# Building Scalable Intelligent Information Systems


Yanlei Diao and
Angelos-Christos Anadiotis

**Ecole Polytechnique**

**Institut Polytechnique de Paris**

Yanlei Diao

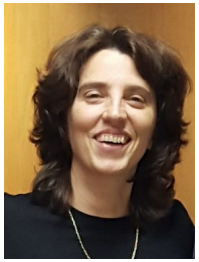

Ioana Manolescu

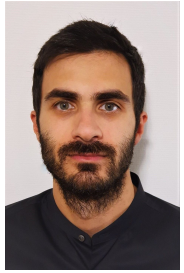

Angelos-Christos Anadiotis

**A range of research projects for cloud and scalable data analytics:**

1. A Unified Data Analytics Optimizer for Cloud Computing **(Diao)**
2. Elastically Scaling Heterogeneous Workloads in Virtualized Servers **(Anadiotis)**
3. Optimizing Big Data Computations: Queries and Algebra in a Single Framework **(Manolescu)**

Inria

## Grand Challenge

Design an **algorithmic foundation** that enables the development of all necessary pillars of **big & fast data analysis**

## Scientific Goals

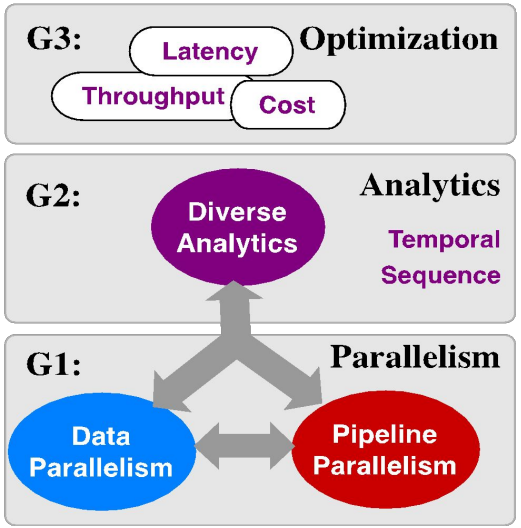


- A novel multi-objective optimization framework
- Distributed, low-latency algorithms for temporal and stream analytics
- System support for scale & latency with maximum degree of parallelism

# 1. A Unified Data Analytics Optimizer for Cloud Computing

Cloud data analytics involves **millions of jobs** and **100's thousands of machines**. What is the best way to allocate resources and execute all the jobs?

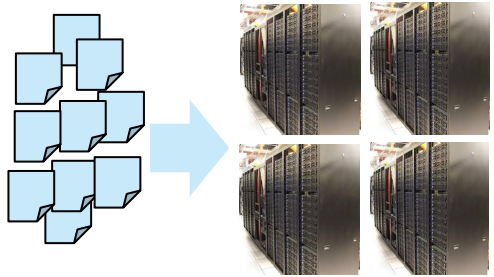

**User Objectives**

**Green Computing**

- *In-situ modeling* of user objectives for analytical tasks based on runtime observations and Deep Learning
- *Multi-objective optimization* automatically adapts cluster and cloud resources to user objectives

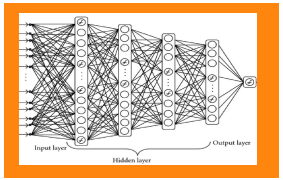


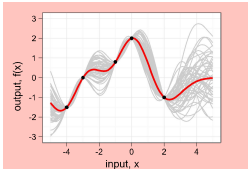


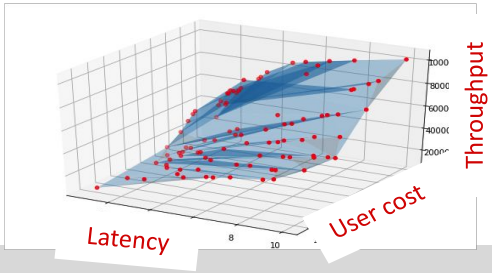


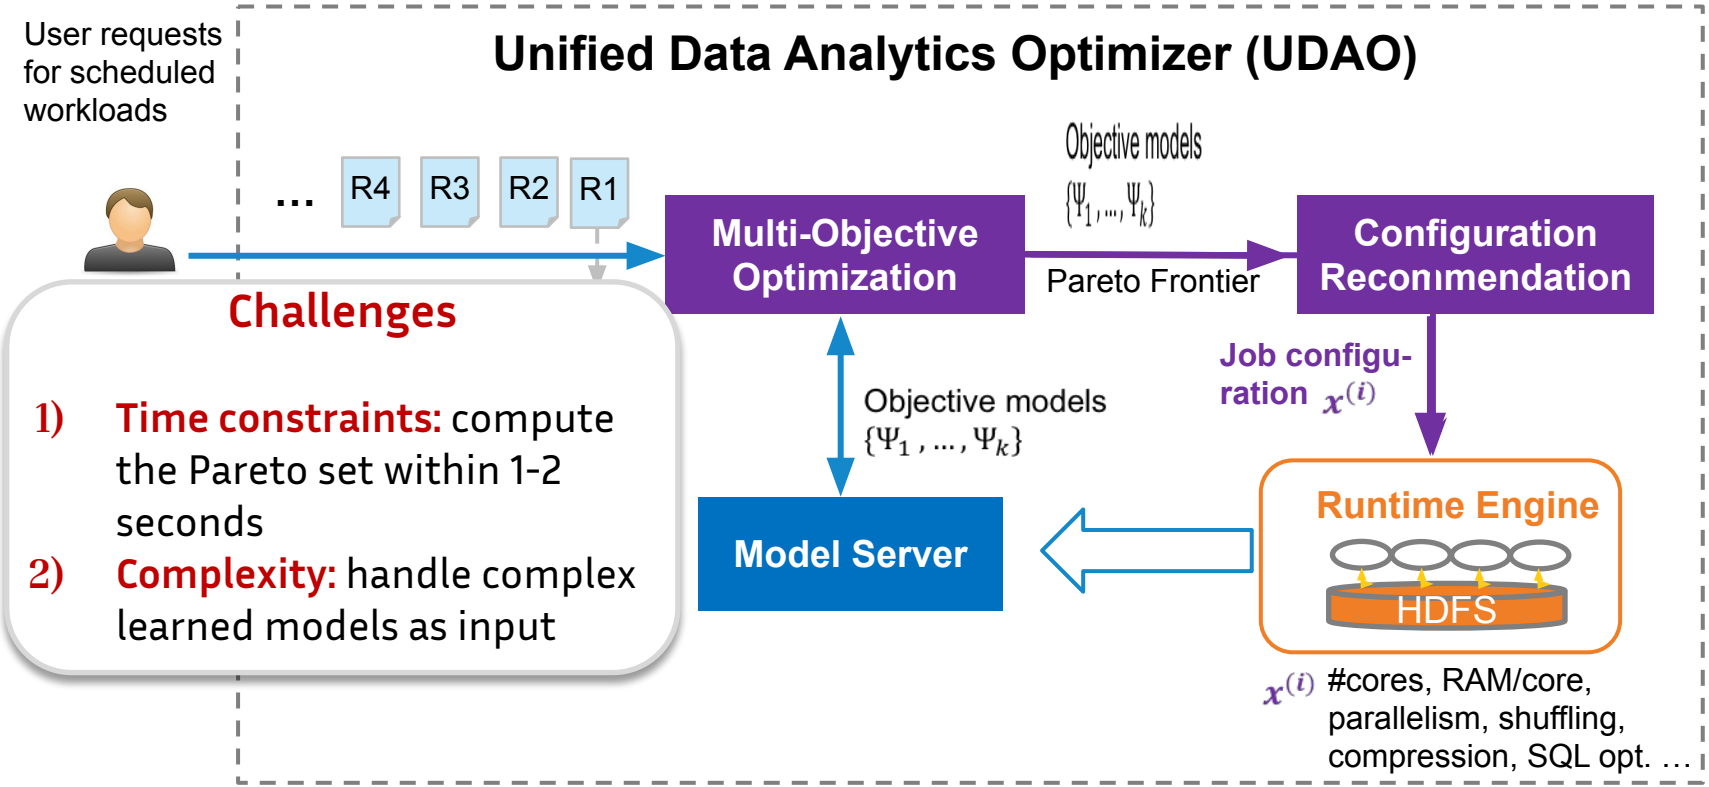
User requests for scheduled workloads
Unified Data Analytics Optimizer (UDAO)
...
R4
R3
R2
R1
Multi-Objective Optimization
Objective models
$\{\Psi_1, \ldots, \Psi_k\}$
Pareto Frontier
Configuration Recommendation
Job configu-ration $x^{(i)}$
Objective models
$\{\Psi_1, \ldots, \Psi_k\}$
Model Server
Runtime Engine
HDFS
$x^{(i)}$ #cores, RAM/core, parallelism, shuffling, compression, SQL opt. …
Challenges
1) Time constraints: compute the Pareto set within 1-2 seconds
2) Complexity: handle complex learned models as input

# Multi-Objective Optimization

Problem III.1. Multi-Objective Optimization (MOO).

$$\arg\min_{x} \; f(x) = \begin{matrix} F_1(x) = \; _1(x) \\ \dots \\ F_k(x) = \; _k(x) \end{matrix} \qquad (1)$$

$$s.t. \quad x \; 2 \; \hat{} \; \checkmark \; R^d$$
$$F_i^L \leq F_i(x) \leq F_i^U, \; i = 1, \dots, k$$

**Pareto optimality**

- Solution set
- Illustrates tradeoffs

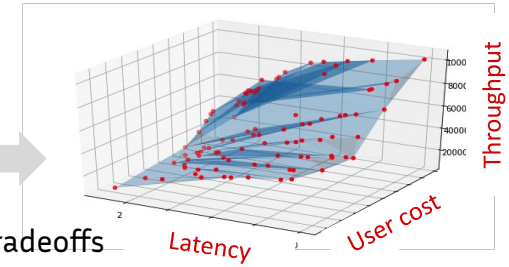


The **Progressive Frontier (PF)** approach transforms MOO into a series of *single-objective constrained optimization* (CO) problems, with each CO returning a Pareto point.

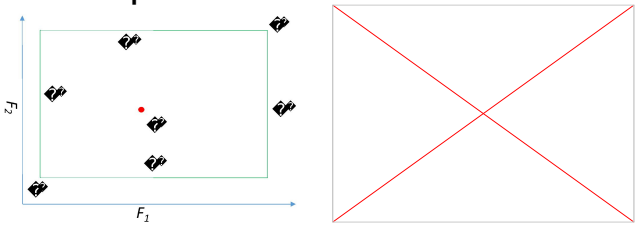


Proposition III.1. If we start the Iterative Middle Point Probes procedure from the initial Utopia and Nadir points, and let it terminate until the uncertain space becomes empty, then in the 2D case, our procedure guarantees to find all the Pareto points if they are finite. In high-dimensional cases, it is guaranteed to find a subset of Pareto optimal points.


❑ Fei Song, et al. Spark-based Cloud Data Analytics using Multi-Objective Optimization. IEEE International Conference on Data Engineering (ICDE), 2021.

# Progressive Frontier Algorithm

1. **An incremental, uncertainty aware algorithm**
   - Incremental: finding just one Pareto optimal point is expensive, so find more points incrementally
   - Uncertainty aware: the next Pareto point is returned from the most uncertain region
2. **Fast, approximate solver for each constrained optimization (CO) problem**
   - Finding a Pareto optimal point is expensive due to MINLP and complex learned models
   - Design a multi-objective gradient descent method using custom loss function
3. **Further use parallel computing to explore multiple subregions simultaneously**

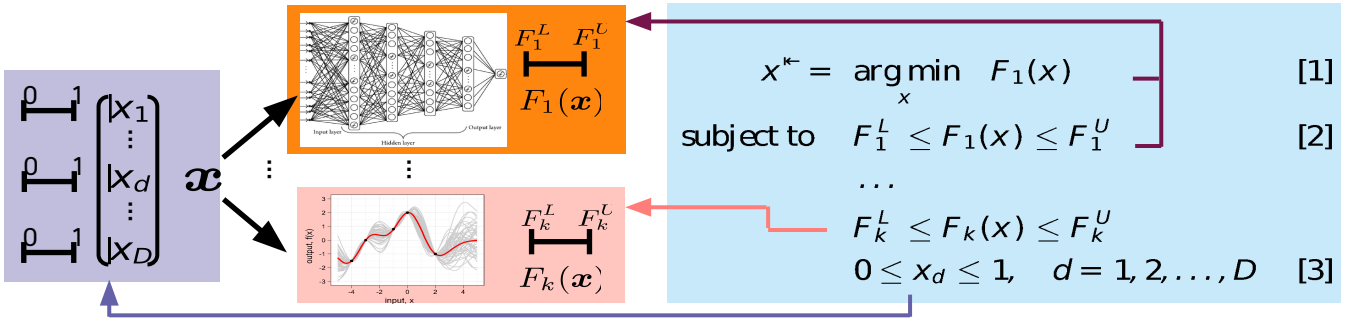

# Evaluation using TPCx-BB Benchmark (24K training & 258 test jobs)

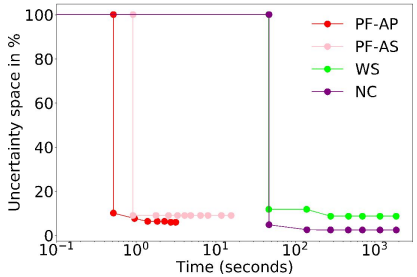


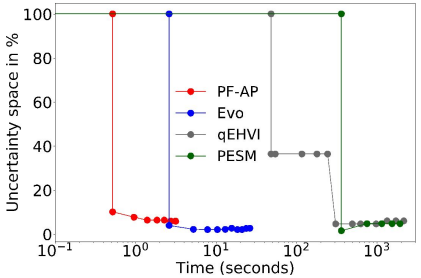


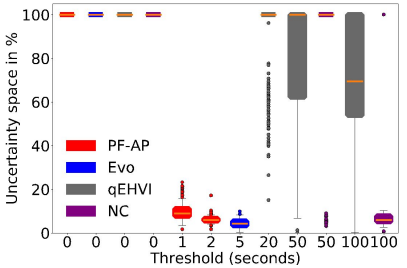


PF outperforms classic MOO, Evolutional, and Bayesian Optimization methods, by recommending from a Pareto set within 1-2 seconds

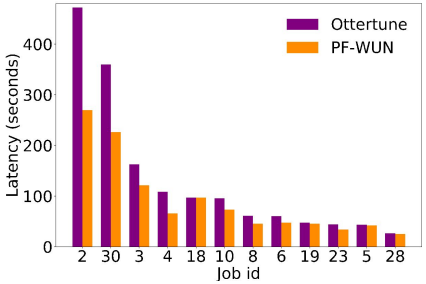


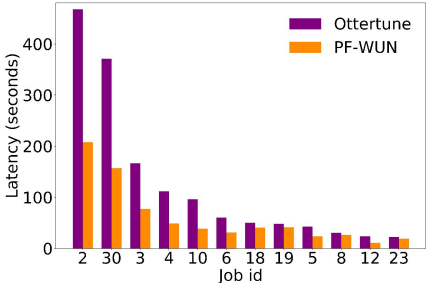


PF outperforms a STOA tuning system with 26%-49% reduction of running time while adapting to different user preferences across objectives

# Stateful data processing

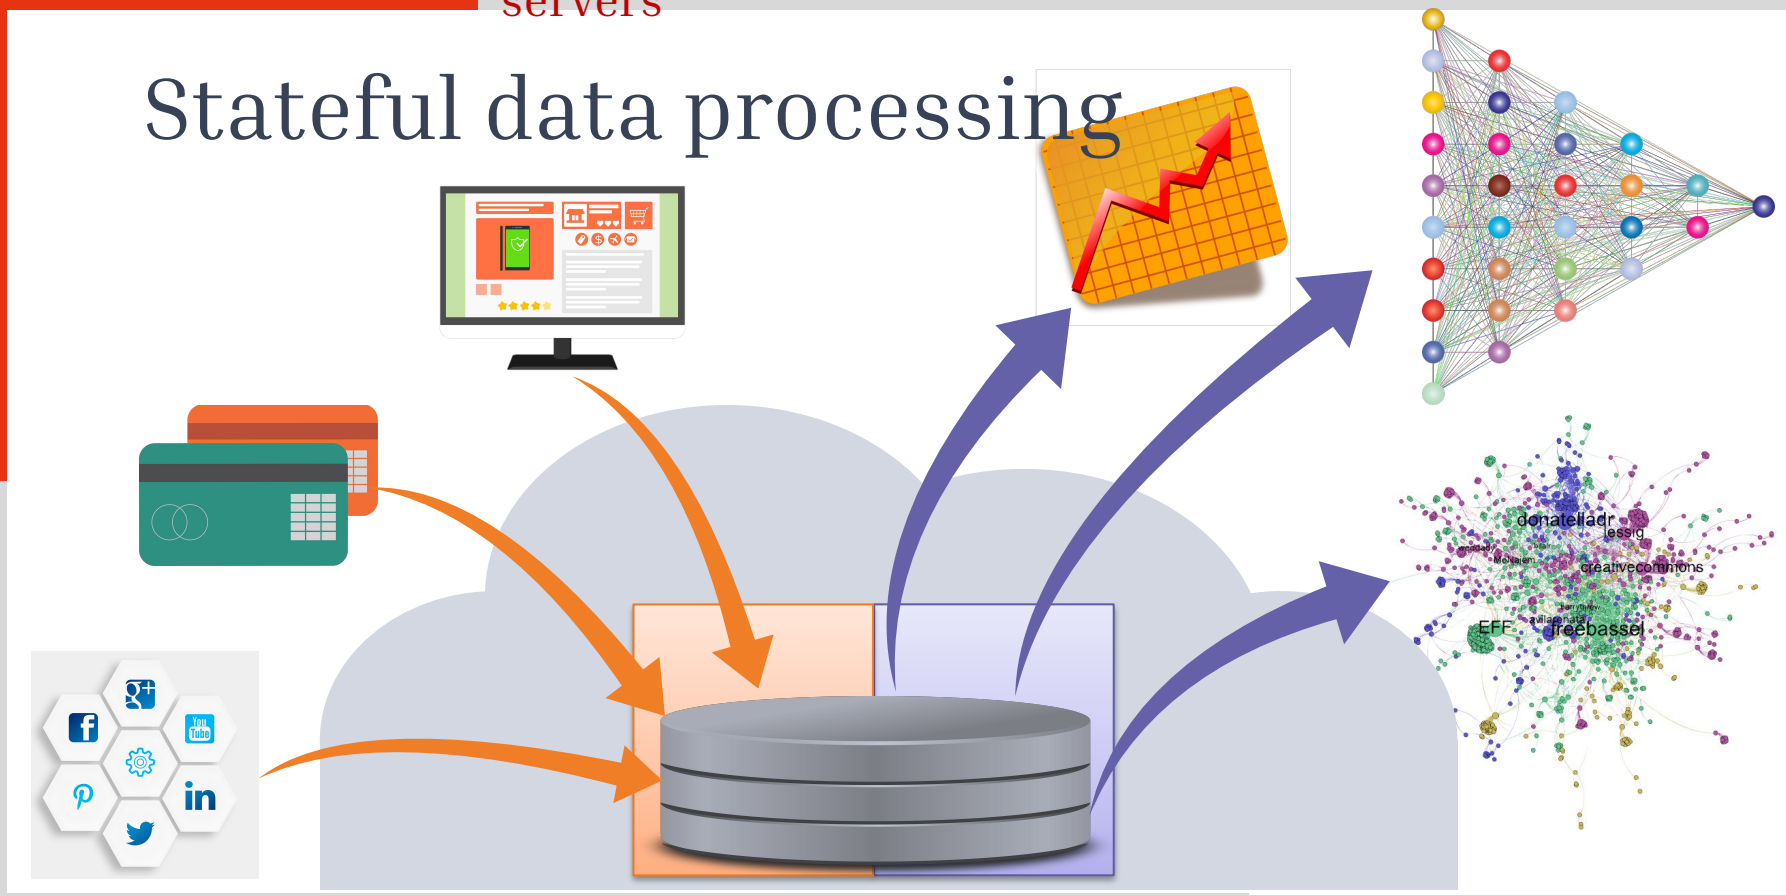

# Elastic scale-up virtualization stack

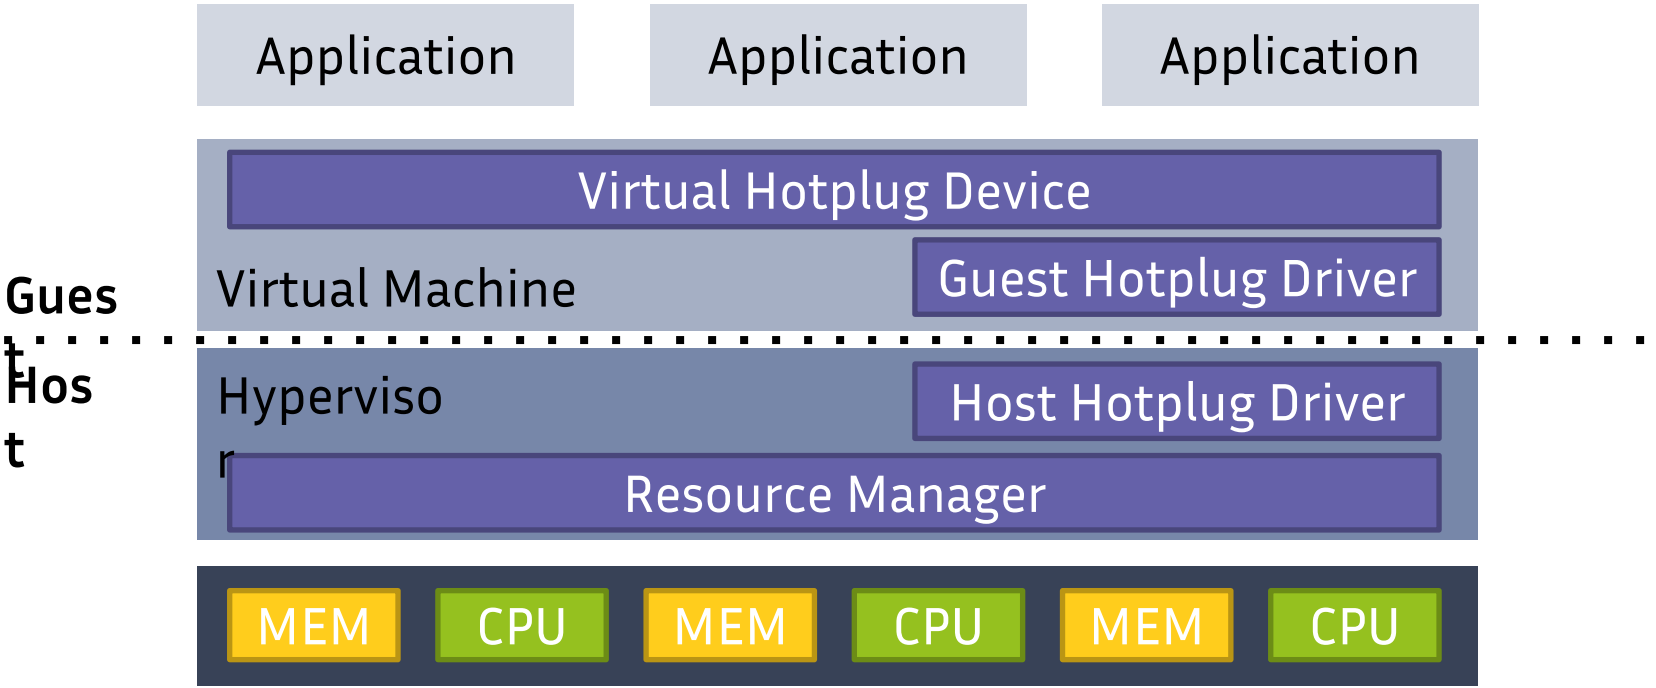

# Thank You!