

# Do Predictors for Resource Overcommitment Even Predict?

Georgia Christofidi

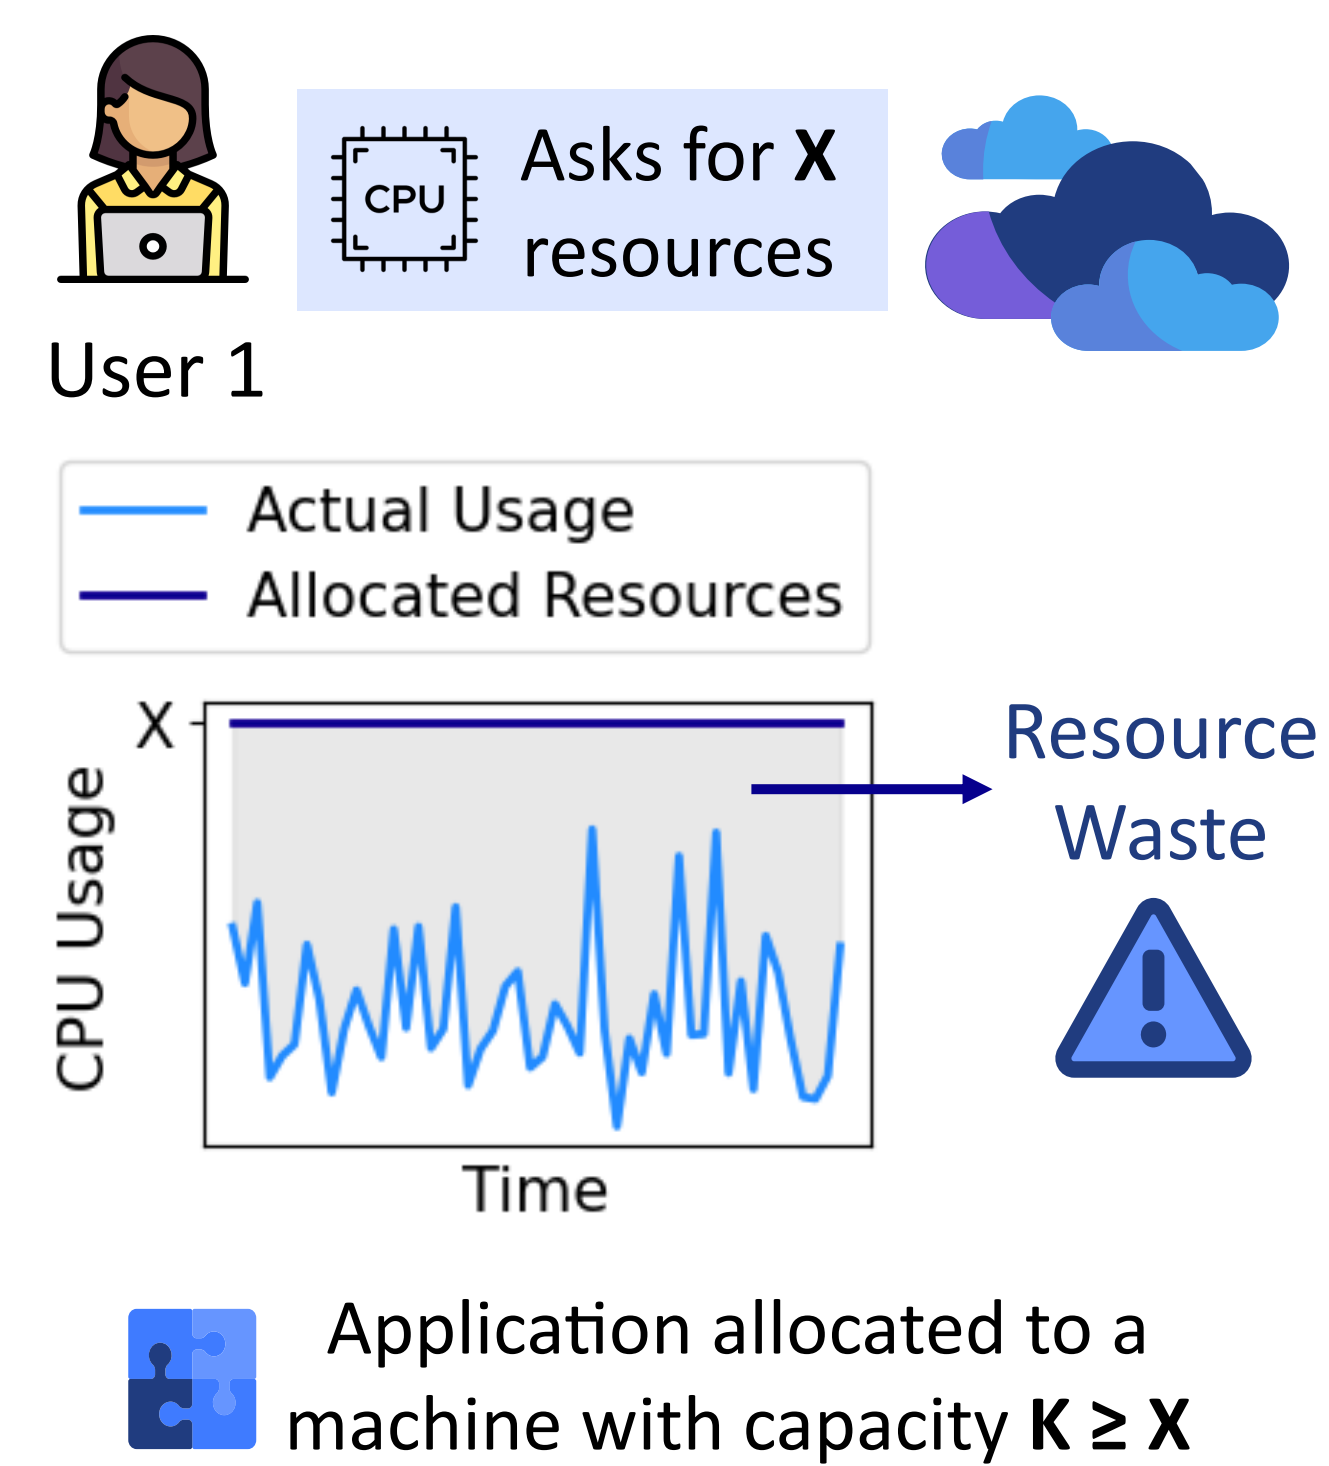
Thaleia Dimitra Doudali

IMDEA Software Institute, Madrid, Spain

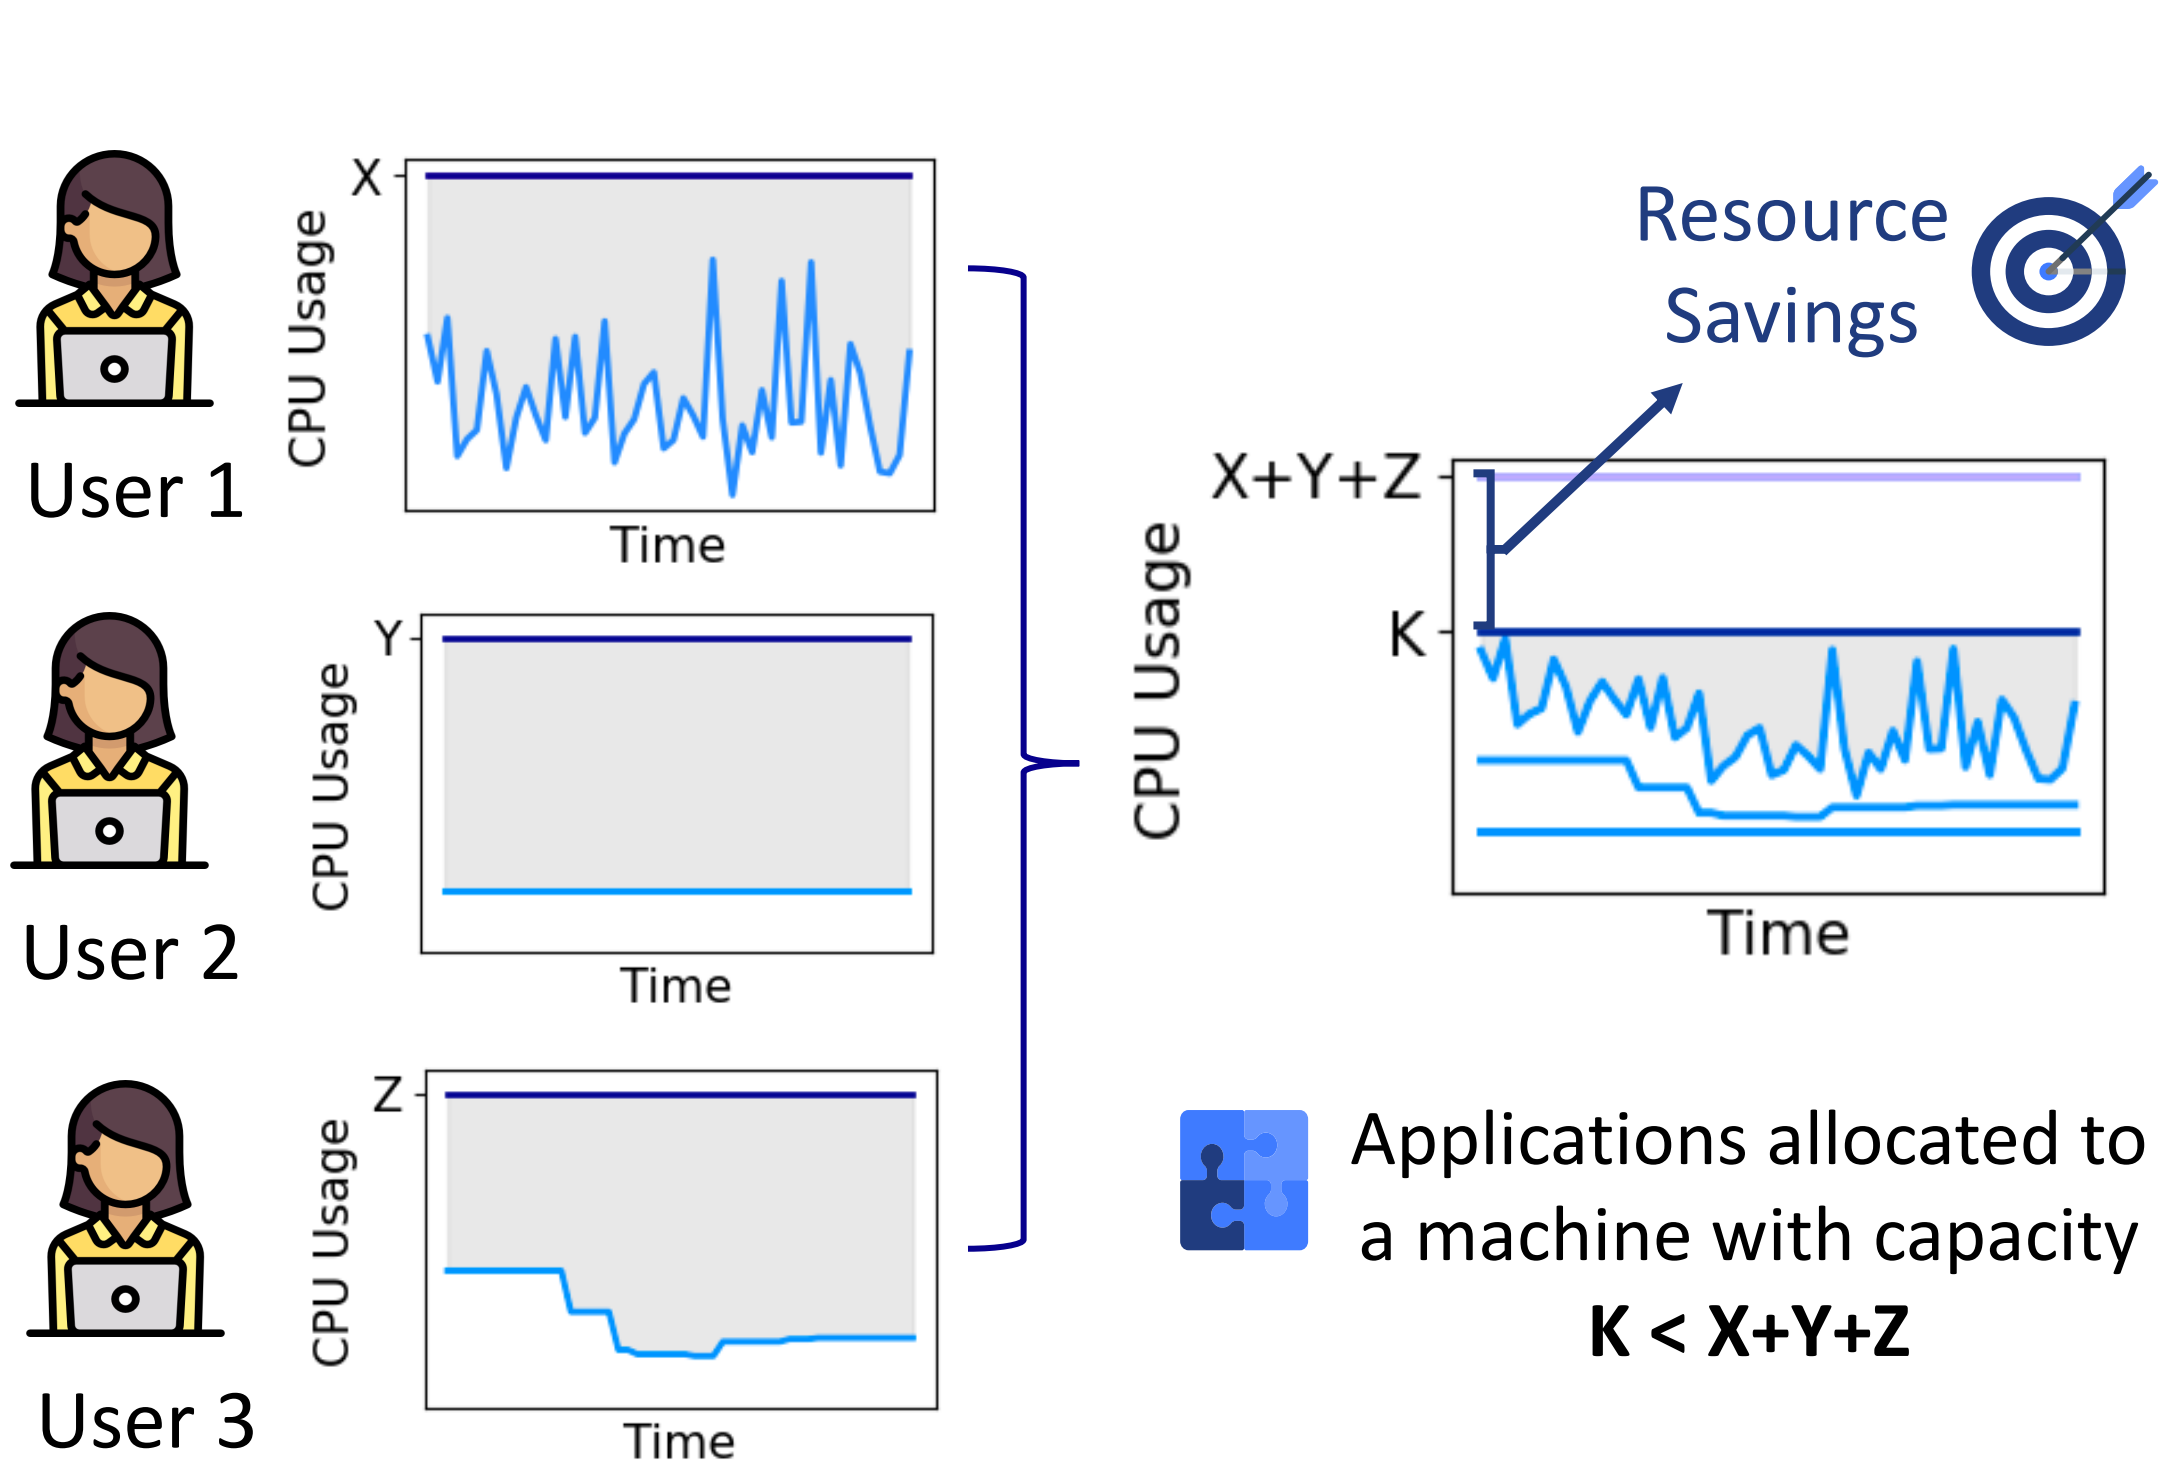


## 1. Problem Space

### 1. Traditional Resource Allocation



### 2. Resource Allocation with Overcommitment



### 3. Benefits & Risks of Resource Overcommitment

- ✓ Minimize resource waste
  - ✓ Improve resource and cost efficiency
  - ✗ Resource usage underestimations
- ↓ Application Performance  
↓ User Experience
- Competing for resources
- SLO Violations

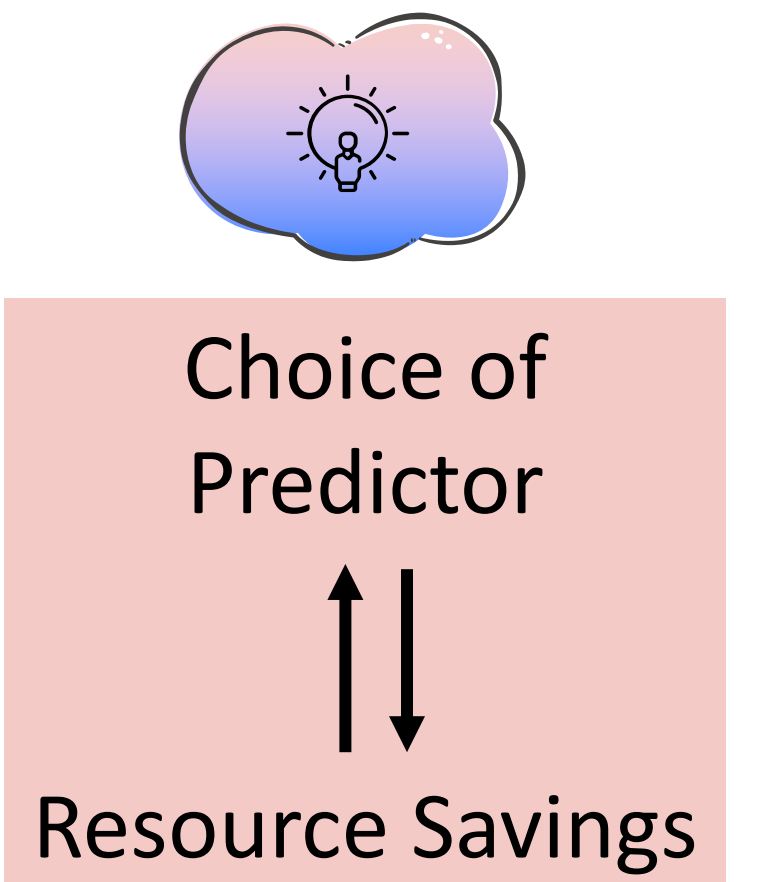

How to realize the benefits despite the risks?  
**Predict future resource usage!**

## 2. Existing Approaches

### Current Predictors

- Future Usage =
1. Borg:  $90\% * U$  (Google Cloud)
  2. Resource Central: sum of the k-th %-ile (Microsoft Azure)
  3. N-Sigma:  $U + N * std(U)$

4. TITTL:  $Max(1, 2, 3)$
- Conservative
- Often overestimates to avoid risk




## 3. Experimental Analysis

**Our Insight:**  
Predictions  $\gg$  Actual Usage  
Predictions  $>$  User-asked Resources (Limit) for 94% of times

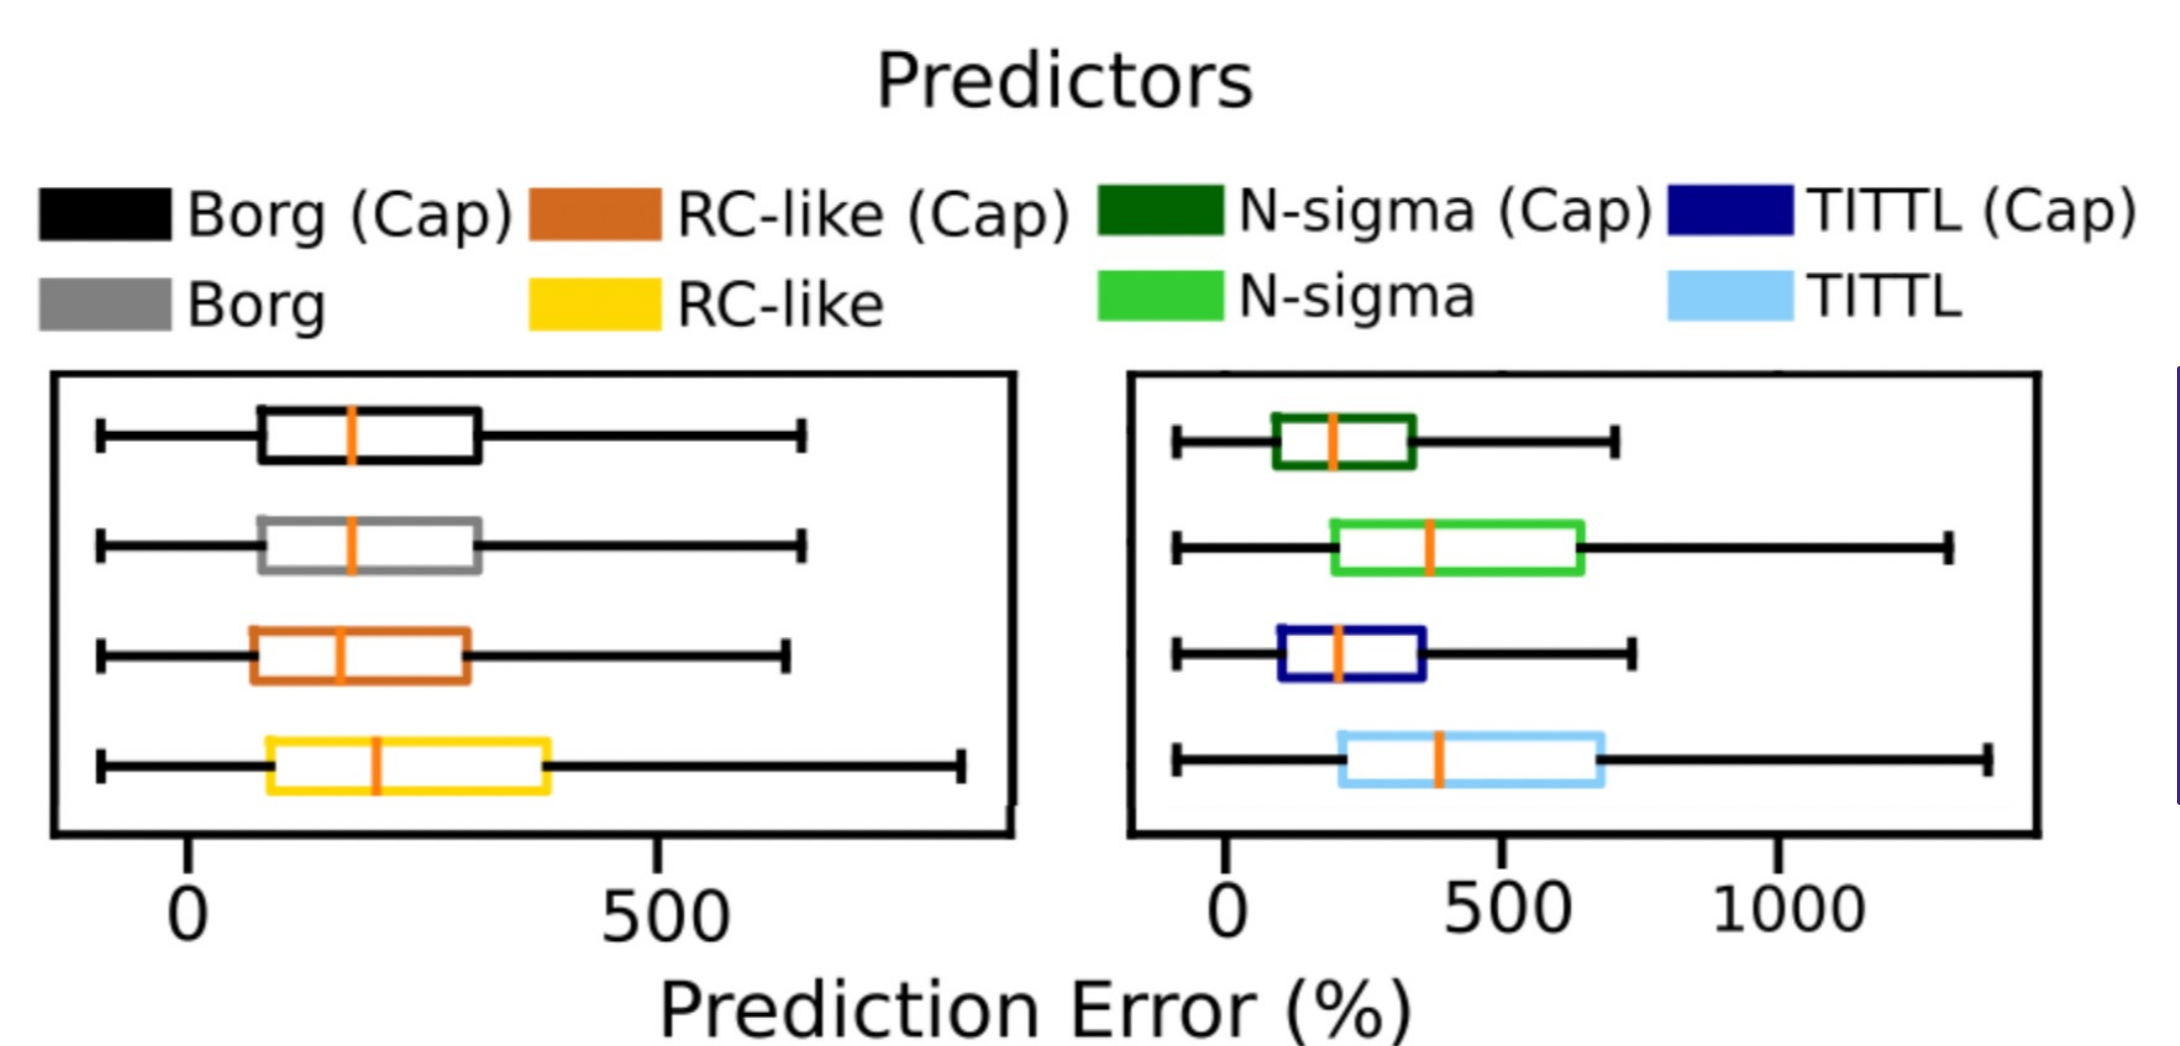
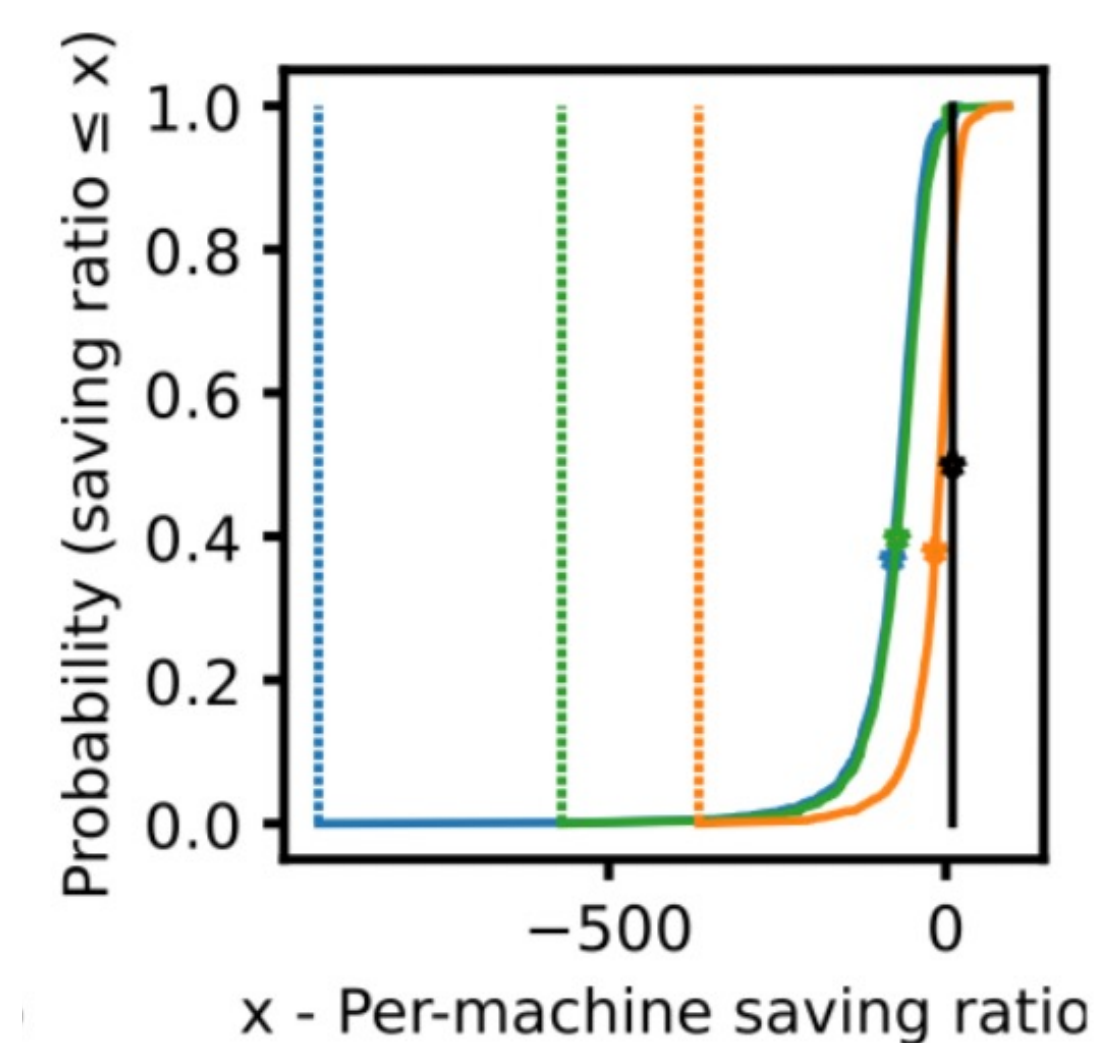
94% of times **no Overcommitment** **No Benefit!**

Does the predictor even predict at all?

Threshold: Limit If Prediction  $>$  Limit: then Prediction = Limit



Savings without the threshold



Current predictors severely **overestimate** resource usage.

## 4. Summary

Current predictors for resource overcommitment have low effectiveness.

- Resource usage overestimation
- Very little resource savings
- Overcommitment rarely achieved

A highly accurate predictor can enable **59% higher resource savings**, on average, compared to current approaches.

### Future Research

#### What predictive model to use?

- Lightweight solutions
- Practical Integration of ML & simple heuristics
- High Prediction Accuracy

We need a predictor that accurately predicts and closes the existing gap in attainable resource savings.

Scan for code and paper:

