



UNIVERSITY OF
CAMBRIDGE

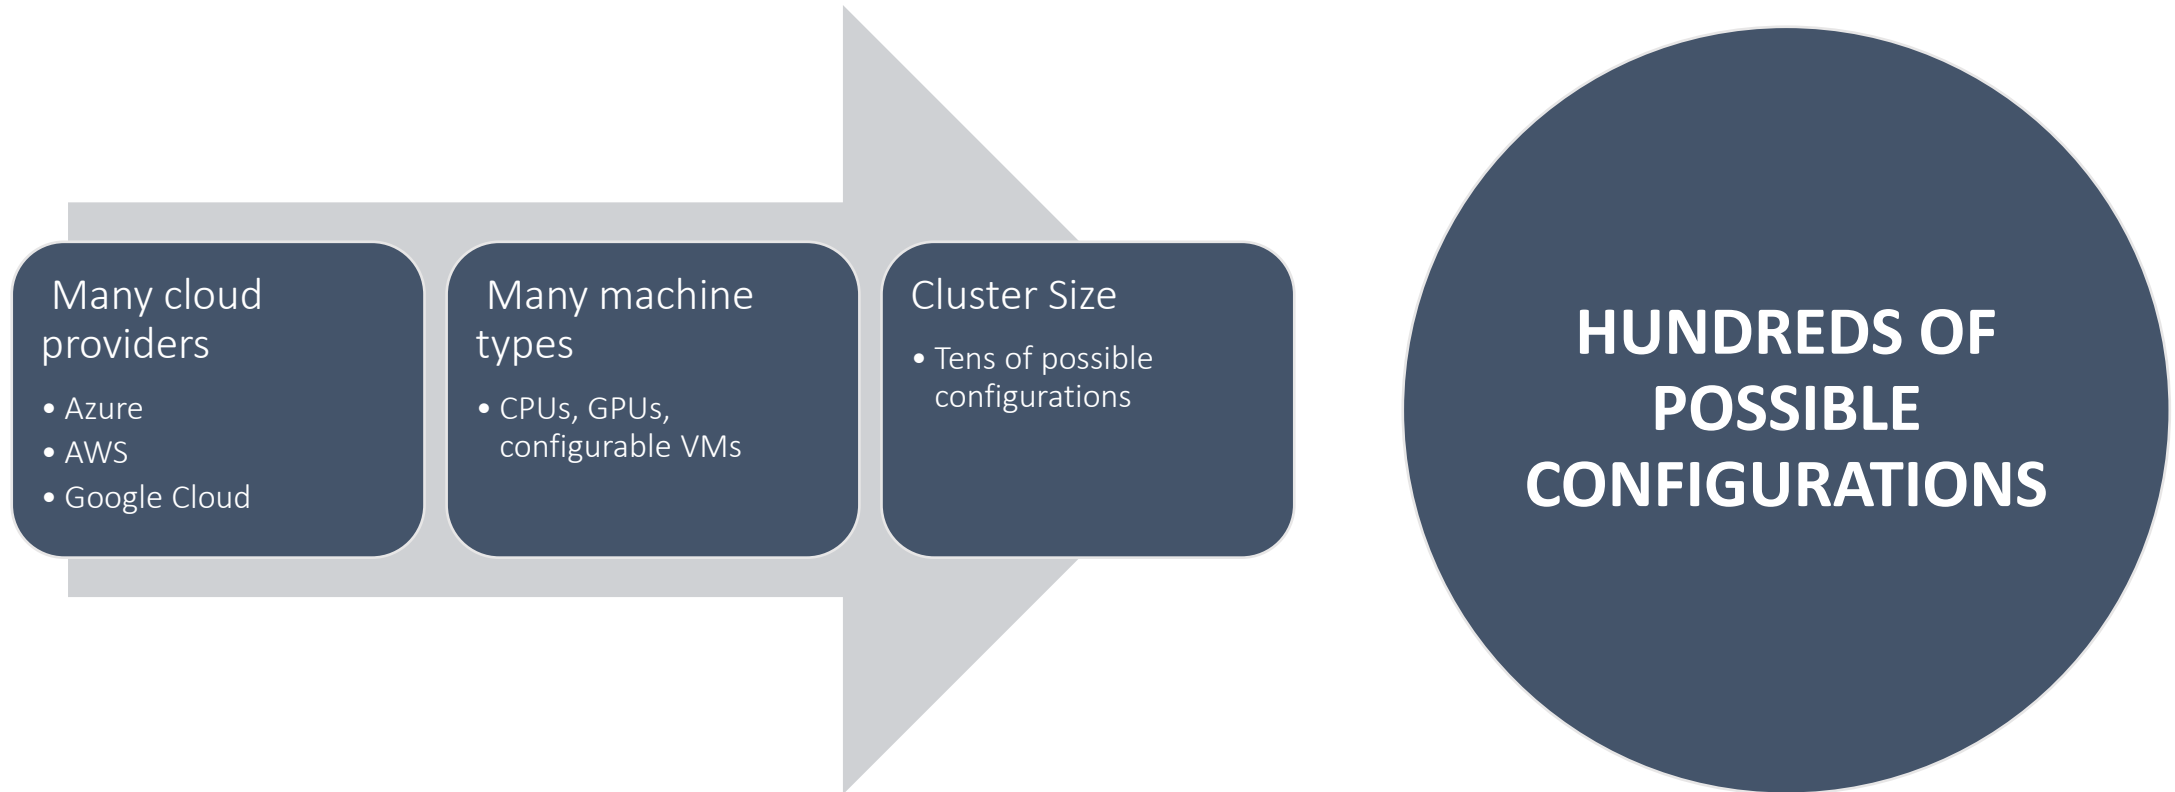
CherryPick: Adaptively Unearthing the Best Cloud Configurations for Big Data Analytics

Authors: Omid Alipourfard, *Yale University*; Hongqiang Harry Liu and Jianshu Chen, *Microsoft Research*; Shivaram Venkataraman, *University of California, Berkeley*; Minlan Yu, *Yale University*; Ming Zhang, *Alibaba Group*

Presented by Teodora



Big Data Analytics jobs are very popular "nowadays" (2017)



Good configuration = Better performance + Lower cost

$$\text{Same cost} \Rightarrow \frac{\text{worst running time}}{\text{best running time}} = 3$$

$$\text{Same performance} \Rightarrow \frac{\text{the most expensive configuration}}{\text{cheapest configuration}} = 12$$



Problem

Find **best cloud configuration** (minimizes **cost** for a given **performance threshold**) for a **recurring jobs**, given its **representative workload**.

40% of the jobs is cloud are **recurring**.

Challenges

Low Overhead

- Run few configuration

Coordinate
descent

Ernest

?

Adaptivity

- Works across all data apps

Try all
configs

High Accuracy

- Close to the global minima

Challenges

Low Overhead

- Run few configuration

Coordinate
descent

Ernest

Bayesian Optimization

Adaptivity


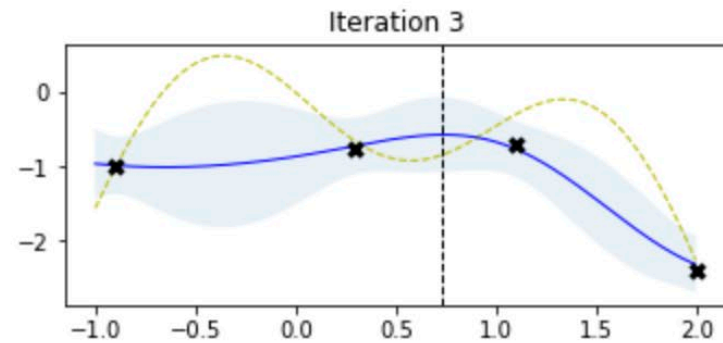

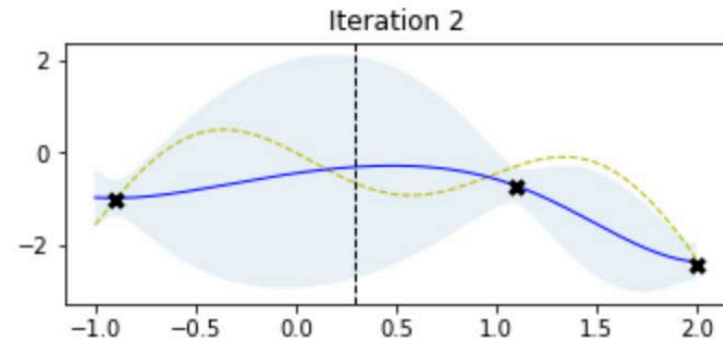

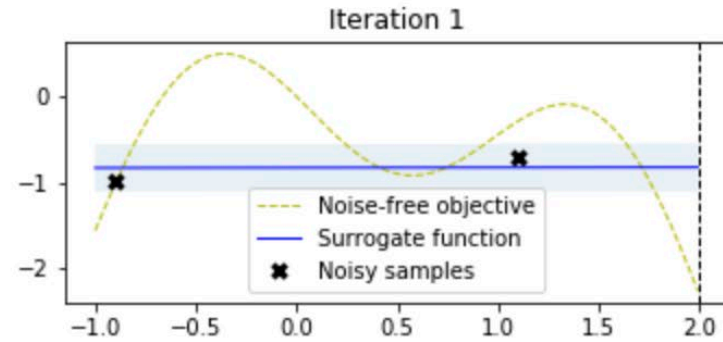
- Works across all data apps

Try all
configs

High Accuracy

- Close to the global minima


Bayesian Optimisation



Cloud has multiplicative noise


- cloud resources are shared, so if we run the same workload, same job, with same configuration, the running time and cost might not be the same





Price per unit of time Time

minimize $C(\vec{x}) = P(\vec{x}) \times T(\vec{x})$ $\tilde{T}(\vec{x}) = T(\vec{x})(1 + \epsilon_c)$
subject to $T(\vec{x}) \leq \mathcal{T}_{max}$ $\tilde{C}(\vec{x}) = C(\vec{x})(1 + \epsilon_c)$



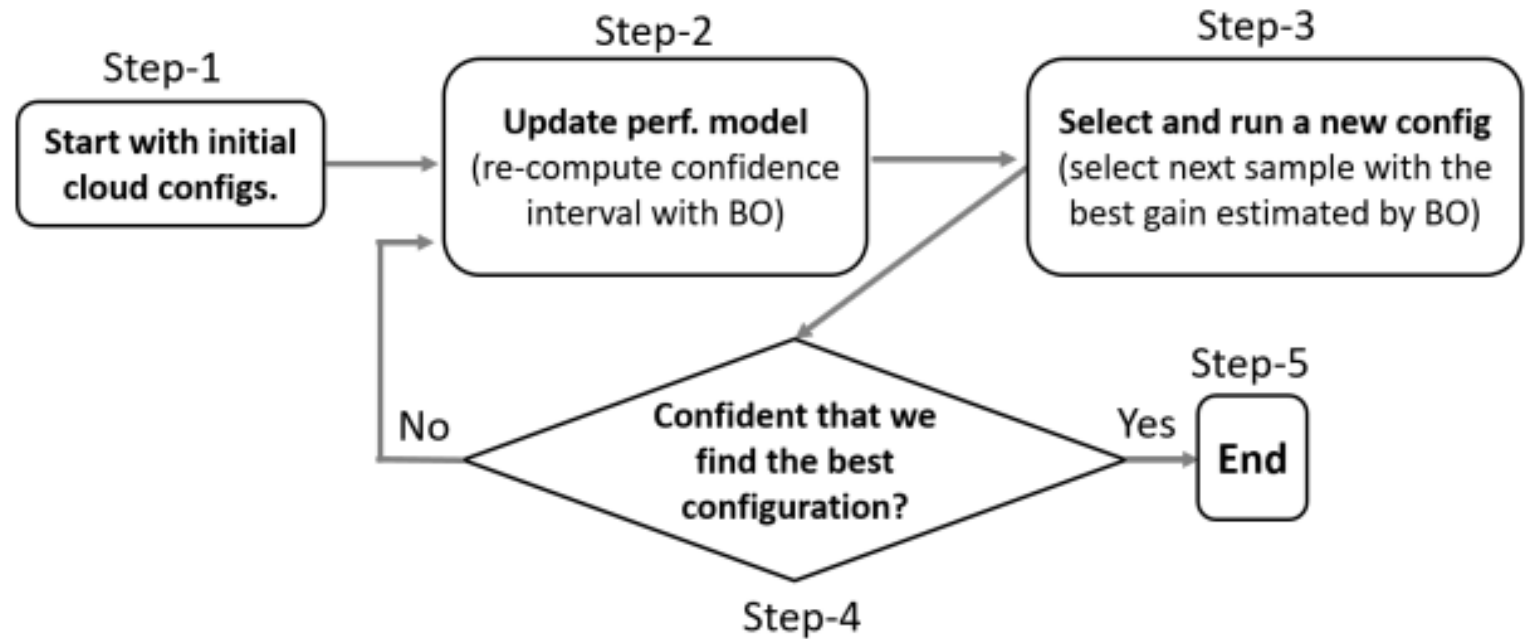
minimize $\log C(\vec{x}) = \log P(\vec{x}) + \log T(\vec{x})$

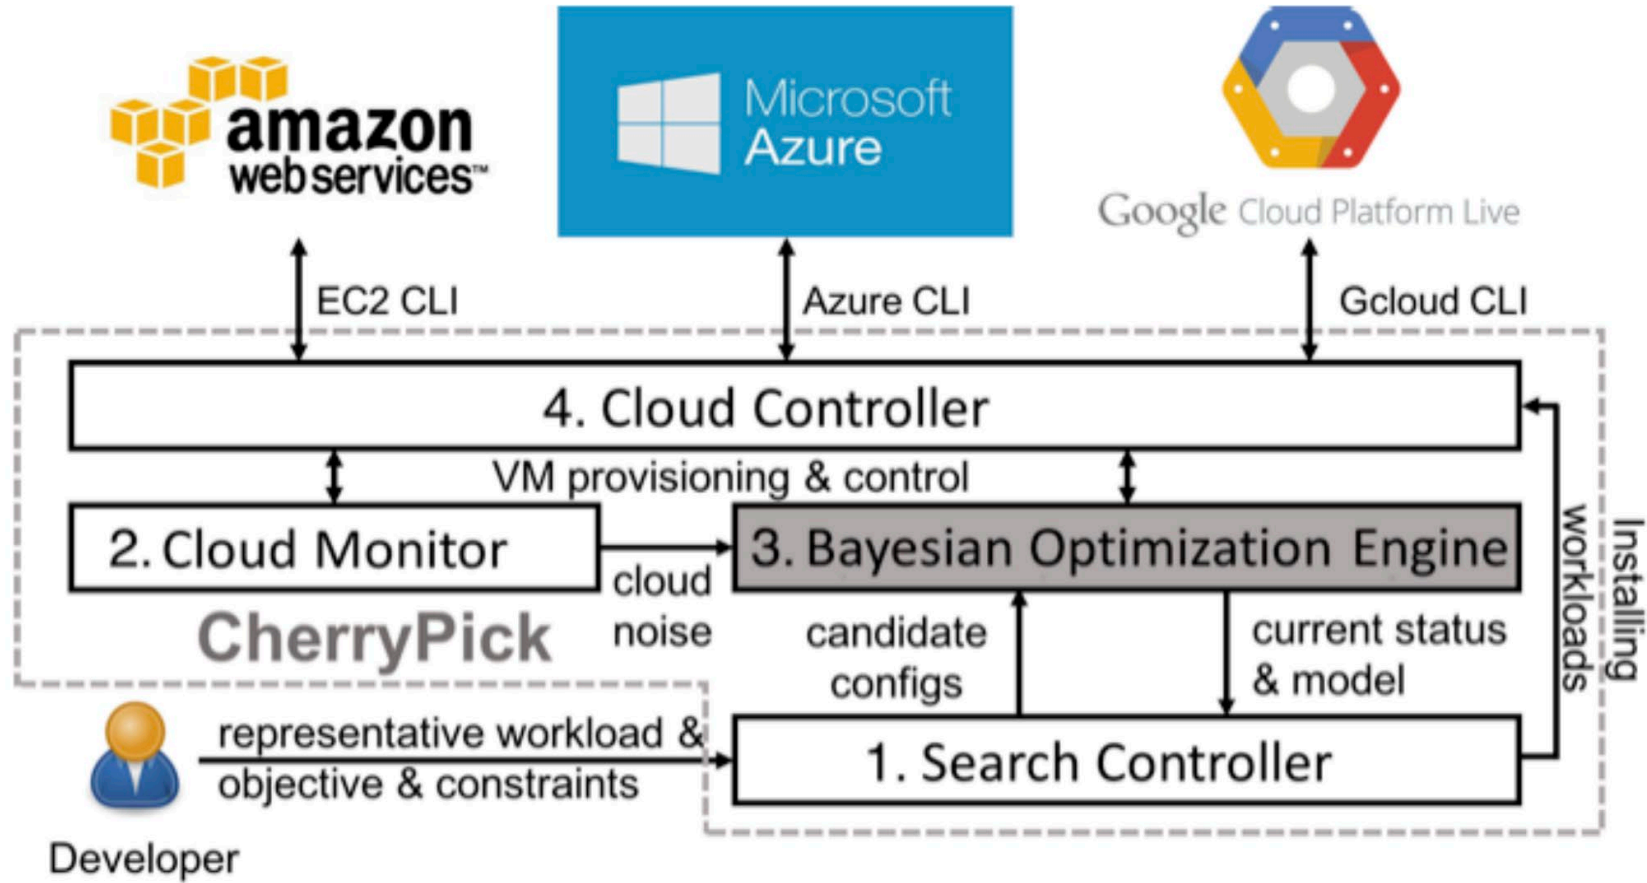
subject to $\log T(\vec{x}) \leq \log \mathcal{T}_{max}$

We use BO to minimize $\log C(\vec{x})$ instead of $C(\vec{x})$

$$\log \tilde{C}(\vec{x}) = \log C(\vec{x}) + \log (1 + \epsilon_c)$$

Workflow







Evaluation

Input:

- Five popular analytical jobs
- 66 reasonable configurations, of four families in Amazon EC2

Objective


- Minimize cost, within a performance threshold

Results:

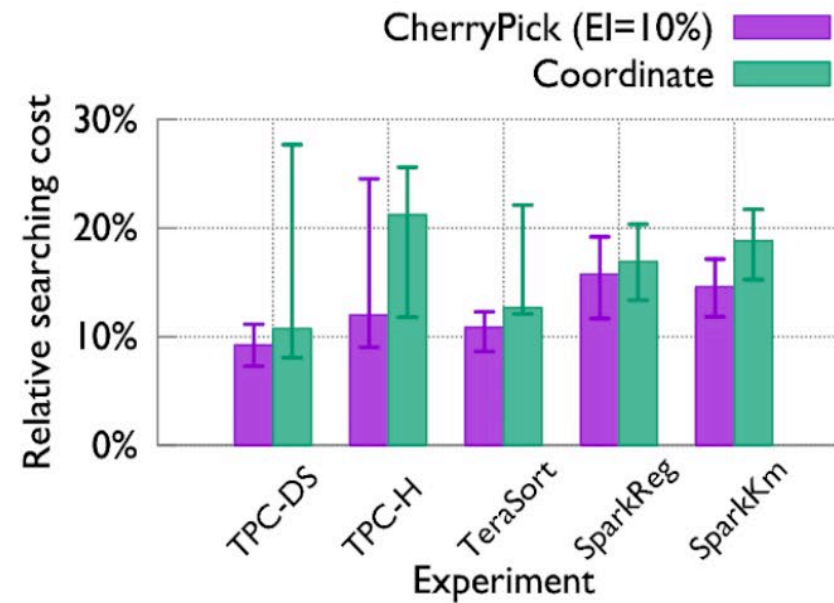
CheryPick = 45-90% to pick optimal solution, otherwise finds a solution within 5%

Alternatives = 75% more time to get to 45% overhead

Results



(a) Running cost



(b) Search cost

Figure 7: Comparing *CherryPick* with coordinate descent. The bars show 10th and 90th percentile.



Criticism

The algorithm is strong but the way they phrase it makes it seem weaker "45-90% chance to find the optimal".

Prior set to GP and **acquisition to Expected Improvement**, a bit restrictive? (ex. conjugate distribution for prior)

Representative workloads are needed. How can one get them?

Does it actually converge to a minima for noisy prior?



Questions?