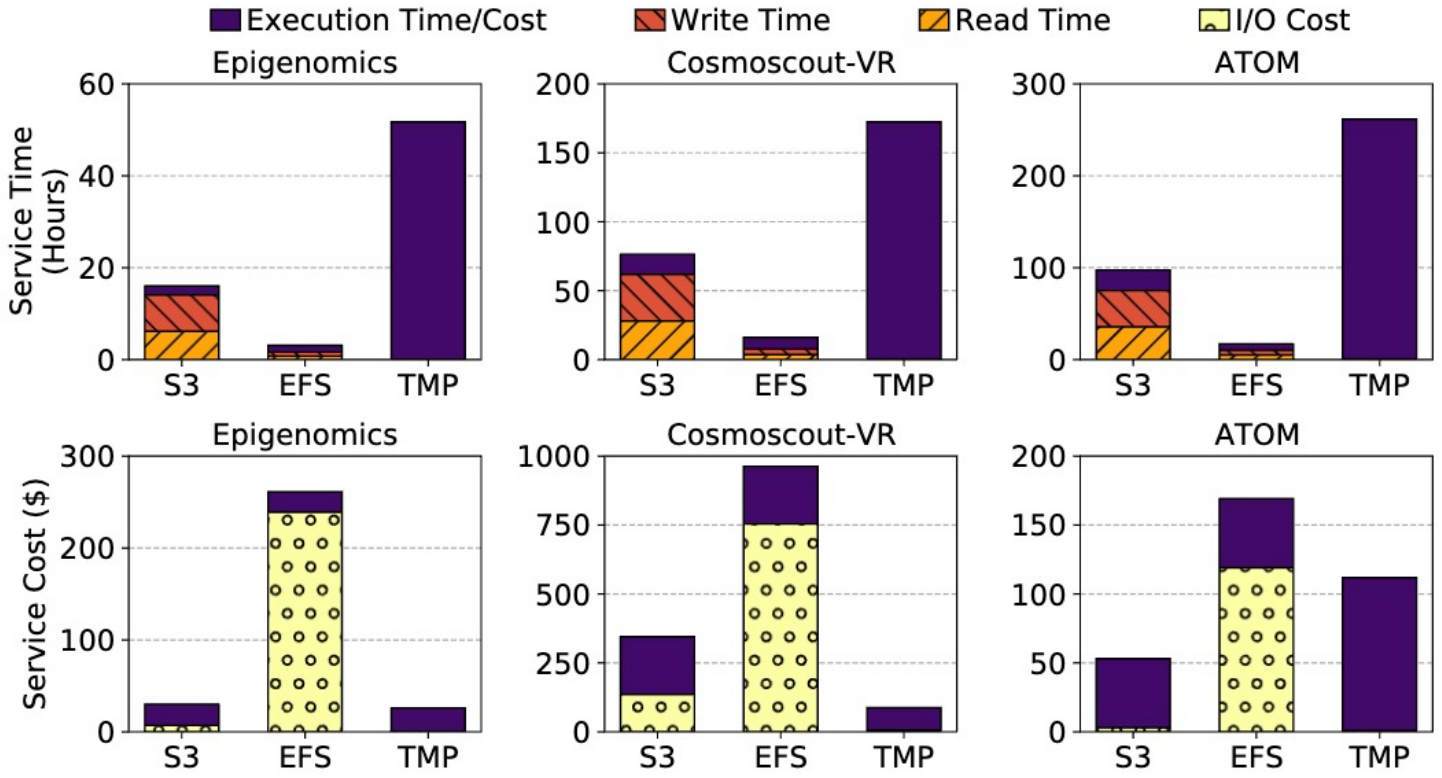

I/O time and cost can be significant

Production Serverless Platforms have Different Tiers of Storage Options

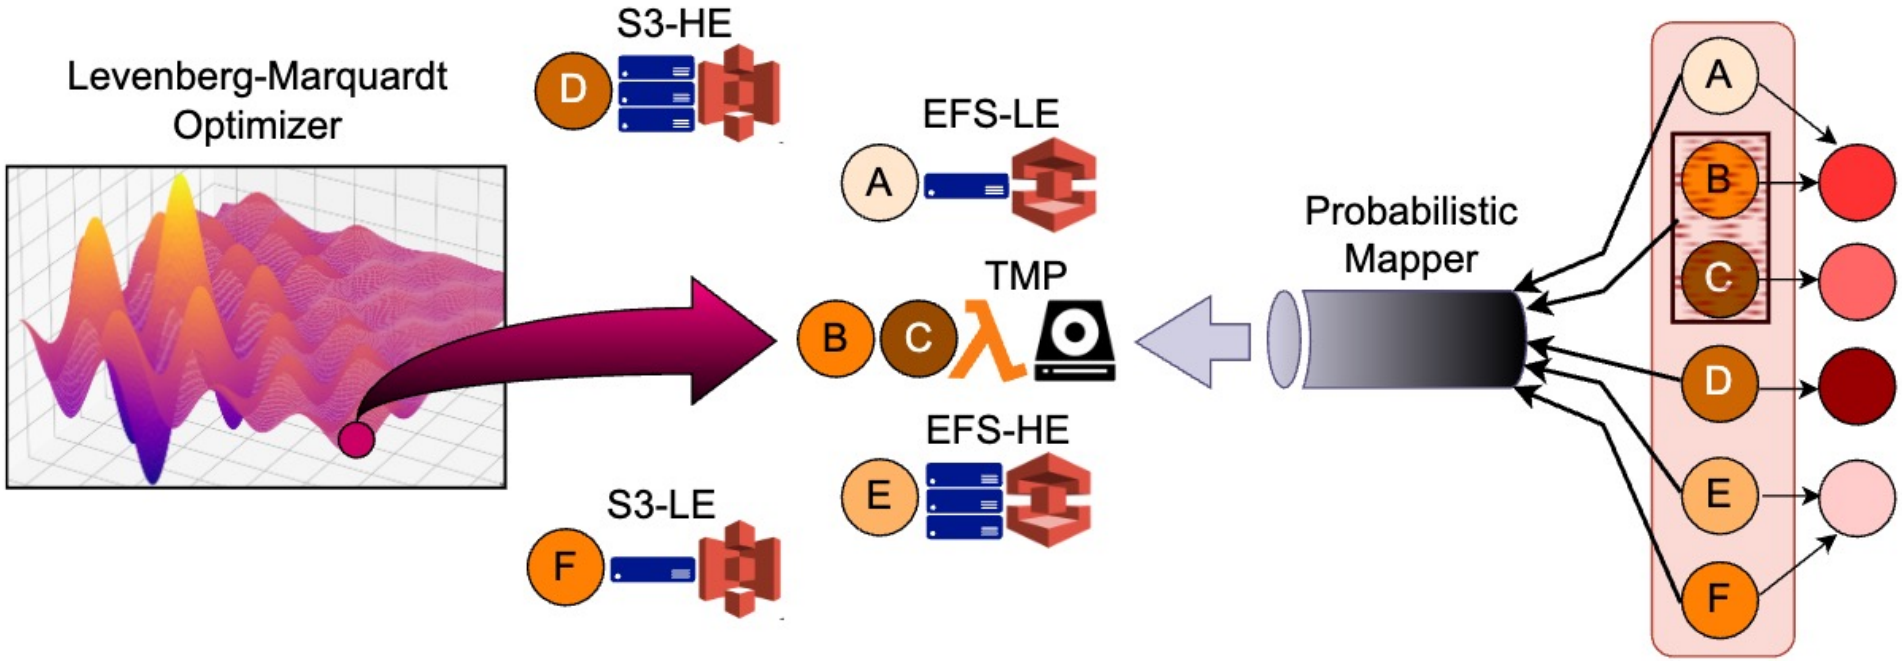

Opportunistically selecting different tiers of storage for different components of a DAG can improve I/O performance

Can we integrate these optimizations in Globus Compute – Federated Function as a Service ?