USENIX ATC '24

Harmonizing Efficiency and Practicability: Optimizing Resource Utilization in Serverless Computing with Jiagu

Qingyuan Liu, Yanning Yang, Dong Du,

Yubin Xia, Ping Zhang, Jia Feng, James Larus, Haibo Chen

IPADS, SJTU; Huawei Cloud; EPFL

BACKGROUND & MOTIVATION

Serverless Computing is Popular

- **Popular cloud paradigm**
  - Users upload the code and platforms are responsible for dev/ops.

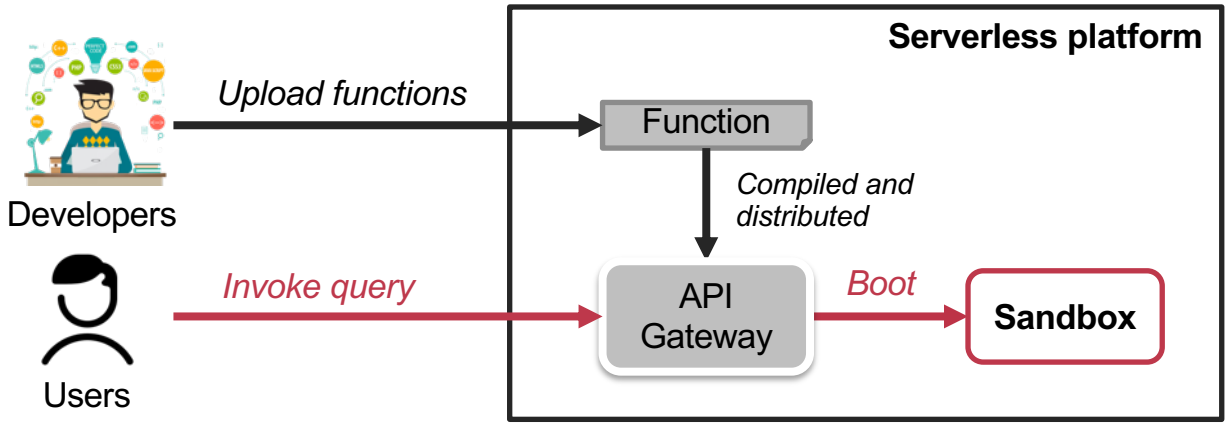

AWS Lambda Microsoft Azure Huawei FunctionGraph Knative

Key Components of Serverless Systems: Scheduler

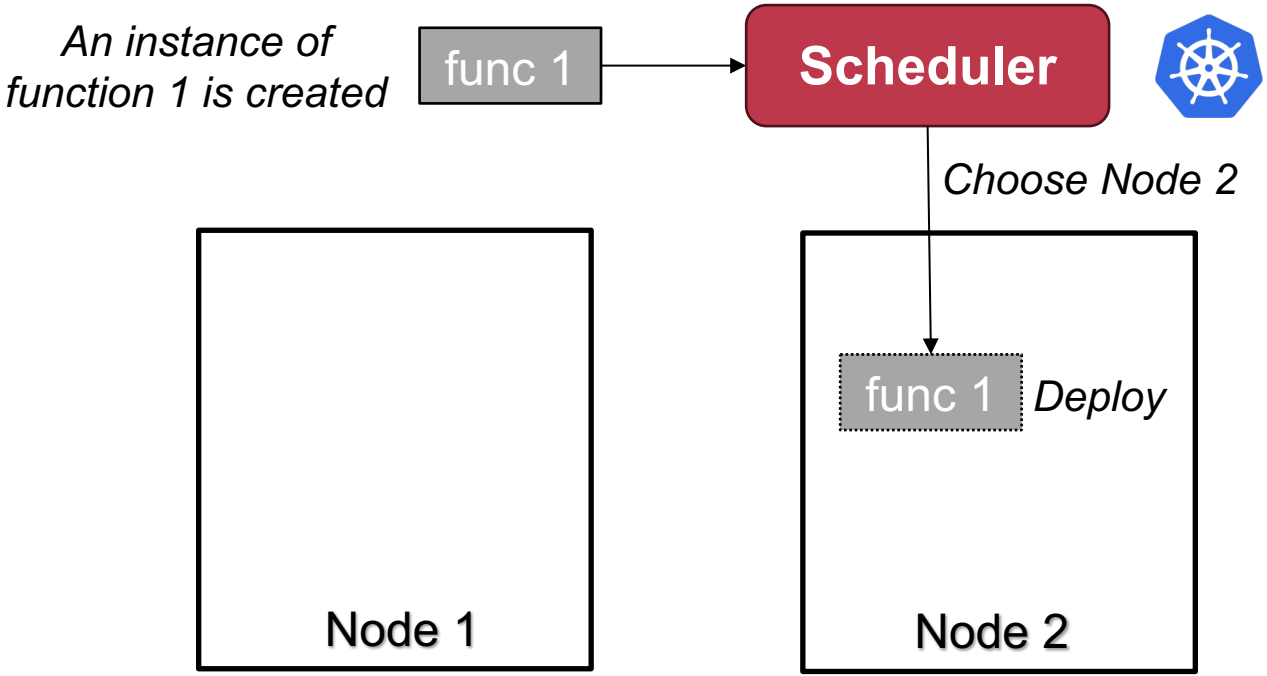

Scheduler: assigning each instance to a right server

Key Components of Serverless Systems: Router

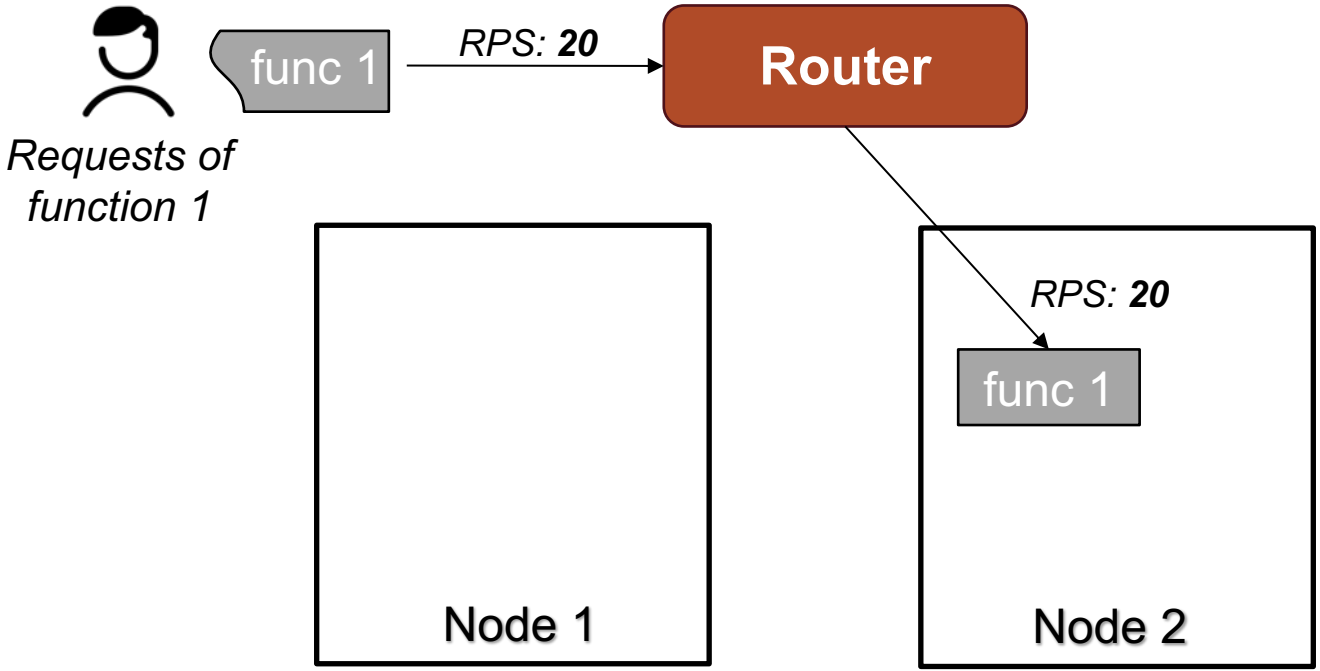

Router: distribute requests to specific instances

Key Components of Serverless Systems: Autoscaler

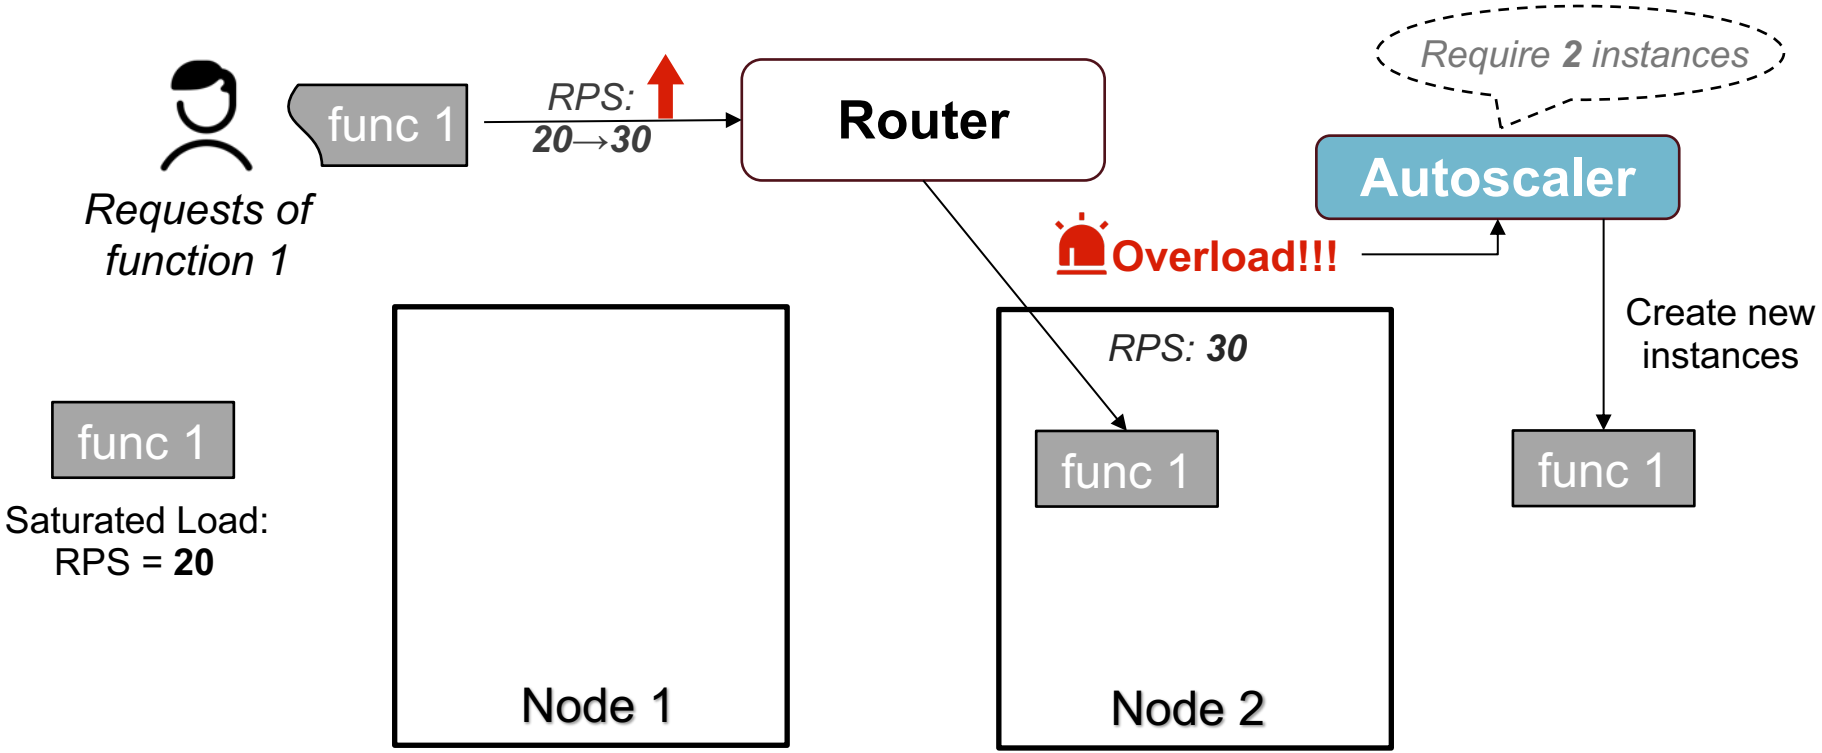

Autoscaler: scaling instances according to user loads

Key Components of Serverless Systems

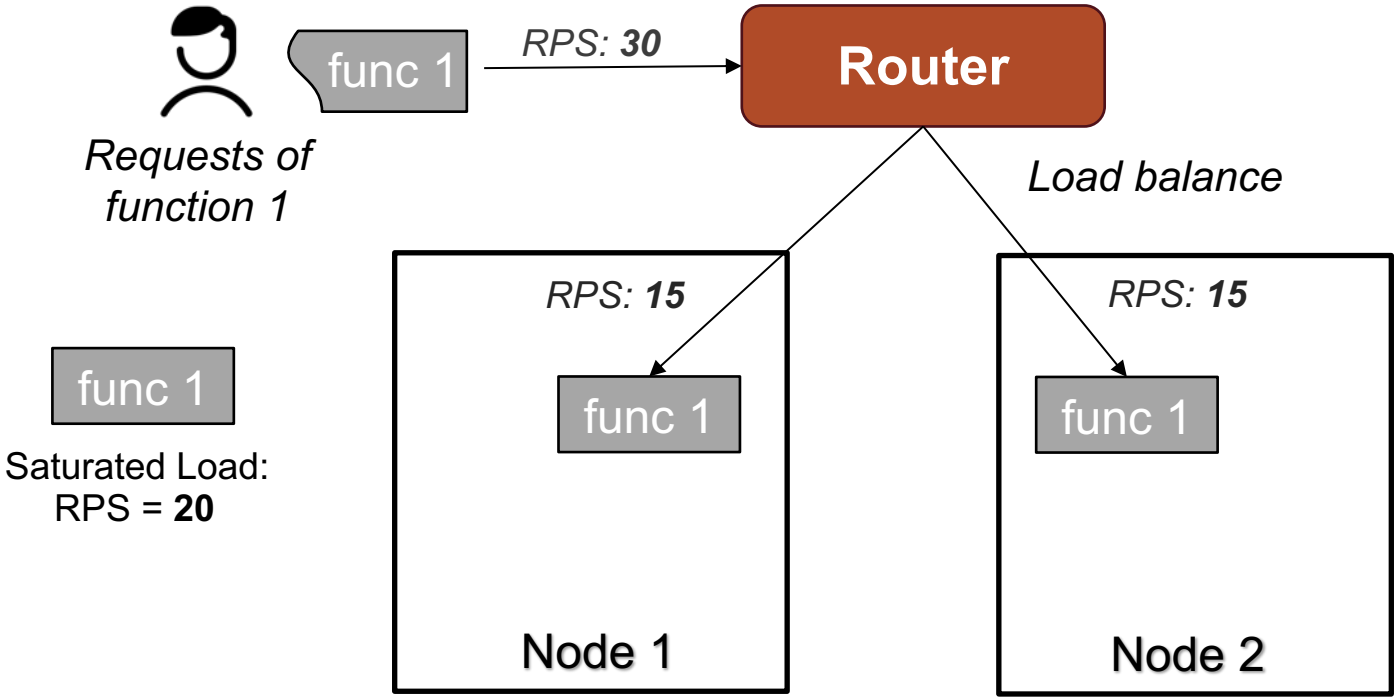

Key Components of Serverless Systems: Autoscaler

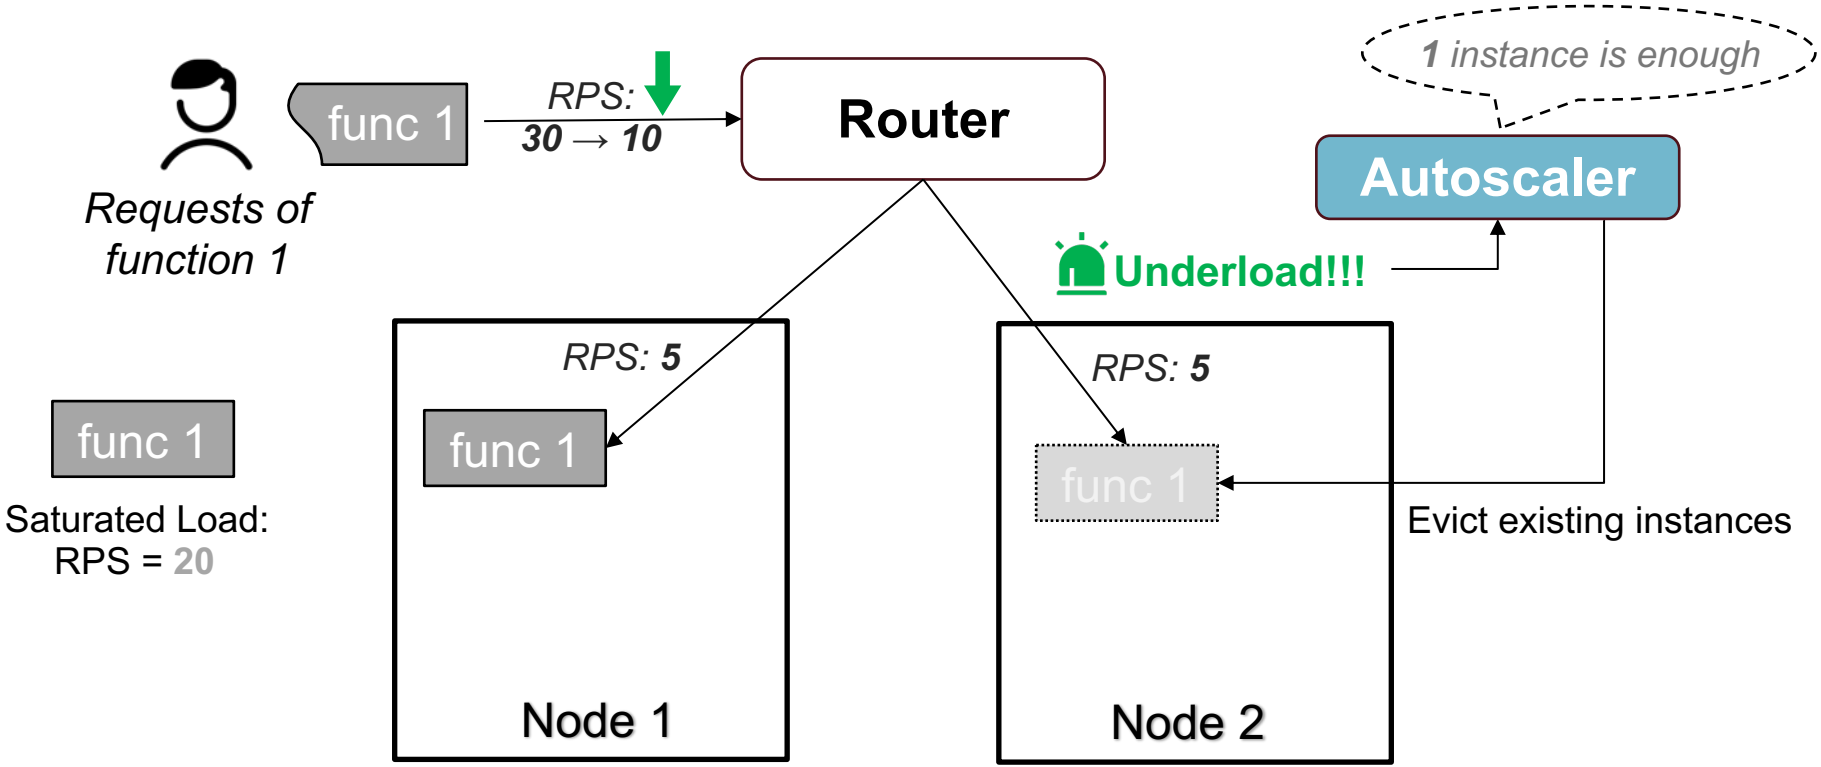

Autoscaler: scaling instances according to user loads

Low Resource Utilizations for Serverless Computing

Resources are under-utilized in serverless computing

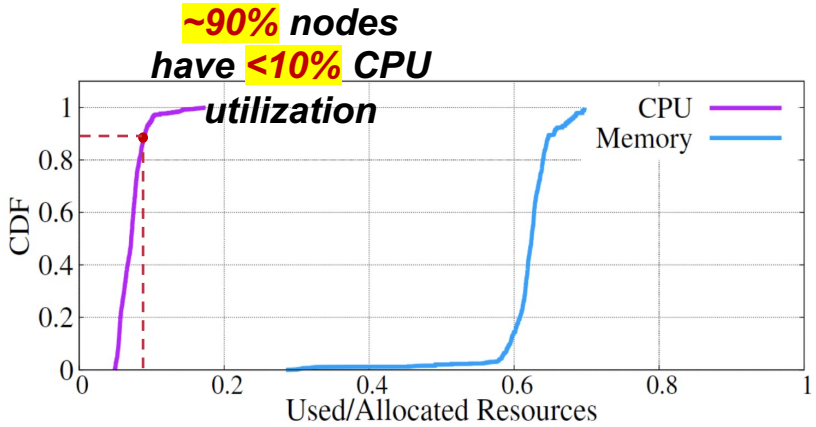

Cost saving is the mostly concerned issue for most organizations (62%)

"+1% resource util, Billons of $ saved"

Identify Two Causes of Resource Wastage

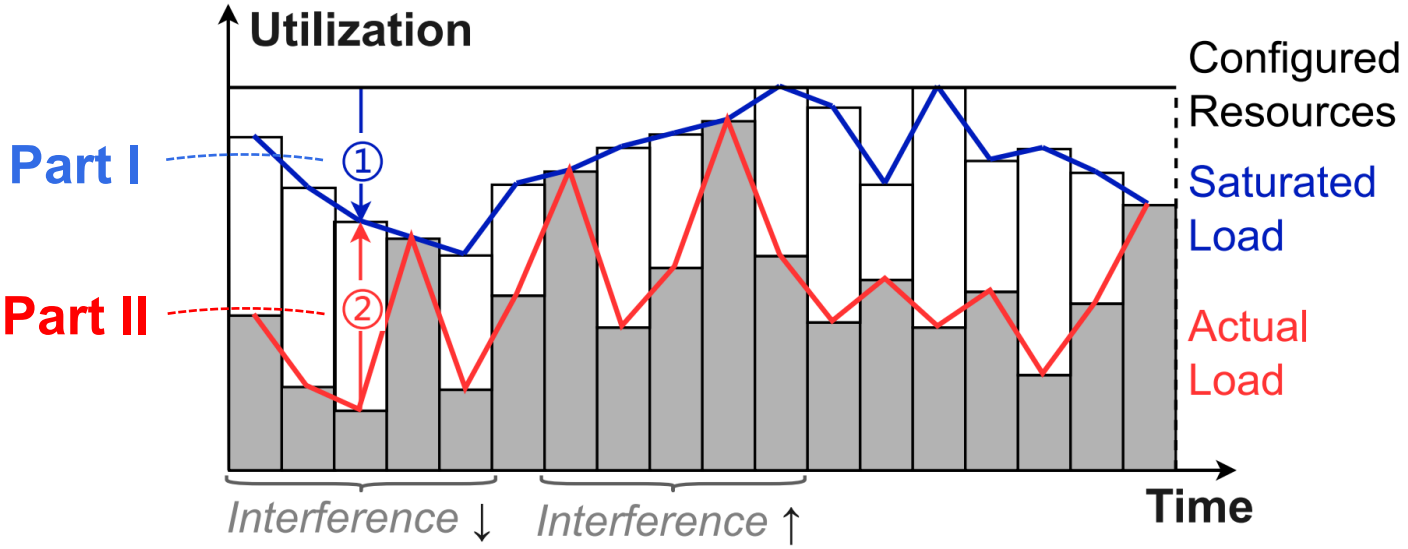

Part I: caused by **resource overprovisioning**

Part II: caused by **load overestimation**

Wastage Part I: Resource Overprovisioning

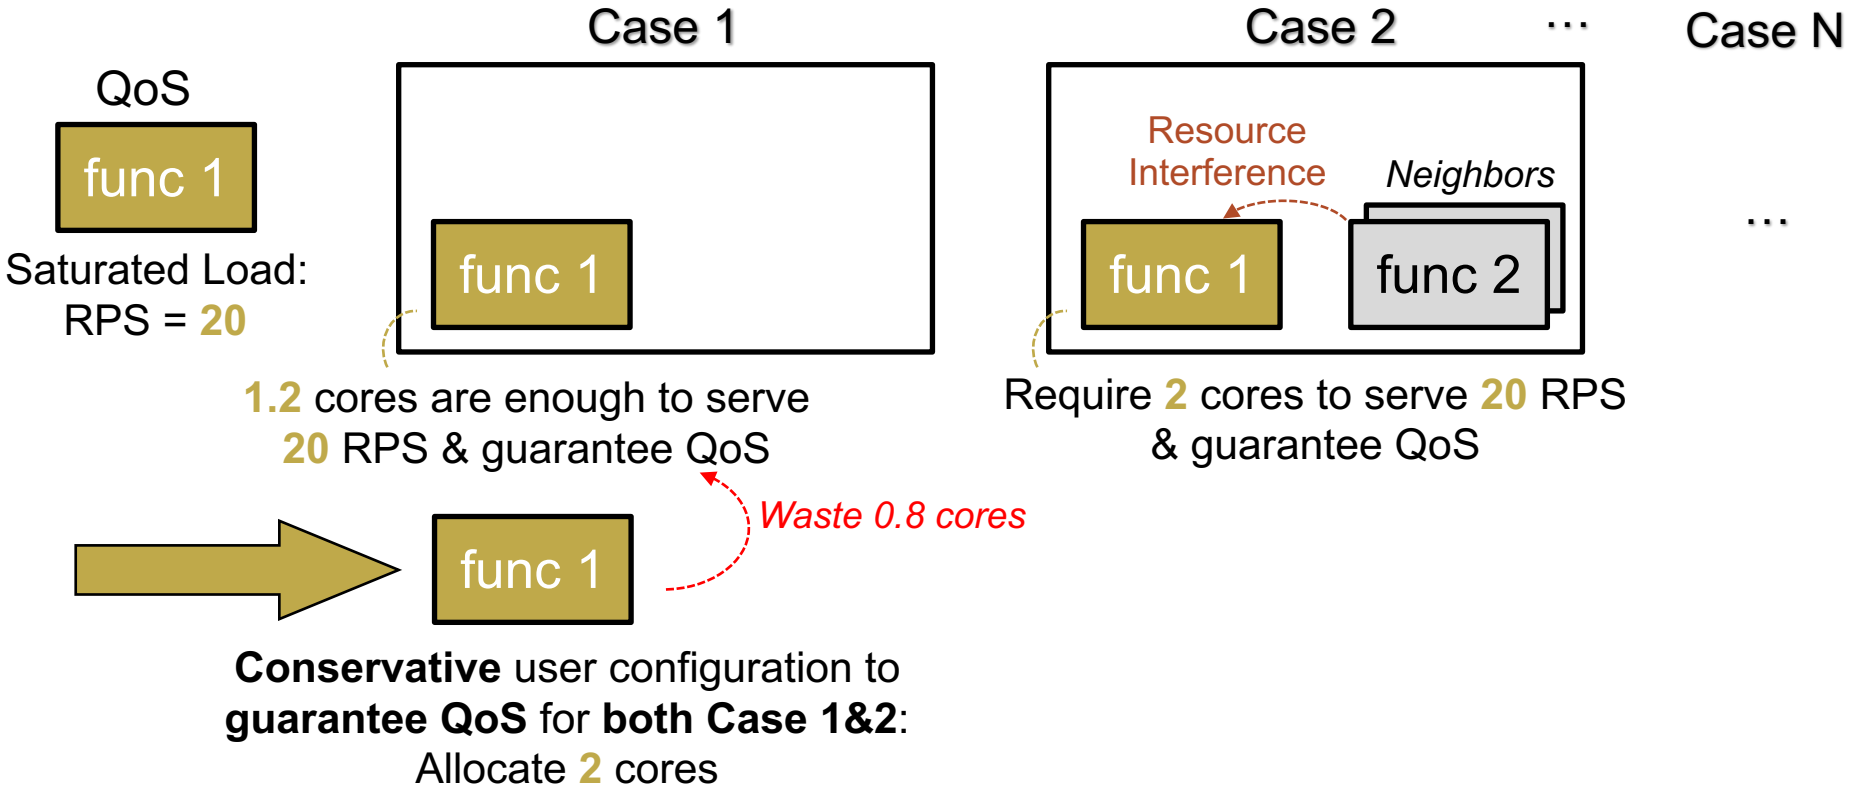

Wastage Part I: Resource Overprovisioning

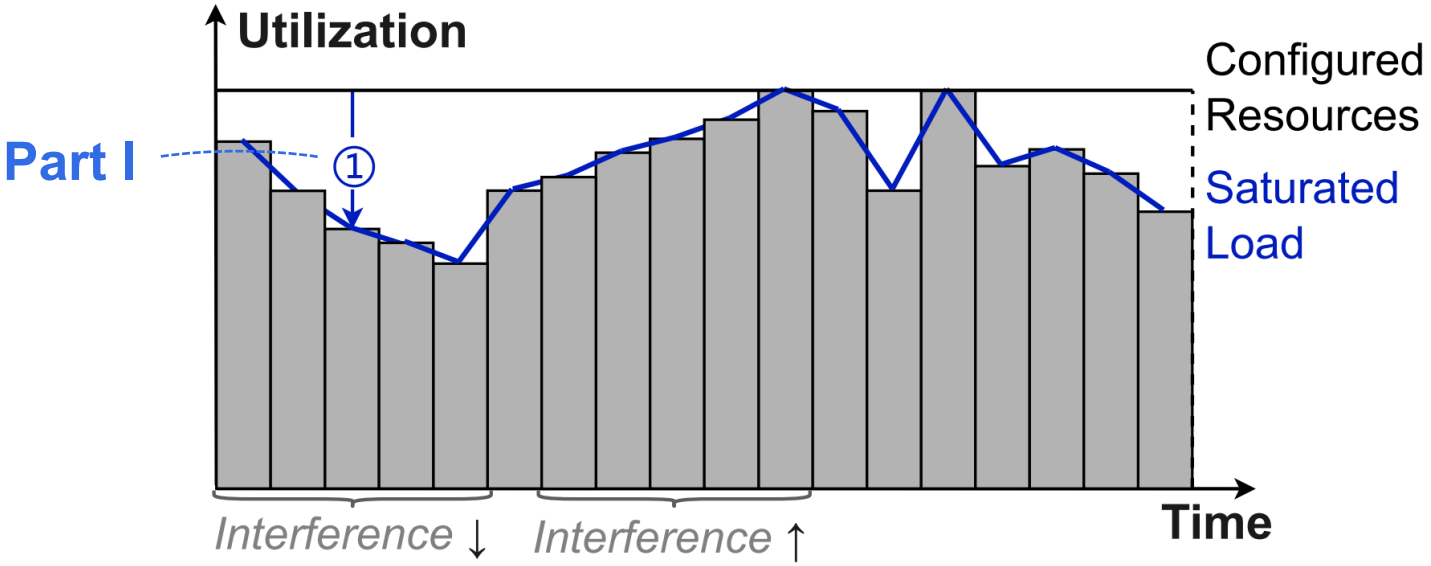

Part I: resources are **overprovisioned** even for *saturated* instances

Wastage Part II: Load Overestimation

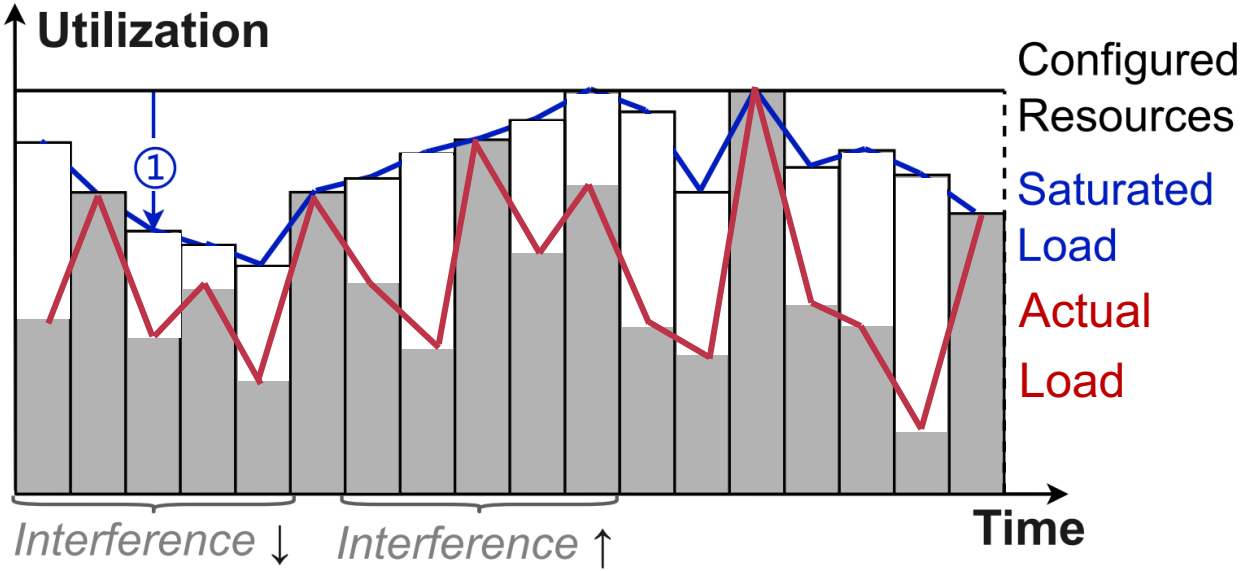

Unpredictable load fluctuation causes

load overestimation

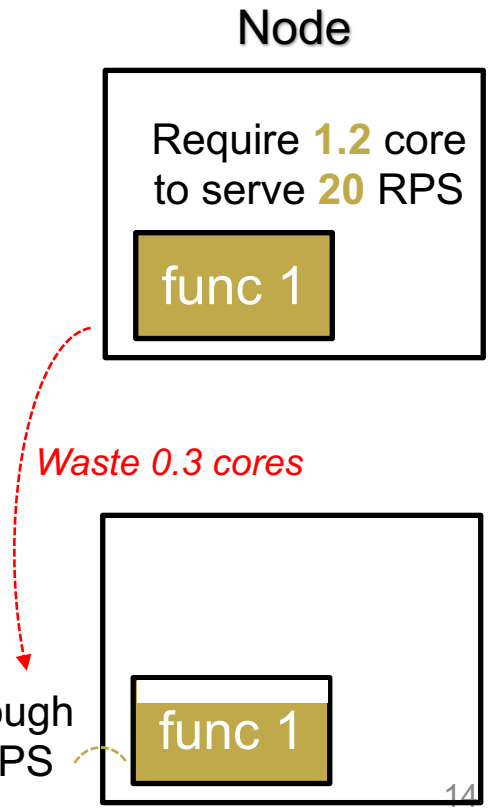

Wastage Part II: Load Overestimation

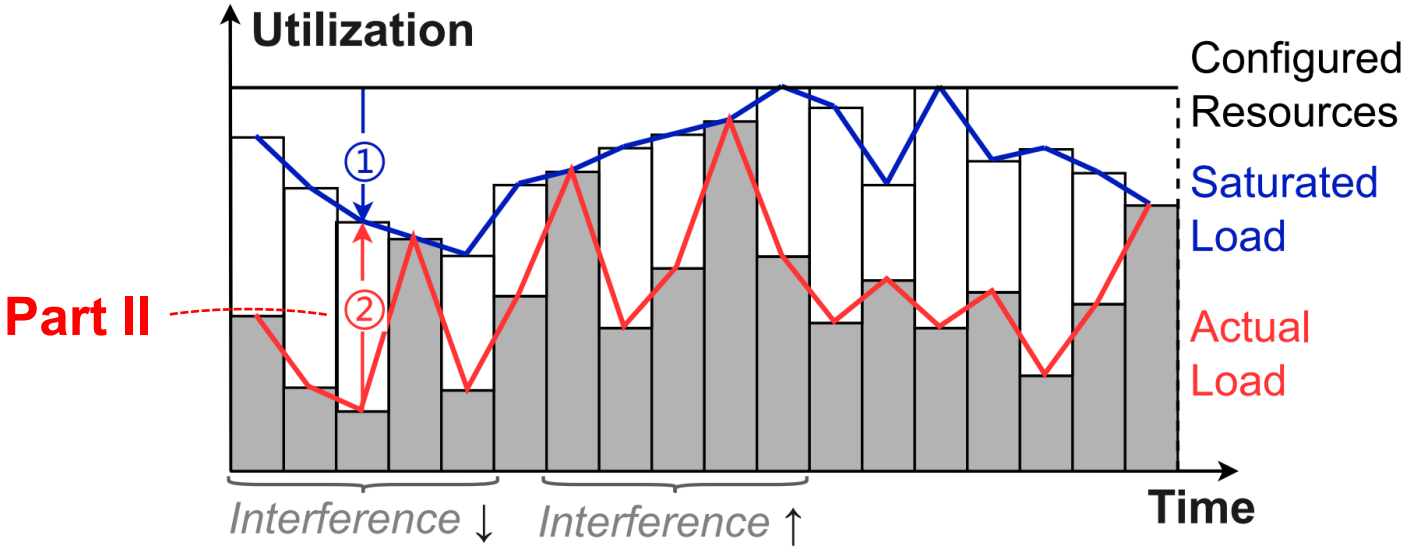

Part II: resources are **overestimated** due to load fluctuation

Challenges to Mitigate the Two Parts of Wastage

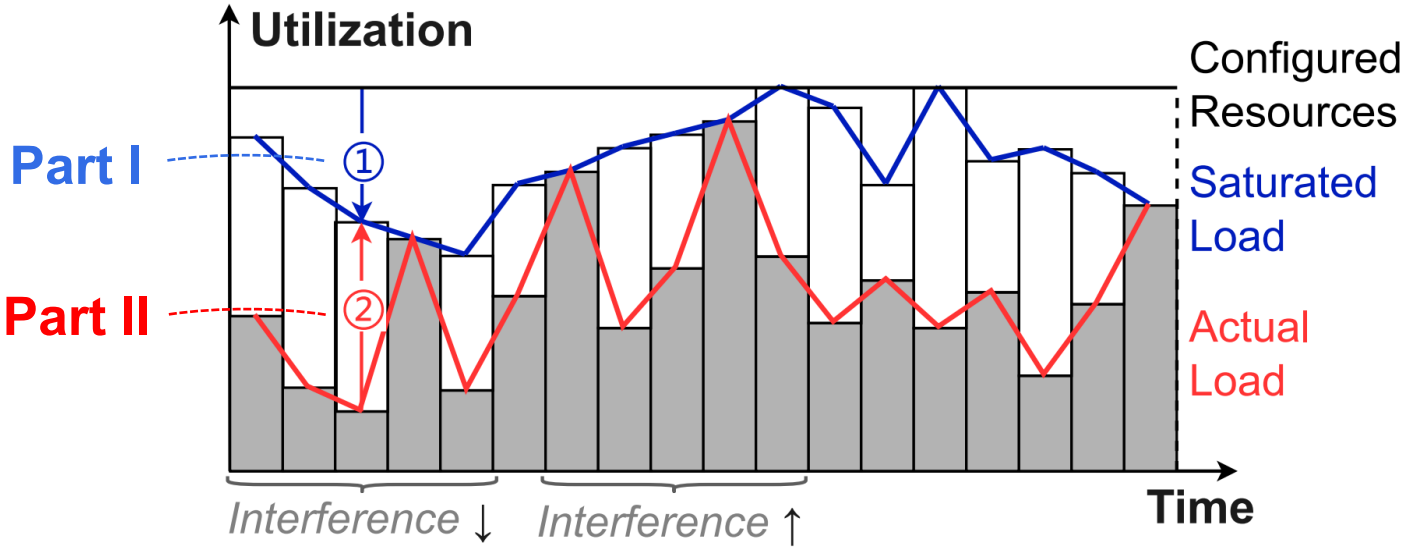

Challenges for prior methods:

Tradeoffs between high effectiveness & low cost

Mitigate Wastage Part I: Overcommitment

Overcommitment: increase deployment density

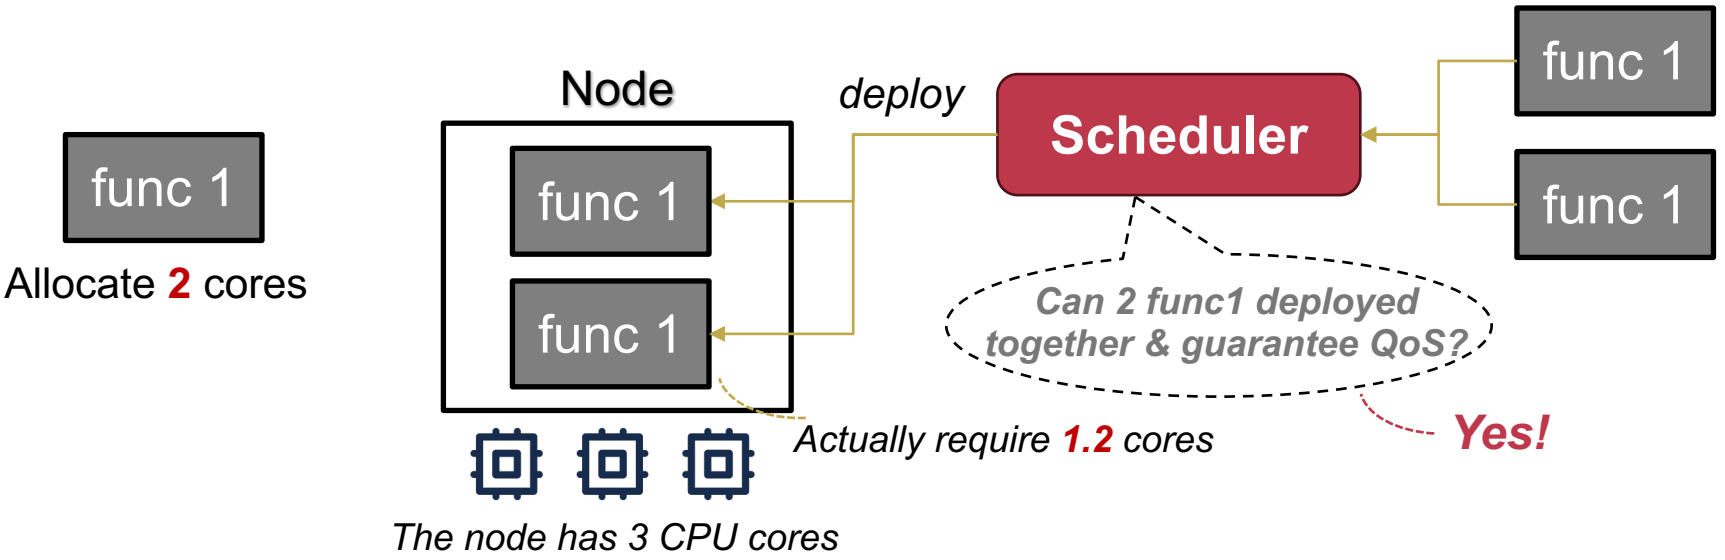

Requirement:

the scheduler should **accurately predict QoS violations**

Challenges of Overcommitment

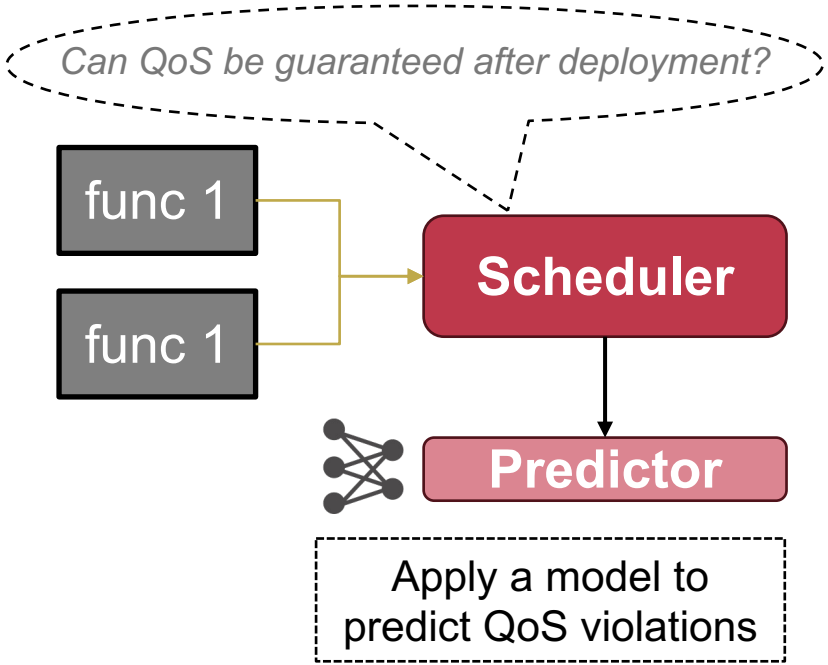

- **Predict with complex models**
  - **Accurate prediction**
  - **Costly (>tens of ms)**
- **Predict with heuristic models or historical information**
  - **Inaccurate prediction** or **unscalable profilings**
  - **Fast (~1ms)**

Challenge I:

Achieve accurate prediction & practical cost (<10ms) simultaneously

Mitigate Wastage Part II: Sensitive Autoscaling

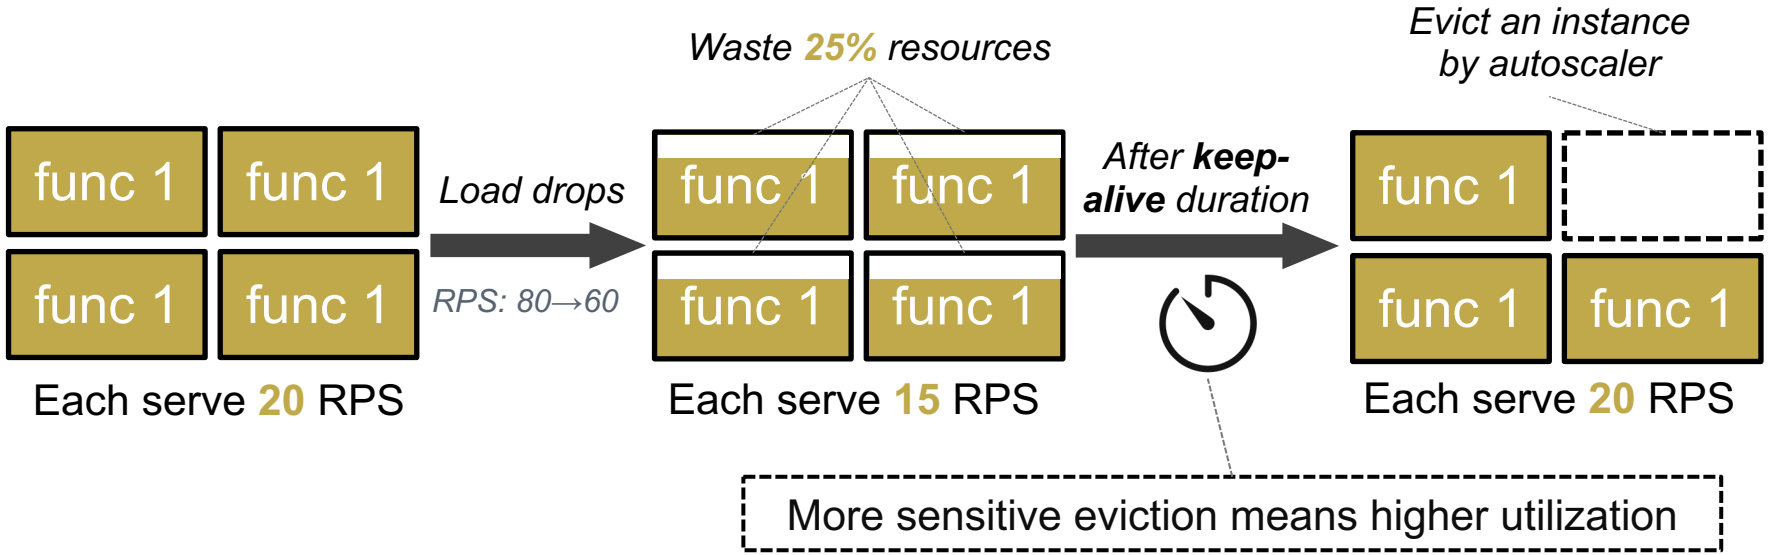

Autoscaling: dynamically reclaim unused resources upon load drops

Challenges of Sensitive Autoscaling

Problem: more sensitive eviction could mean more cold starts

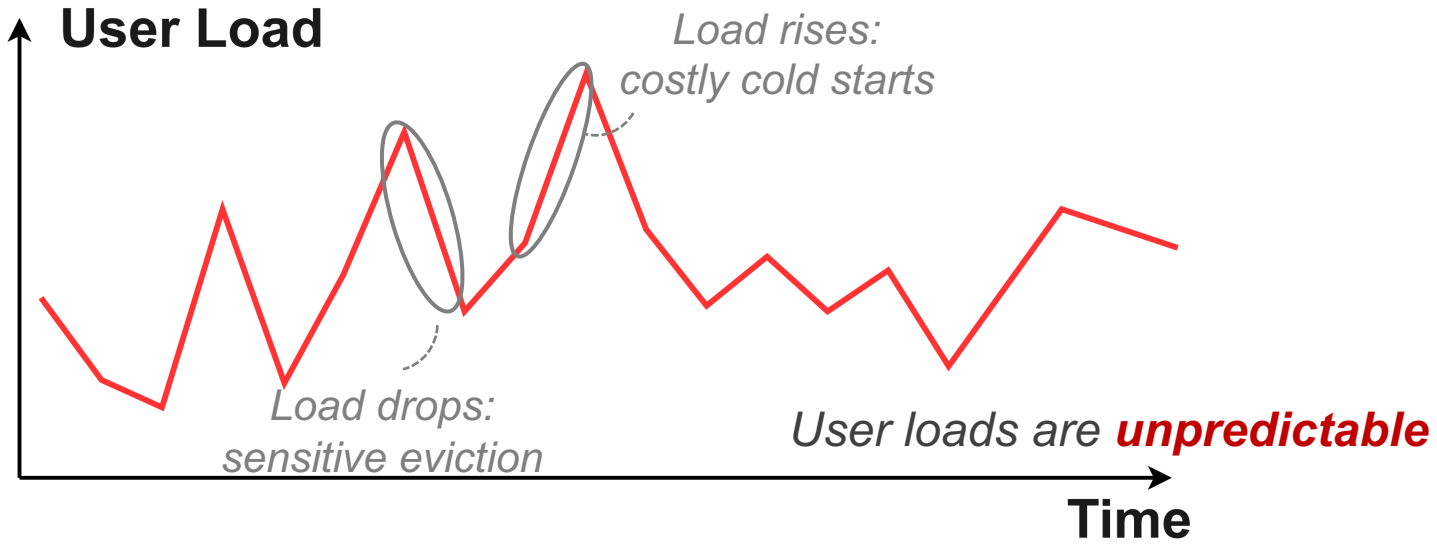

Challenge II:

Achieve high utilization and low cold start costs simultaneously

Jiagu: Two Designs to Break the Tradeoffs

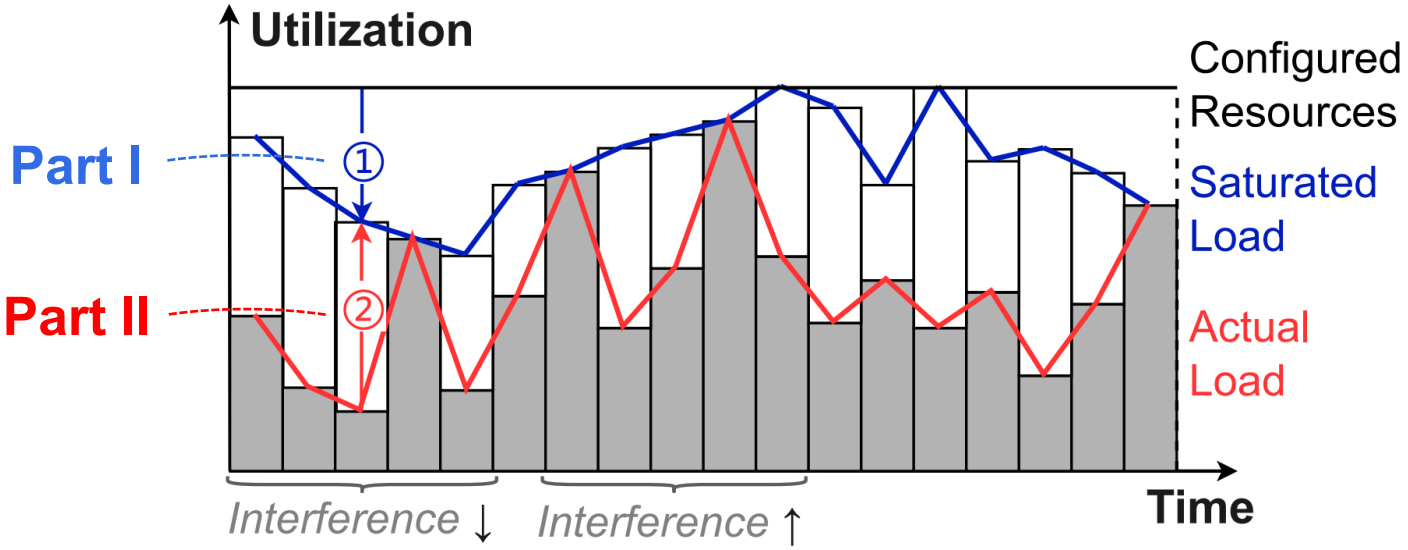

Design I for Part I: pre-decision scheduling

Design II for Part II: dual-staged scaling

1. Achieve both efficiency and performance for overcommitment

DESIGN I: PRE-DECISION SCHEDULING

Insight I: Decouple Prediction and Decision Making

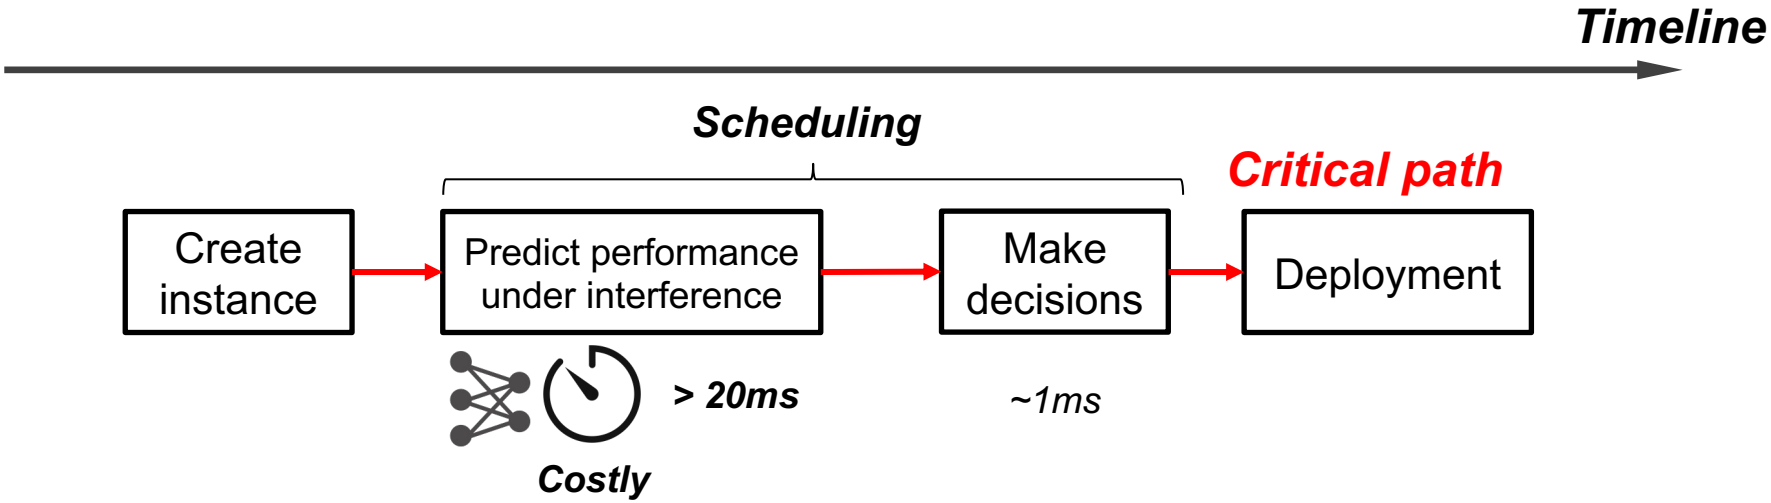

Costly predictions are on the critical path

Insight I: Decouple Prediction and Decision Making

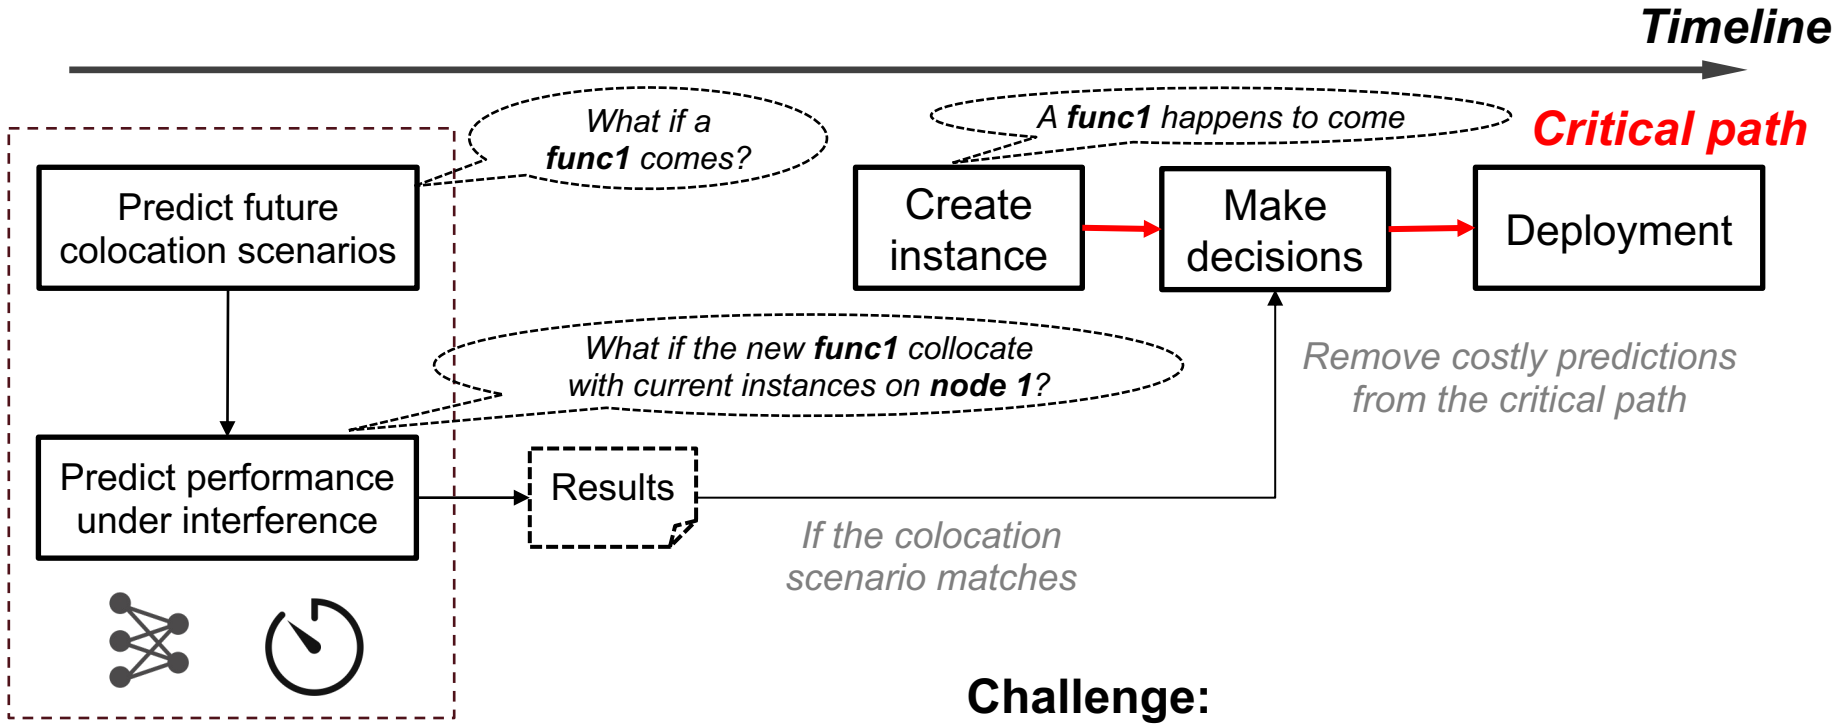

Challenge:
impossible to traverse all possible colocation environments

Serverless Highly-replicated Nature

Serverless instances are **highly replicated**

~56% instances are of functions that have >12 replicated instances

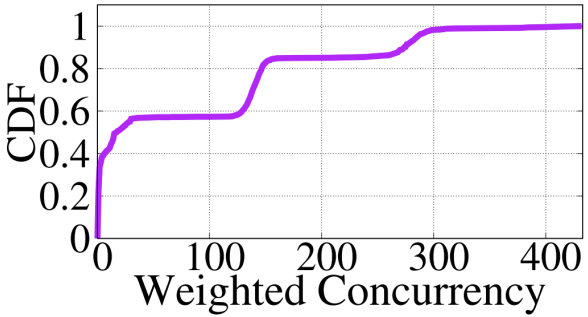

(a) Weighted instance concurrencies of functions.

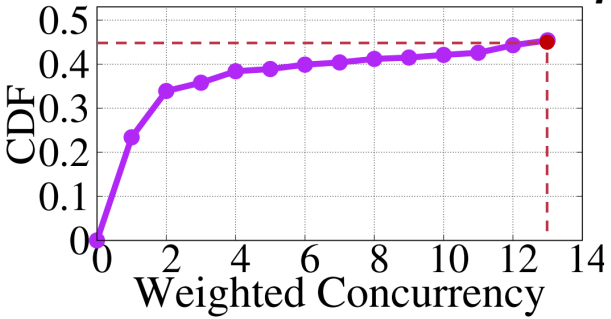

(b) Weighted instance concurrencies (<13) of functions.

Insight I: Decouple Prediction and Decision Making

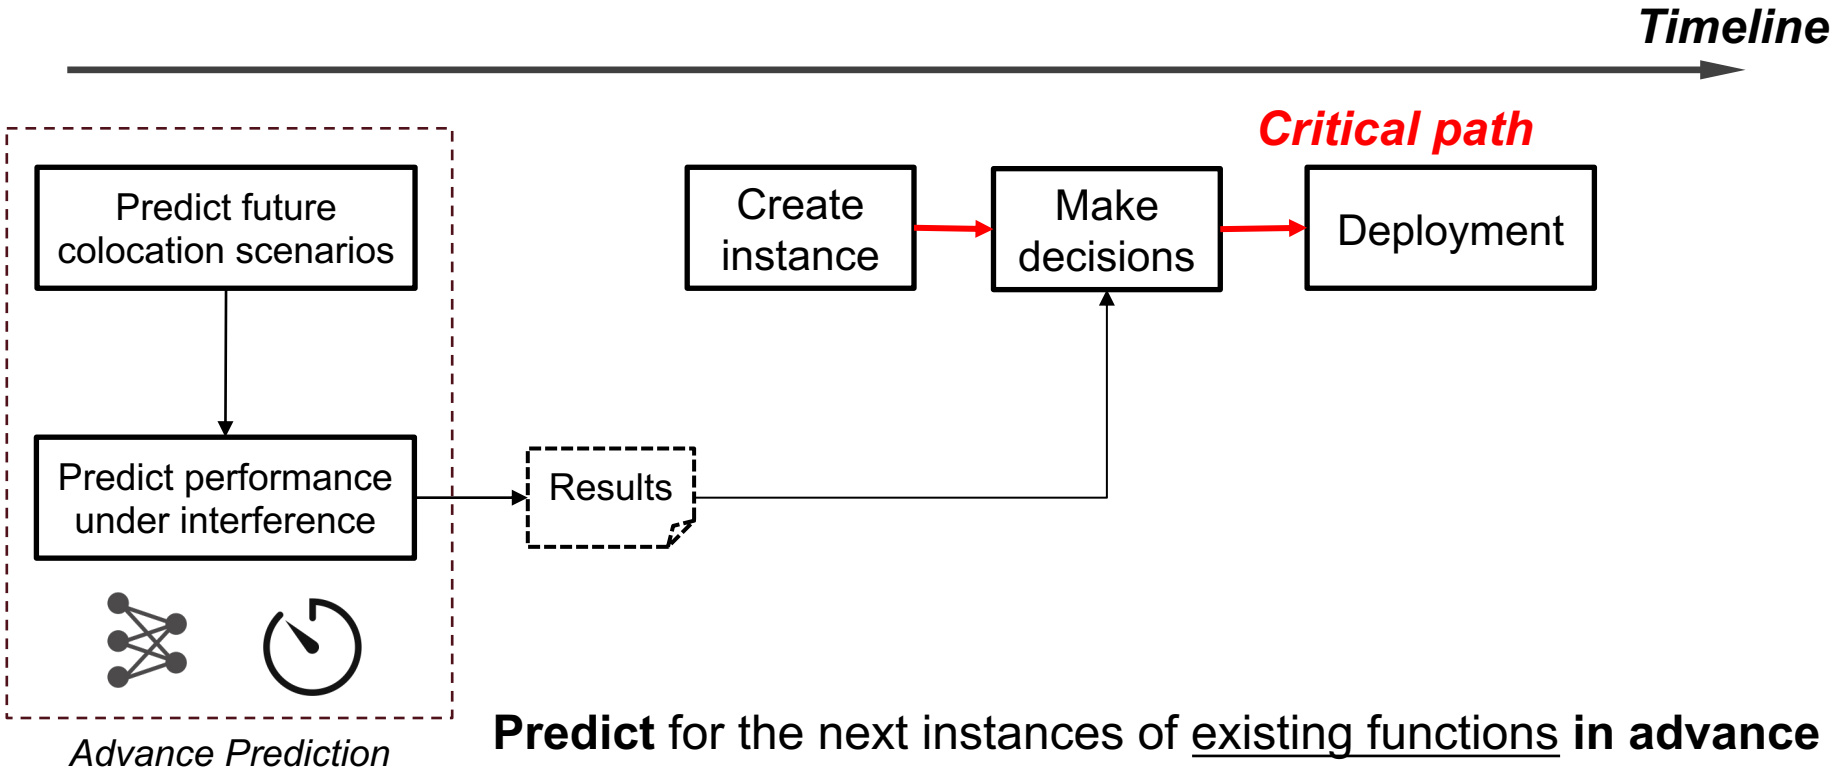

Predict for the next instances of existing functions **in advance**

Advance Prediction to Construct the Capacity Table

Predict for the next instances of existing functions **in advance**

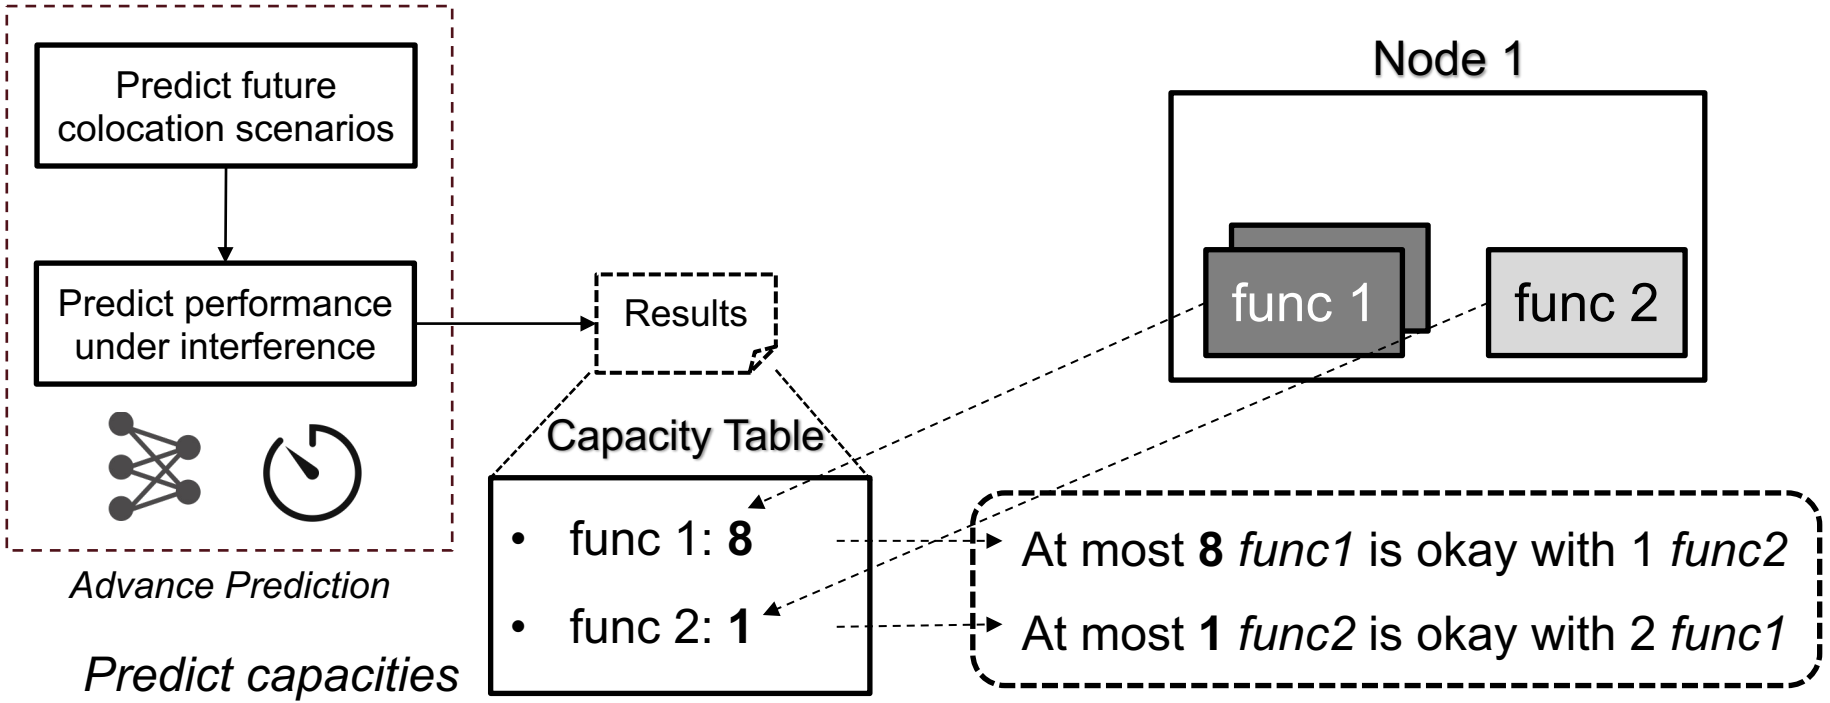

Pre-decision Scheduling: Basic Idea

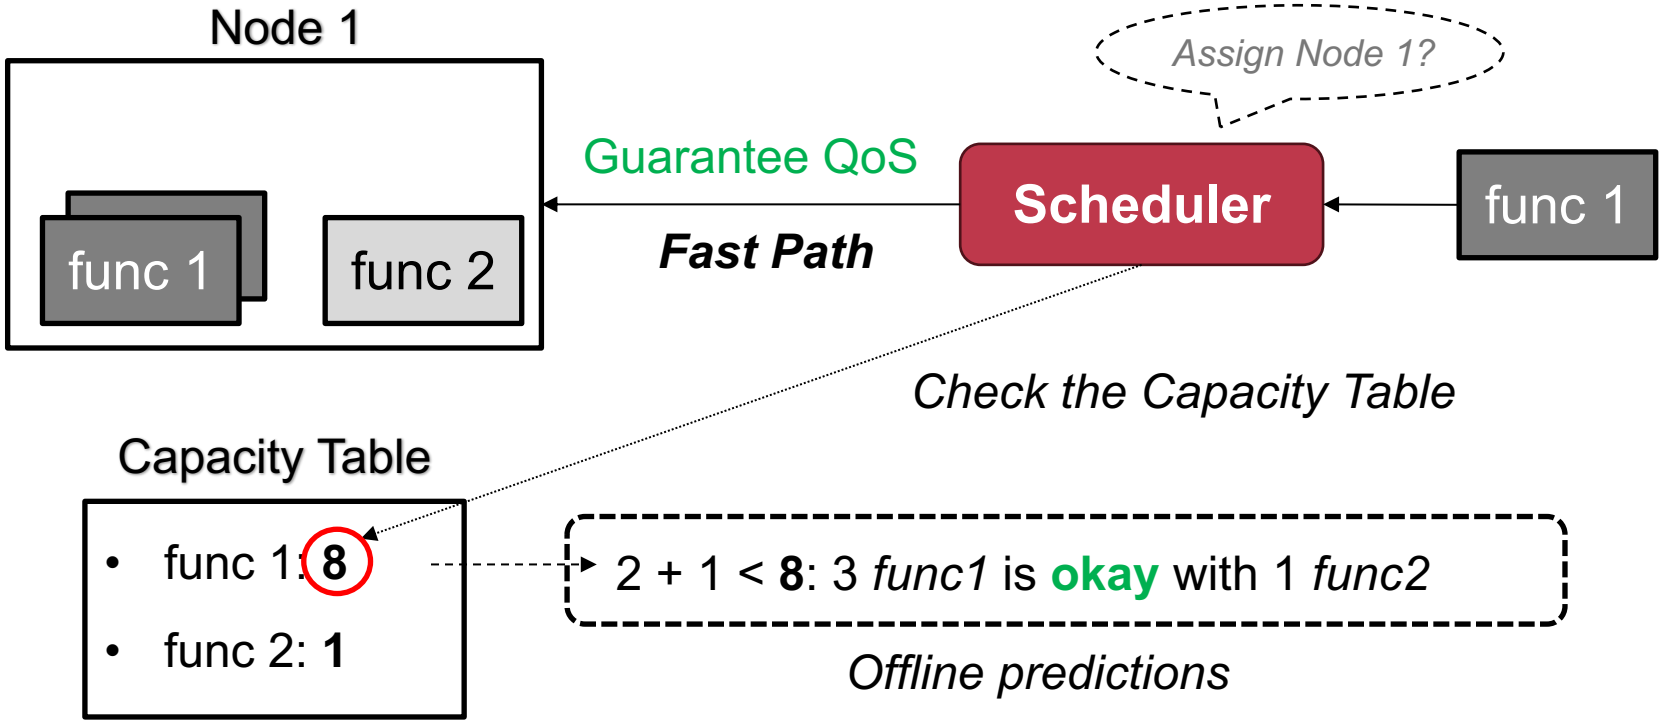

Pre-decision Scheduling: Basic Idea

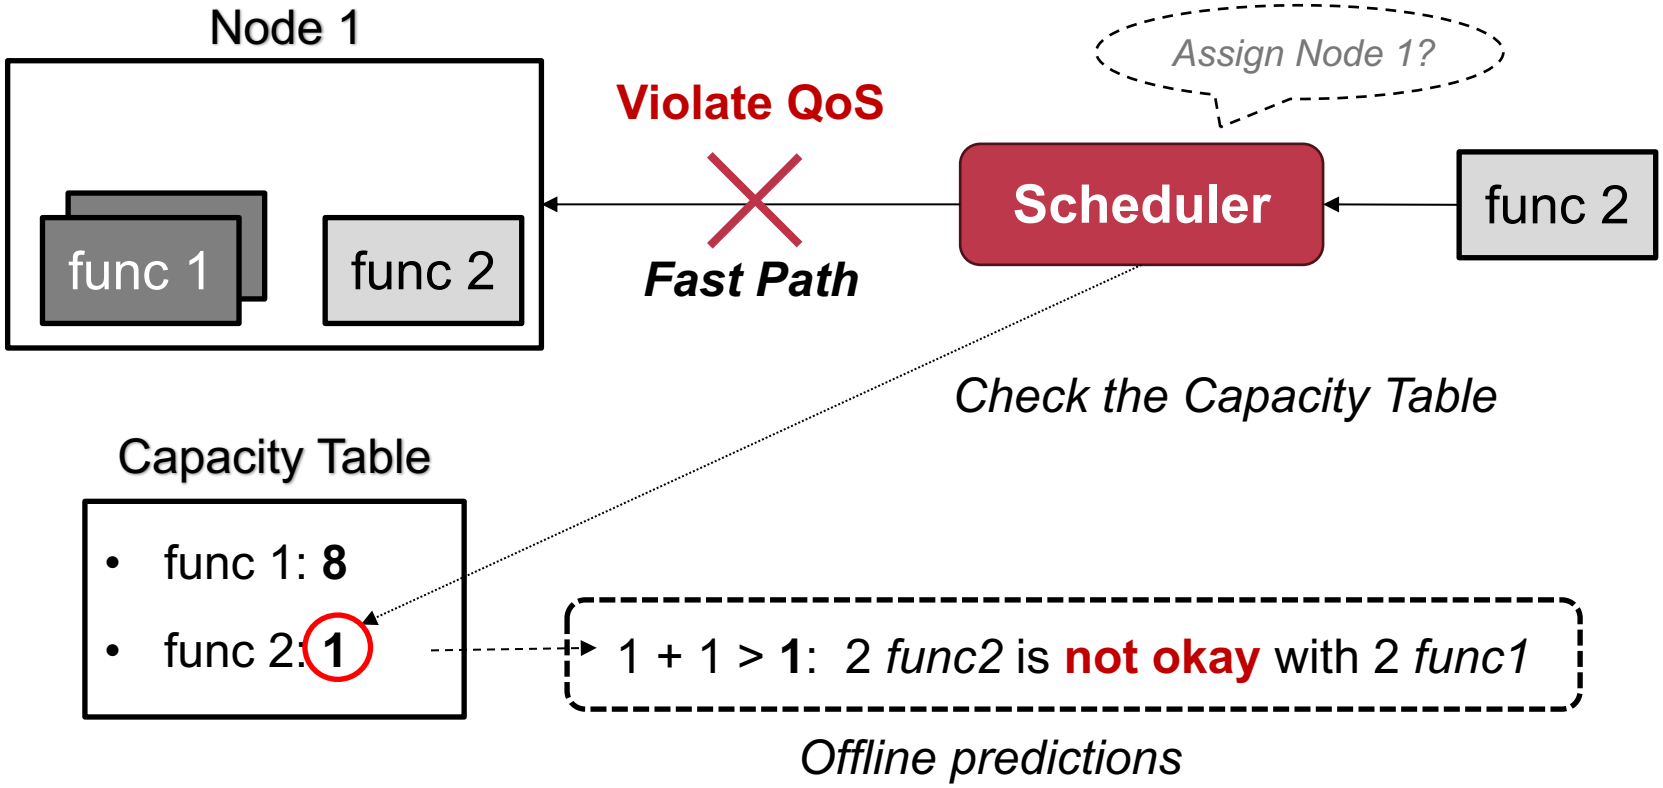

Pre-decision Scheduling: Basic Idea

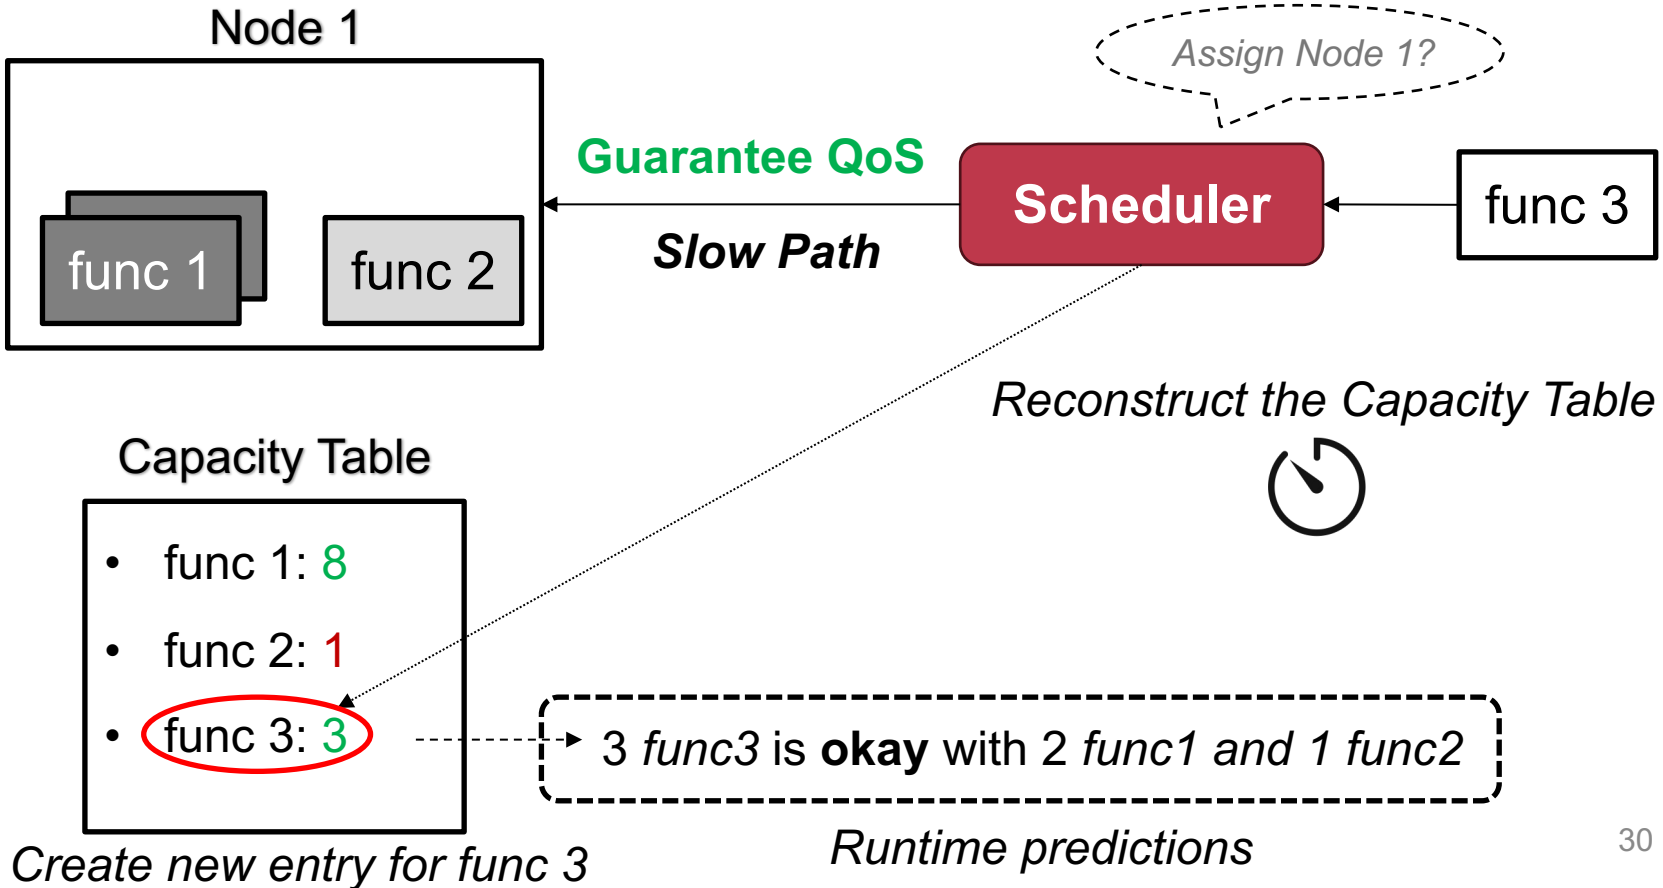

Pre-decision Scheduling: Basic Idea

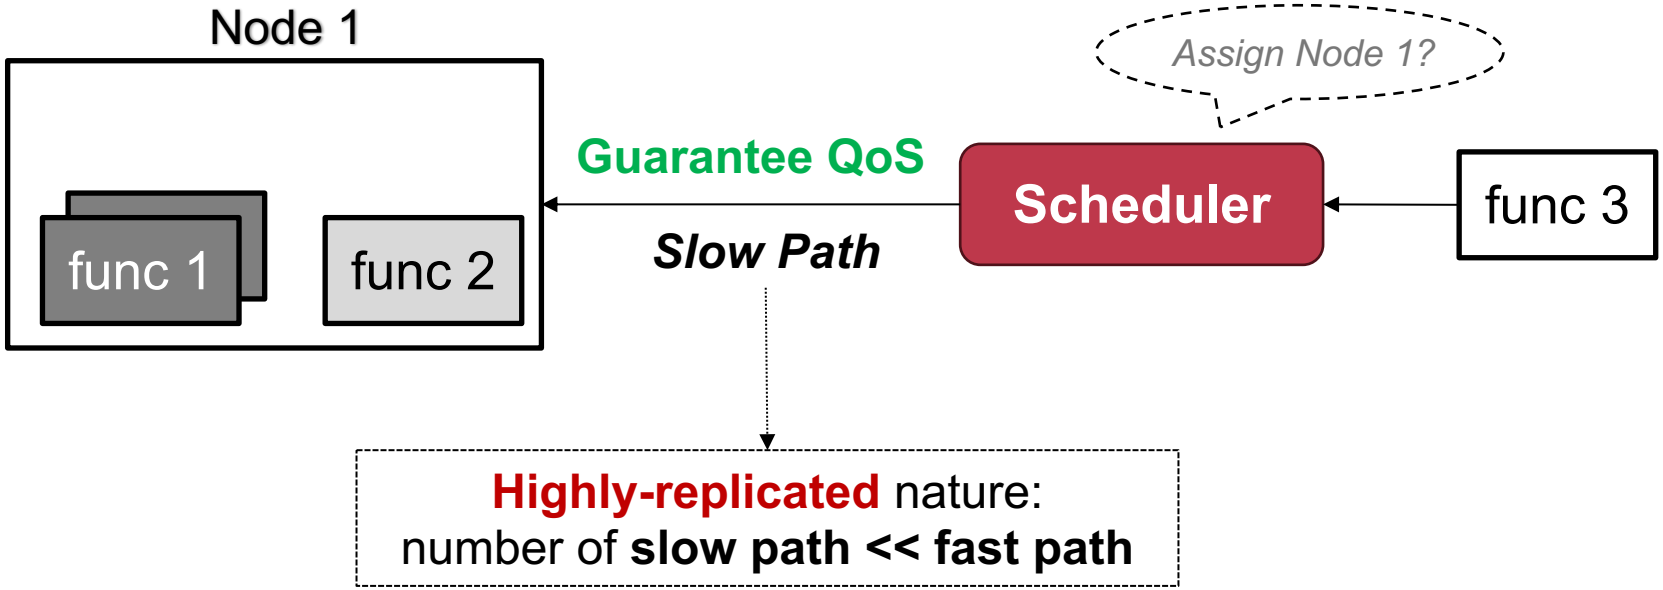

Asynchronous Update

- **Capacity table:**
  - Require timely update: the colocation environment is constantly changing
- **Asynchronous update:**
  - Keep the capacity table up-to-date
  - Prevent the updating from introducing prediction overhead in the critical path

(Details in the paper)

Calculate the Capacity

Features of co-located functions

$$P_{A\cup\{B,C,...\}} = RFR\{P_A, R_A, C_A, \overbrace{R_B, C_B, R_C}, C_C, ...\}$$

Solo-run performance | Profiles | Concurrency

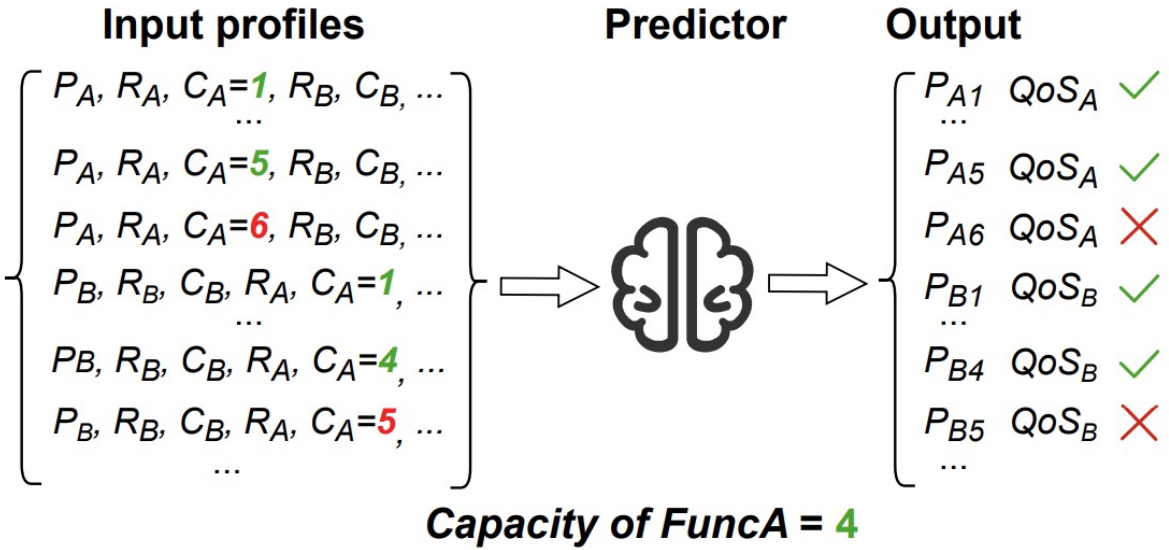

Find the **maximum QoS-guaranteed concurrency:**
(Details in the paper)

Pre-decision Scheduling: Put It All Together

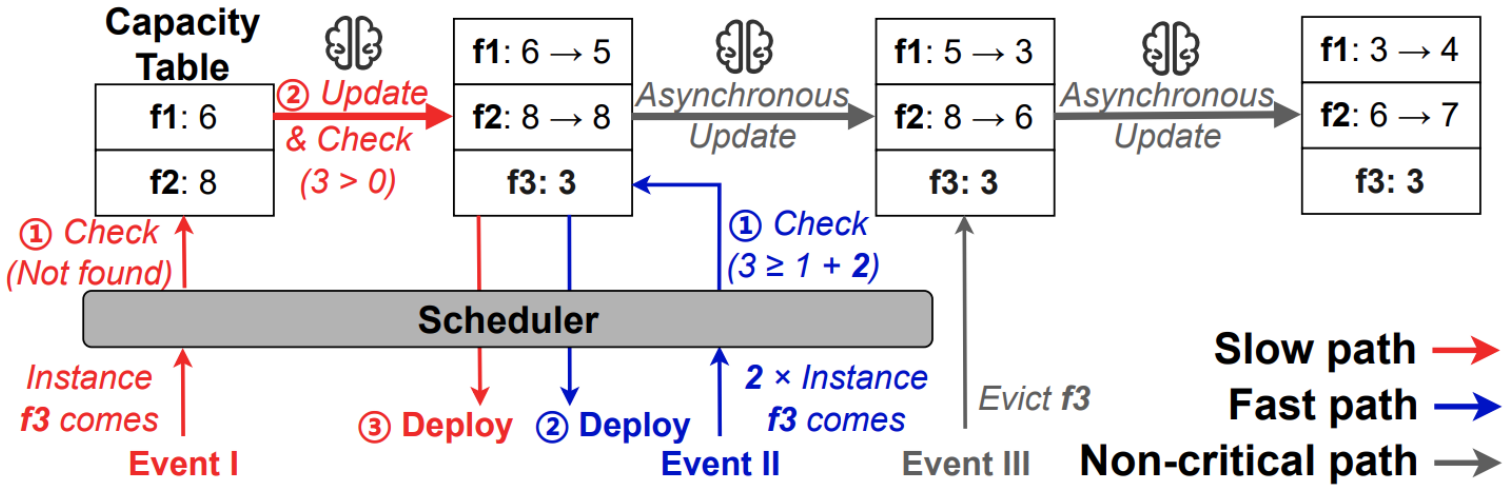

(Details in the paper)

2. Achieve sensitive autoscaling for high utilization without additional cold starts

DESIGN II: DUAL-STAGED SCALING

Insight 2: Decouple Resource Releasing and Instance Eviction

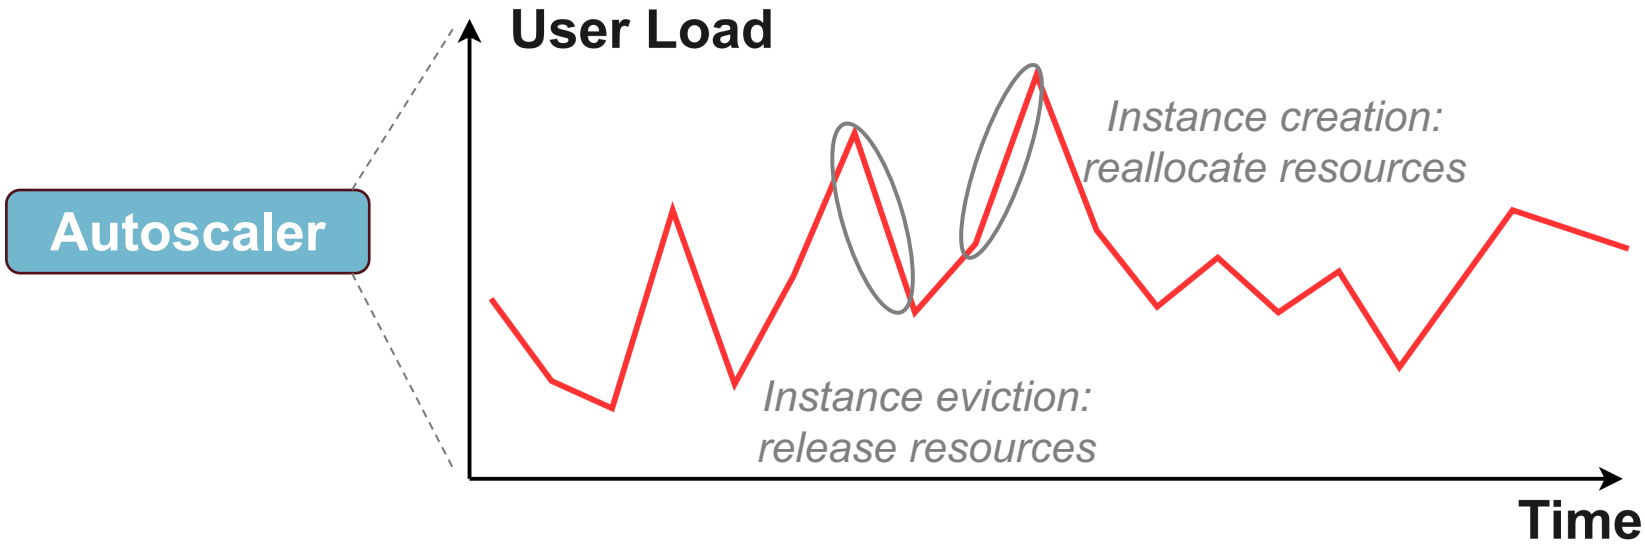

Root cause of overheads:

The coupling between **resource allocation/release** and **instance creation/eviction**

Dual-staged Scaling: Basic Idea

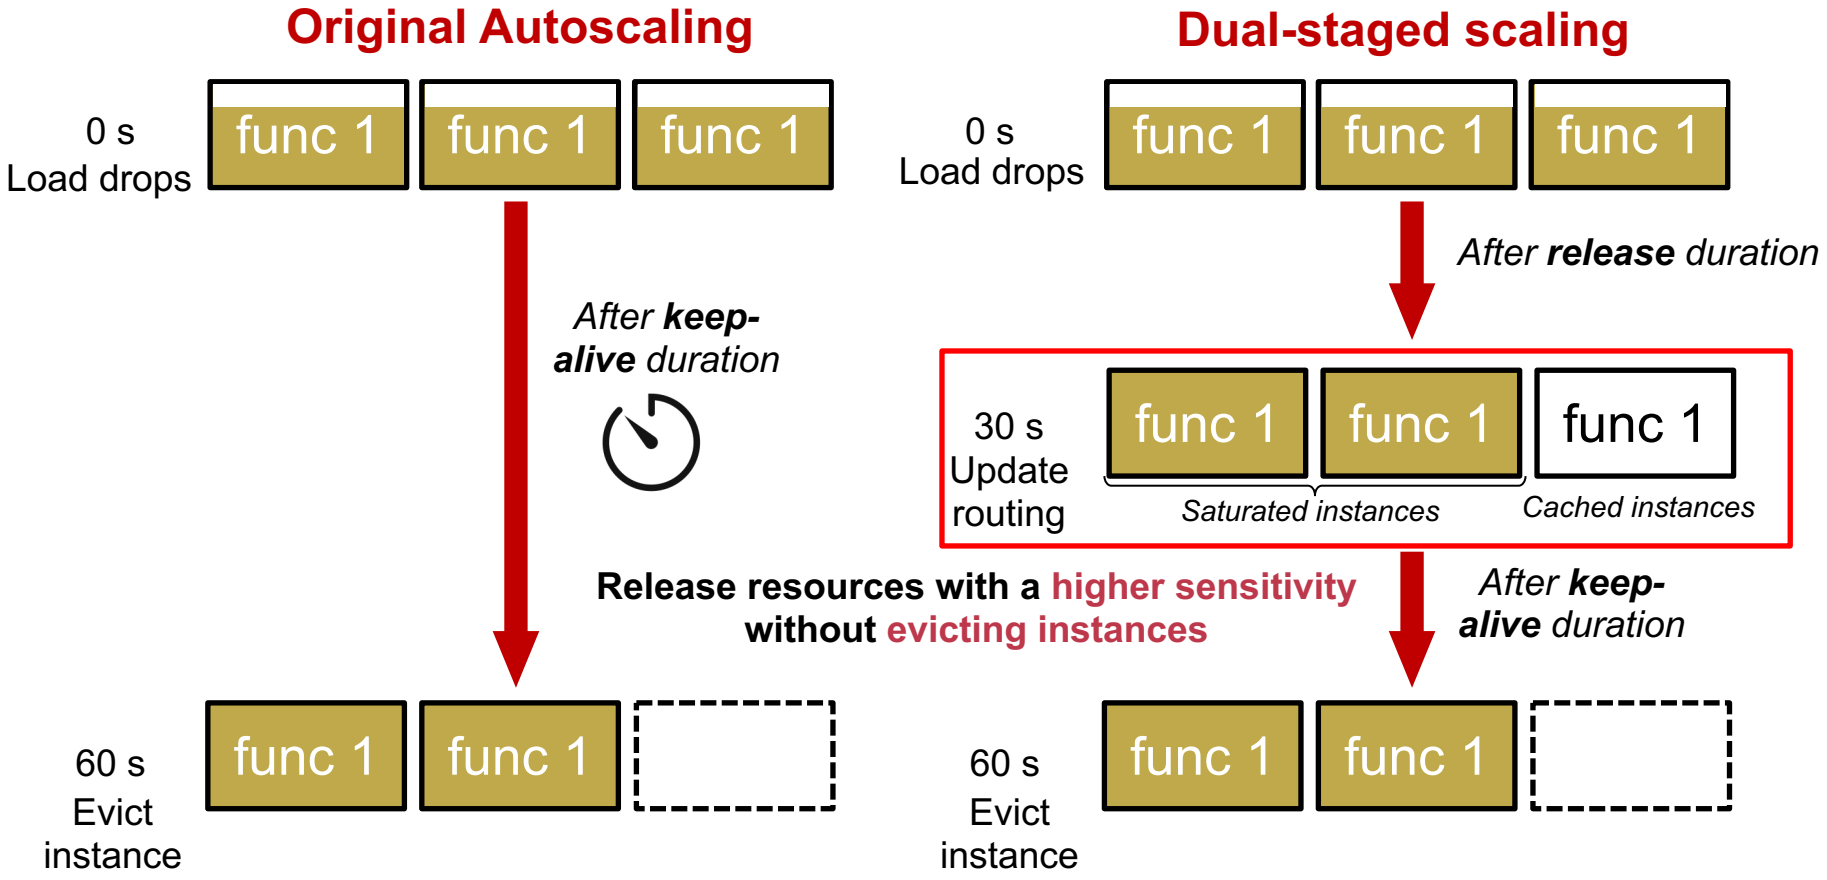

Dual-staged Scaling: Basic Idea

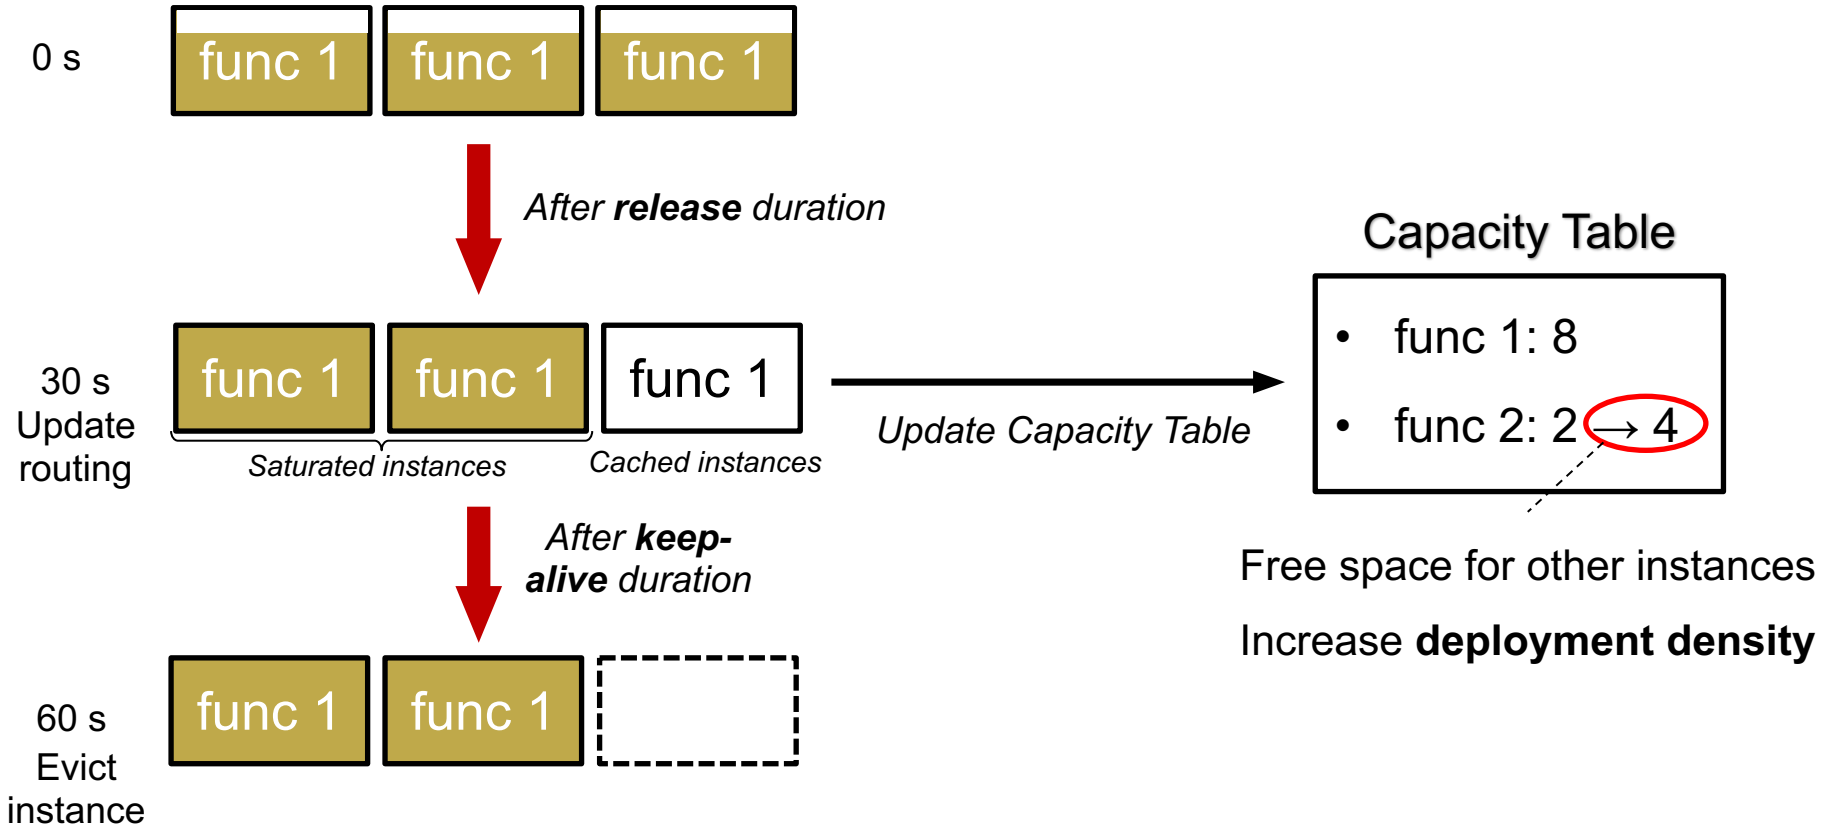

Dual-staged Scaling: Logical Cold Start

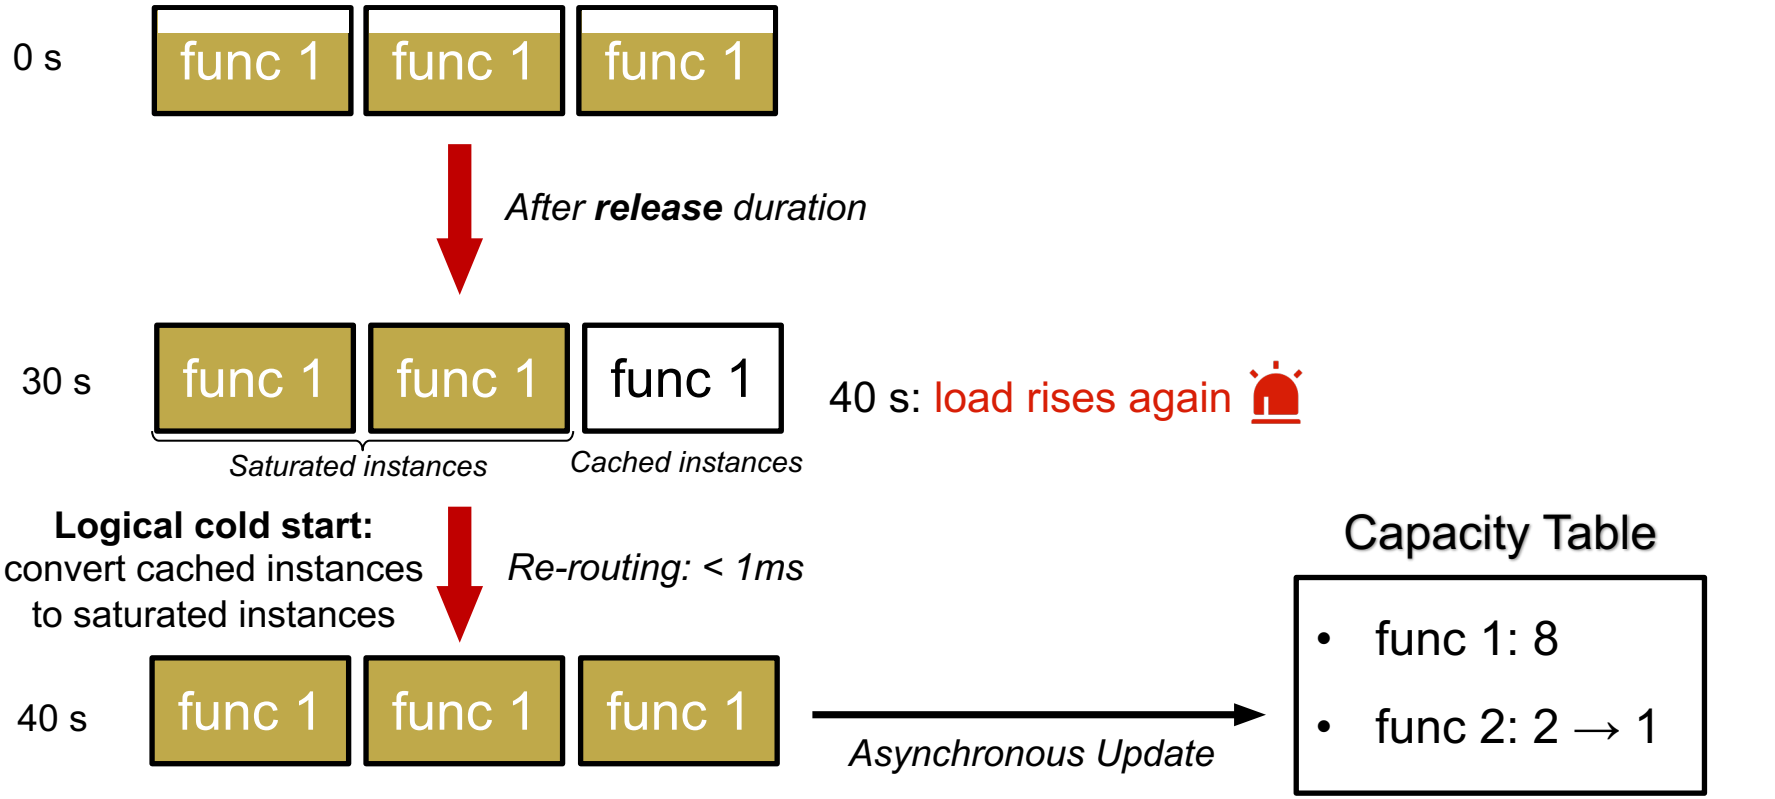

Dual-staged Scaling: On demand Migration

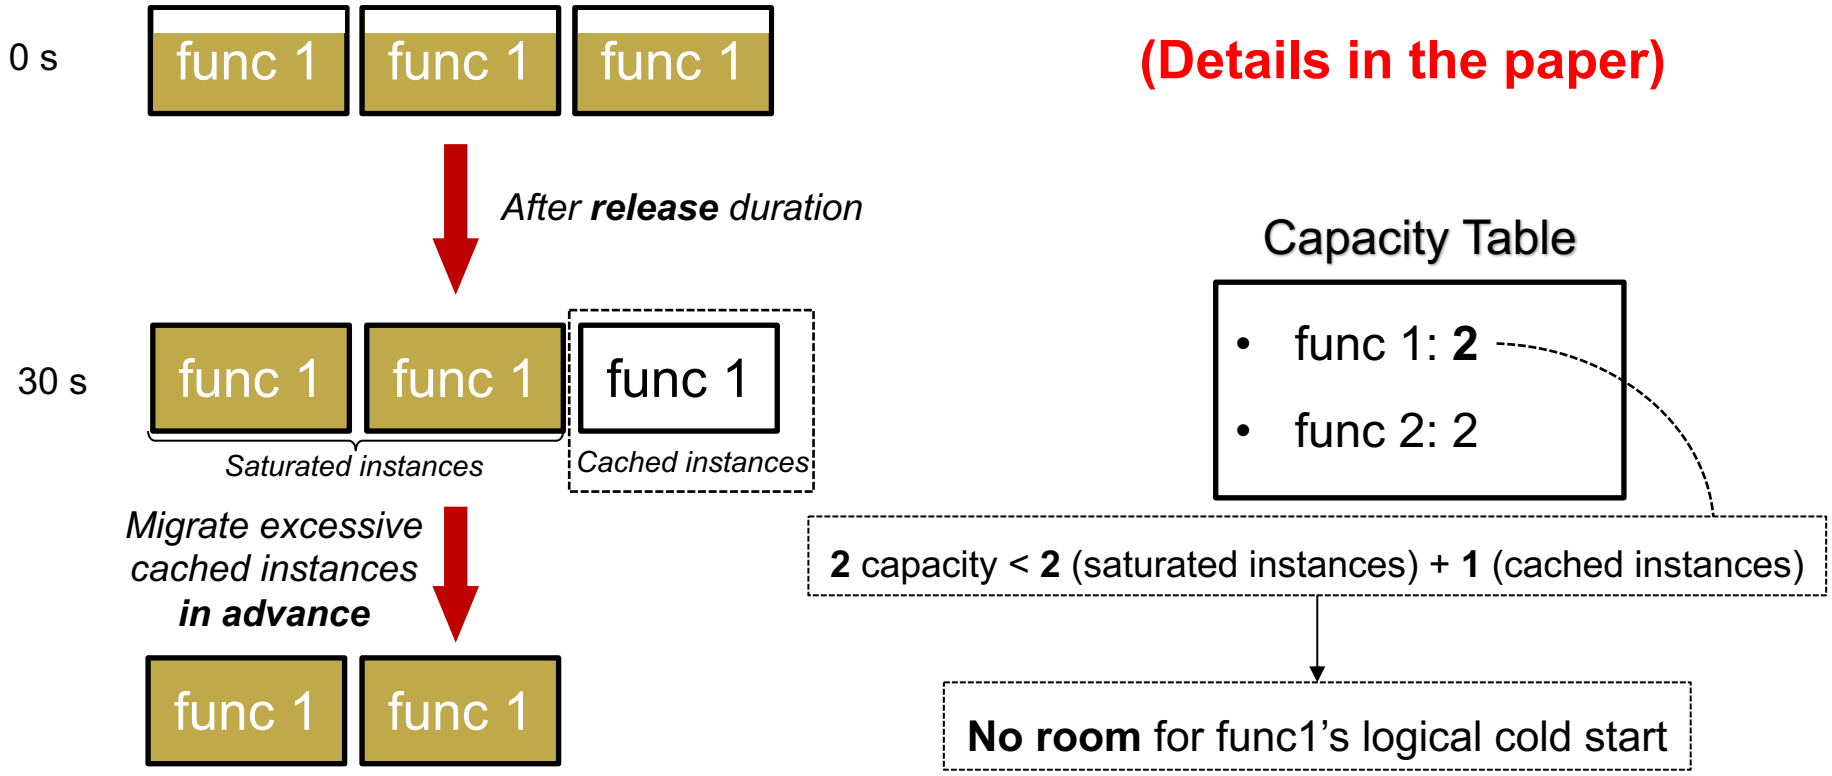

EVALUATION

Evaluation: Effective Scheduling with **Practical Cost**

Optimize scheduling costs

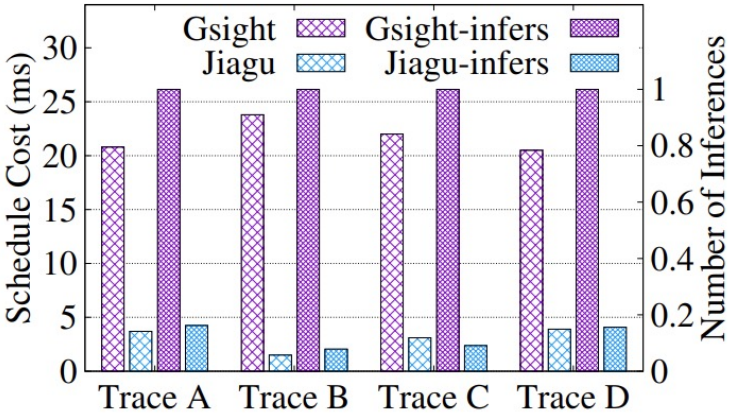

Reduce **83.8%–92.1%** inferences on critical path

81.0%–93.7% lower scheduling costs

Optimize total cold start costs

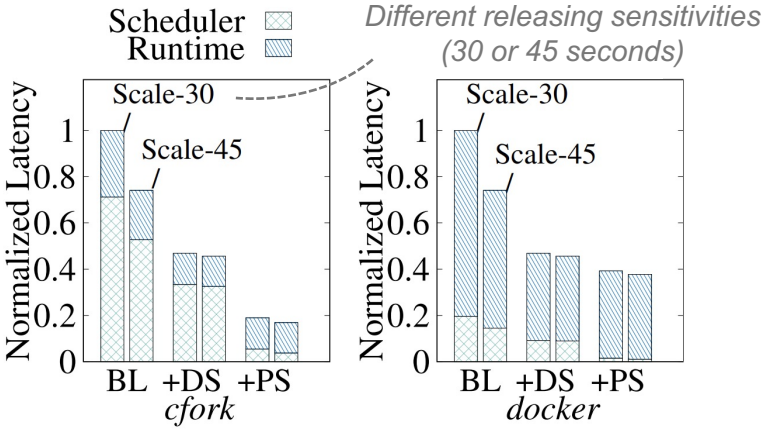

Dual-staged Scaling (DS):

reduce the number of cold starts

Pre-decision Scheduling (PS):

reduce the cost of each cold start

Evaluation: Effective Scheduling with Practical Cost

Optimize resource utilization

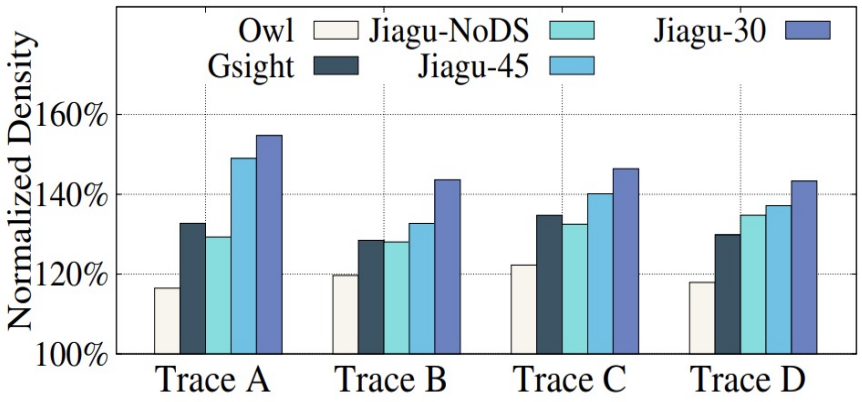

Up to **22%** higher deployment density than Gsight

38.3% higher deployment density than Owl

Ensuring QoS with accurate prediction

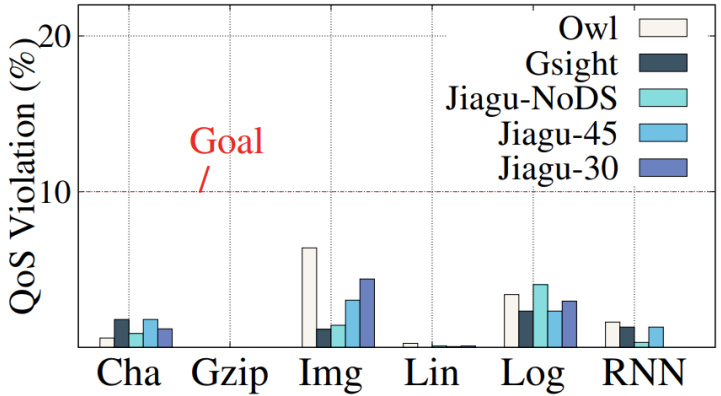

QoS violation rate meets the goal

Conclusion

- **Jiagu: optimize resource utilizations of serverless platforms**
- **Pre-decision scheduling: reduce resource overprovisioning**
  - **Overcommitment:** effective scheduling with accurate performance prediction
  - Reduce the prediction cost in the scheduling critical path
- **Dual-staged scaling: reduce load overestimation**
  - Achieve high resource utilization with sensitive autoscaling
  - Eliminate the side effect of incurring additional cold starts

Q&A / Contact Us

- **Qingyuan Liu**
  - PhD student, IPADS Lab, SJTU
  - liu_qy@sjtu.edu.cn
  - Personal Website (Chinese):
    - https://ipads.se.sjtu.edu.cn/zh/pub/members/qingyuan_liu/

- **Dong Du**
  - Assistant Professor, IPADS Lab, SJTU
  - dd_nirvana@sjtu.edu.cn
  - Personal Website:
    - https://dongd.info/