

ORACL: Optimized Reasoning for Autoscaling via Chain of Thought with LLMs for Microservices

Haoyu Bai¹, Muhammed Tawfiqul Islam¹, Minxian Xu², Rajkumar Buyya¹

Abstract—Modern microservice and serverless applications require autoscaling to balance resource utilization and Quality of Service (QoS). Since effective scaling decisions often depend on root-cause analysis (RCA), controllers must reason over a large and evolving decision space. Existing approaches either require costly per-deployment training or depend on brittle hand-crafted rules. We propose ORACL, **Optimized Reasoning for Autoscaling via Chain of Thought with LLMs for Microservices**, an LLM-guided framework for joint RCA and resource allocation. ORACL converts runtime telemetry into natural-language state sketches and uses an instruction-tuned, reinforcement-learning-regularized LLM to identify root causes, prune candidate actions, and recommend safe resource adjustments. We curate, profile, and annotate over 80GB of traces and metrics drawn from open-source microservice workloads to train this model. Experiments show that ORACL improves RCA accuracy by 15% over state-of-the-art baselines, accelerates training by up to 24×, and improves QoS by 6% in short-term scenarios. These results demonstrate that LLM-guided reasoning provides generalizable, interpretable, and actionable resource management for cloud-native systems.

Index Terms—LLM-based resource management, Microservices, Root Cause Analysis, Orchestration, Quality of Service

I. INTRODUCTION

The emergence of complex microservice architectures represents a transformative shift in software system design and management, significantly impacting the technological landscape. These modern applications, composed of lightweight and scalable containerized microservices, greatly enhance scalability, agility, and resilience. Microservices [1] are typically loosely coupled, deployed across diverse physical nodes globally, and managed by various service providers [2]. Some functionalities may even reside on privately hosted residential machines due to privacy concerns. Such deployments present diverse configurations, varying considerably in computational capabilities and network latencies [3]. Users often migrate between different cloud platforms or virtual machine types to optimize for energy efficiency, cost savings, or enhanced service performance. Additionally, microservices evolve rapidly, with frequent functional and structural updates driven by developers striving to maintain competitiveness in dynamic software markets.

QoS in microservice applications, often measured by metrics such as average response time [4], is significantly affected

by both the number of container instances and the resources allocated to each. Flexible deployment strategies and pay-as-you-go pricing models give users substantial control over resource utilization. However, accurately identifying existing and potential bottlenecks and their root causes is essential for adjusting resources appropriately. In practice, many operators over-provision resources for specific applications to meet service-level objectives (SLOs), but this can lead to substantial resource waste [5]. Such inefficiency increases operational costs by unnecessarily consuming virtual machine capacity and contributes to environmental impacts. Consequently, efficient RCA and autoscaling techniques are critical to balancing QoS requirements with budget constraints.

Traditional approaches to cloud incident RCA still rely heavily on on-call engineers (OCEs) manually collecting and correlating heterogeneous telemetry, such as logs, metrics, and traces, from multiple applications and services [6]. In large-scale cloud environments, the sheer volume and heterogeneity of this data make it difficult for OCEs to rapidly identify the small subset of signals that are truly relevant to a given incident, thereby delaying effective diagnosis and remediation. Even when automation is available, the overall process remains labor-intensive and error-prone, as engineers must sift through noisy, sometimes incomplete information, interpret complex system behavior, and iteratively infer likely root causes. As a result, despite substantial advances in data-driven and artificial intelligence-based techniques for operations and monitoring, OCEs still invest significant manual effort in analyzing incident data and determining the underlying causes of failures [7].

Current autoscaling methodologies predominantly utilize rule-based systems or deep neural networks with extensive parameterization [8]. Rule-based methods are highly dependent on manually defined rules, which may fail to adequately capture complex interactions within microservices due to vast combinatorial complexities. Conversely, deep learning-based approaches require large training datasets and often take extended periods, sometimes weeks, to converge [9]. The associated training latency can degrade initial system performance, and the data collection process itself can be a burden, typically taking tens of seconds to minutes per allocation. Changes to deployment environments or the introduction of new deployments require redesigning neural networks, as shifts in state and action spaces require prolonged retraining. Additionally, these resource-intensive neural network models typically demand significant GPU resources, creating accessibility challenges for small-scale enterprises.

Moreover, existing technologies often overlook significant domain-specific prior knowledge, such as YAML specifica-

¹Quantum Cloud Computing and Distributed Systems (qCLOUDS) Lab, School of Computing and Information Systems, The University of Melbourne, Australia.

²Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, China.

tions, microservice topologies, service functionalities, relationships within service chains, and specific user requirements. Understanding domain-specific characteristics, for example, higher I/O and memory demands in Online Transaction Processing (OLTP) systems, CPU-intensive operations in Online Analytical Processing (OLAP) applications [10]. However, current methodologies generally neglect this critical knowledge, thereby relying on inefficient, purely neural-network-driven exploratory processes. Leveraging human domain expertise can substantially reduce the complexity of exploration, leading to more efficient and practical autoscaling and RCA solutions.

To overcome these limitations and leverage recent advancements in LLMs [11]–[13], we propose a semantic-aware resource allocation framework that integrates various cloud-oriented tools, allowing the LLM agents [14] to use different cloud-native tools for automated reasoning and resource optimization. Our method comprehensively models microservice characteristics and employs cloud-specific chain-of-thought reasoning techniques [15] to optimize resource allocation more effectively.

The main **contributions** of this paper are:

- 1) We design a novel framework that enables LLMs to directly interact with cloud clusters by transforming raw cloud states into semantically meaningful representations, allowing the model to perform automated reasoning and resource adjustment. Based on this design, we construct and annotate a large-scale prompt–response dataset (over 80 GB) that captures engineers’ expertise in diagnosing and analyzing cloud-related issues, serving as a foundation for training and evaluation.
- 2) We propose a unified algorithm that connects RCA and Resource Allocation through an analogy-based formulation. This approach embeds a training pipeline that optimizes the model’s intermediate reasoning process itself, enabling the framework to generalize across diverse microservice-based applications.
- 3) We conduct comprehensive testing and validation of the proposed model and policies using real-world traces and a dedicated experimental platform. These evaluations thoroughly assess the effectiveness and performance of both RCA and autoscaling for microservices.

The rest of the paper is organized as follows. Section II provides background and related work on both RCA and autoscaling for microservices. Section III discusses the limitations of prior work and provides an illustrative experiment highlighting these shortcomings, thereby motivating a system architecture designed to overcome them. Section IV introduces the overall architecture and high-level design of each component. Section V details the implementation and underlying mathematical formulations. Section VI presents a comprehensive experimental evaluation. Section VII concludes our work and outlines potential directions for future research.

II. RELATED WORK

In this section, we review root-cause diagnosis and orchestration techniques for microservices. Prior studies [16],

[22], [23], [27] demonstrate that accurate diagnosis plays a critical role in microservice-based systems, as it can either guide orchestration decisions or be reformulated directly as a scheduling problem. Existing methods often rely on anomaly-injection datasets to learn mappings from failure patterns to corrective actions, after which the control procedure is recast as a resource-allocation or scheduling task. This workflow highlights the tight coupling between diagnosis quality and orchestration efficiency in microservice environments.

A. RCA Problems

In microservice-based cloud systems, RCA serves as the foundation for reliability engineering by enabling precise fault localization, noise reduction, causality-driven autoscaling, and proactive failure containment. Without accurate RCA, resource management and self-healing actions risk reacting to symptoms rather than actual bottlenecks, leading to instability and inefficient resource usage. The seminal work Murphy [16] introduced a probabilistic graphical model based on a *Markov Random Field (MRF)* to diagnose performance anomalies from loosely defined metric associations, eliminating the need for an explicit acyclic dependency graph. To enhance reliability and efficiency, BARO [17] integrates Bayesian Online Change Point Detection (BOCPD) with statistical hypothesis testing, unifying anomaly detection and RCA in a single, robust framework.

Interpretable causal reasoning has also advanced. Chain-of-Event (CoE) [18] learns weighted dependencies among metrics, logs, and traces through an event-level causal graph, achieving over 85% top-1 accuracy and producing explicit causal chains for human-aligned interpretability. Compared with Murphy’s on-the-fly inference, CoE leverages historical incident data to improve accuracy on recurring failures. Meanwhile, SLIM [20] introduces a lightweight, interpretable rule-based RCA model designed for imbalanced or novel failure data. SLIM generates concise, human-readable *IF–THEN* rules optimized for F1-score and can adapt online, achieving high interpretability and efficiency for continuously evolving systems.

Recently, RCACopilot [19] introduced LLMs into Microsoft’s incident management workflow, combining structured telemetry with unstructured data such as alerts and tickets to generate natural-language RCA explanations. It follows a retrieval-augmented generation (RAG)-style approach by retrieving the top-k most similar historical incident prompts as contextual evidence. Its large-scale deployment shows how RCA can be coupled with human-like reasoning, contextual retrieval, and adaptive summarization.

B. Autoscaling

In microservice architectures, autoscaling dynamically adjusts the number of service replicas based on real-time workload and performance signals, ensuring resource efficiency, service stability, and consistent quality of experience under fluctuating demand.

TABLE I: A comparison of related works with our proposed system

| Framework | Task-Level | | Data-Level | | Reasoning-Level | | System-Level | |
|---------------------------|------------------------|------------|-------------------|------------------|-----------------|---------------|----------------|--------------|
| | Root Cause/ Bottleneck | Scheduling | Unstructured Data | Interpretability | Reasoning | Optimized CoT | Unifying Tasks | Transferable |
| Murphy [16] | ✓ | | | ✓ | ✓ | | | |
| BARO [17] | ✓ | | | ✓ | ✓ | | | |
| CoE [18] | ✓ | | ✓ | ✓ | ✓ | | | |
| RCACopilot [19] | ✓ | | ✓ | ✓ | ✓ | | | ✓ |
| SLIM [20] | ✓ | | | ✓ | ✓ | | | |
| DRPC [21] | | ✓ | | | | | | |
| FIRM [22] | ✓ | ✓ | | | | | ✓ | |
| AWARE [23] | ✓ | ✓ | | | | | ✓ | ✓ |
| IntentContinuum [24] | ✓ | | | | ✓ | | | |
| Kubeintellect [25] | ✓ | | | | | | | |
| Analytical [26] | | ✓ | ✓ | ✓ | ✓ | | | |
| ORACL (Our Method) | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |

1) Rule-Based and Deep Learning-Based Approaches:

Autoscaling for microservices spans rule-based heuristics and learning-based controllers. Threshold mechanisms (e.g., AWS Auto Scaling and the Kubernetes HPA) trigger scale-out/in when utilization crosses manually set limits but are workload-agnostic and error-prone [28]. To improve robustness, ProScale [29] combines short-horizon Simple Moving Average forecasting with a greedy allocator and a decision-tree regressor to predict the required instance count under bursty traffic [29]. LSRAM further frames service-level-objective (SLO)-aware provisioning as a convex optimization that minimizes resource cost subject to tail-latency constraints, and solves it via gradient descent, accelerated by a vector-projection/cosine-similarity approximation for distributing partial SLO budgets [30].

Learning-based methods cast scheduling as sequential decision making. Elarnng applies UCB1 multi-armed bandits to select VM counts based on CPU utilization trends [31]. GRAF builds hierarchical service graphs and uses workload prediction with graph neural networks for dependency-aware allocation, albeit with non-trivial graph maintenance overhead in highly dynamic settings [32]. Deep RL systems enrich workload representations: FIRM identifies critical chains using SVMs and optimizes resources via an actor-critic agent [22], while AWARE adds meta-learning and bidirectional RNN encoders to accelerate convergence [23]. However, their synchronous optimization along critical paths can overreact under asynchronous fluctuations. DRPC addresses this by coupling imitation learning with TD3 and decomposing policies into lightweight node-local submodels, reducing inference cost and improving responsiveness at scale [21].

2) *LLM Based Approaches*: LLMs are increasingly applied to autonomic system management across intent-driven control, resource scheduling, and multi-agent orchestration. For passive resource allocation, IntentContinuum and KubeIntellect embed LLMs in runtime loops to monitor signals, diagnose SLO violations, and coordinate remediation across heterogeneous Kubernetes clusters [24], [25]. For proactive scheduling, LLMs act as priors or perform action pruning within reinforcement learning (RL) to stabilize task and job-shop allocation [33]. At the edge, the Edge Computing Power Network (ECPN) demonstrates that LLM-guided offloading and allocation reduce energy consumption, accelerate convergence, and en-

hance privacy under dynamic conditions [34]. Beyond single controllers, multi-agent schemes improve scalability; for example, the work in [27] proposes a hierarchical design that couples low-level planner-executor agents with higher-level managers, coordinating via protocols and message queues. ServiceOdyssey integrates runtime state tracking with LLM-driven curriculum design, execution planning, and knowledge curation, enabling progressive self-learning and increasingly sophisticated microservice orchestration [35].

C. Critical Analysis

Although representative methods have made valuable contributions, notable limitations remain. Existing methodologies for microservice resource allocation have made significant advancements but remain limited, particularly in flexibility, adaptability, and interpretability. Rule-based approaches rely on manually configured thresholds and trigger conditions, inherently limiting their effectiveness due to the combinatorial complexity and NP-hard nature of autoscaling problems. Human-defined rules struggle to encompass the exhaustive permutations and interactions among diverse microservices, leading to suboptimal decisions and restricted applicability across varied scenarios.

Neural networks and reinforcement learning methods similarly face critical constraints, primarily due to their rigid structural designs and heavy reliance on substantial training datasets. These methods necessitate extensive model redesign and retraining whenever microservice configurations evolve, posing significant challenges during frequent updates or early deployment stages. Furthermore, slow convergence rates in actor-critic frameworks restrict rapid adaptation to changes in user behavior or cloud environments. The "black-box" nature of deep learning models additionally diminishes interpretability, reducing user trust and limiting transparency in resource allocation decisions.

Recent works demonstrate the promise of LLMs in cloud management, but most existing approaches still formulate the problem as directly feeding system states, whether encoded as textual descriptions or vector embeddings, into an LLM and relying on generic pretrained knowledge for reasoning. As a result, their outputs are often limited to coarse-grained recommendations, failing to capture cloud-specific workload characteristics and complex contextual signals. This reliance

on static knowledge and fixed deduction patterns leads to poor generalization in dynamic and heterogeneous cloud environments; the underlying motivations are elaborated in Section III.

III. MOTIVATION AND SYSTEM ARCHITECTURE

In this section, we present a motivating experiment, derive the key obstacles it reveals, and outline a high-level architecture to address them.

Fig. 1 reports the convergence behavior of a representative DRPC trained on the Train-Ticket microservice application [36]. The application is deployed on a five-node cluster (one master and four workers); each node has 8 vCPUs and 8 GB of memory. The deployment comprises approximately 40 microservices (e.g., *user*, *station*, *price*, *route*). The reward jointly considers resource utilization and latency. Without synthetic-data augmentation, roughly 19 days of traces are required for DRPC to reach usable performance.




Fig. 1: Convergence of the DRPC Algorithm

Guided by this observation and the analysis in Sec. II-C, we target two core obstacles: (i) limited generality under topology churn and (ii) poor sample efficiency at the microservice scale. We additionally aim to (iii) improve interpretability and (iv) unify RCA and orchestration.

Our approach draws inspiration from two sources: (i) DeepSeek-R1 [37], which explicitly optimizes intermediate reasoning steps, and (ii) the question-driven, decision-tree-like workflow that cloud engineers typically follow when diagnosing complex environments.

The key idea is to encode application-agnostic microservice semantics in natural language that mirrors how practitioners diagnose and act. Because a microservice call-graph can be expressed textually, the representation is *universal*, rather than tied to a single, static application. By designing a structured output schema, we unify RCA and orchestration into a single loop and treat the reasoning procedure itself as the decision policy—analogue to optimizing the rule set in a rule-based system. This design yields interpretable rationales, greater robustness to configuration and topology changes, and reduced sample complexity and training cost.

Building on these ideas, we propose **ORACL** and its high-level architecture as shown in Fig. 2. Cluster telemetry (e.g., latency, requests per second) is normalized and appended to a system prompt that concisely describes the deployment and relevant call-path context. The aggregated prompt is provided

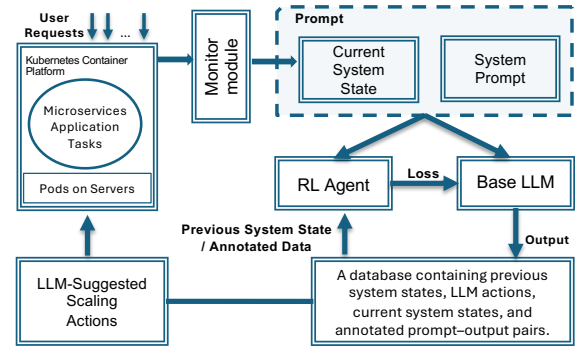


Fig. 2: ORACL System Architecture

to an LLM, which returns a recommended resource allocation or remediation action; the action is then executed on the cluster. In parallel, a reinforcement learning (RL) agent collects annotated prompt–output pairs together with historical and current system states. Each executed action induces a new system state, which we merge with the prior state to form a state transition. These transitions, along with the annotated prompts, constitute the training data used to fine-tune the base LLM, thereby closing the loop among diagnosis, decision making, and learning.

IV. SYSTEM DESIGN

In this section, we present a system design for our approach, organized according to the Monitor–Analyze–Plan–Execute framework over a shared knowledge base (MAPE-K), as illustrated in Fig. 3. The system comprises four primary components: (1) a tracing module that continuously observes and records system state; (2) an interface module that exposes essential capabilities the LLM cannot directly perform; (3) an LLM agent that aggregates contextual evidence to produce structured outputs and incorporates a reinforcement-learning module to support both online and offline learning.

A. Interface Module

Given the practical limitations of LLMs, the interface module orchestrates a three-stage workflow. **Data Collector** captures and normalizes system telemetry and call-path context. **Workload Predictor** forecasts near-term workload based on these observations. **Processor API** fuses the signals into a structured *system-state prompt* and delivers it to the LLM agent. The **Interaction Interface** then verifies the returned actions before execution.

1) *Data Collector*: LLMs have limited numerical sensitivity; for example, they may struggle to reliably distinguish small but important differences in metrics (e.g., CPU utilization of 80% vs. 85%). We therefore discretize continuous values into qualitative bins (e.g., *Low/Medium/High*) and encode derived efficiency indicators using the same scheme. LLMs also lack an intrinsic notion of healthy versus abnormal states, so we attach explicit thresholds or statistical baselines to guide reasoning. Finally, because LLMs cannot directly observe real-time behavior, we integrate external monitoring

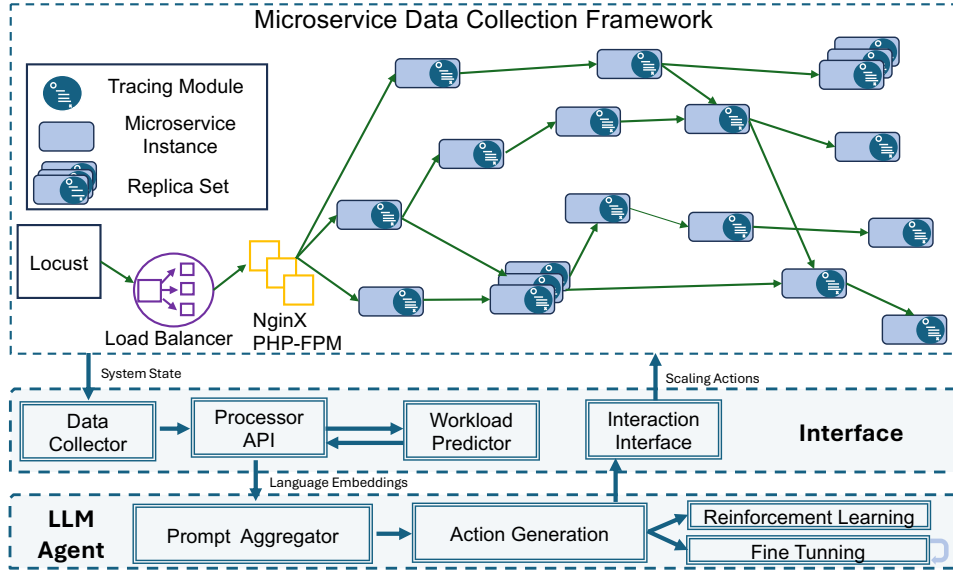


Fig. 3: ORACL System Prototype Diagram

tools (e.g., Prometheus for metrics and Jaeger for tracing¹) to provide structured inputs for LLM-based decision making.

2) *Workload Prediction Unit*: Real-world request streams exhibit strong temporal and spatial correlations, making workloads often predictable [38]. However, accurate time-series prediction remains challenging for general-purpose LLMs without substantial domain-specific fine-tuning. To avoid entangling numerical forecasting with higher-level reasoning, we used a dedicated forecasting model [38] for workload prediction and provided only the predicted workload, along with its confidence when available, as part of the prompt.

3) *Interaction Interface*: Because the LLM cannot directly manipulate cluster resources, all actions are mediated by an external execution API. Despite fine-tuning and reinforcement-learning-based optimizations, hallucinations and other high-confidence yet incorrect or conflicting outputs may still occur. To mitigate this risk, we apply a lightweight rule-based verification layer before execution, which rejects obviously unsafe decisions, such as reducing resources for already constrained or overutilized microservices. In our experiments, this checker is deployed in the production environment to filter evident hallucinations and preserve system reliability.

B. LLM Agent

This clause specifies the LLM Agent. The Agent comprises: (i) a Prompt Aggregation Module (PAM) that encodes objectives, constraints, and context into a structured prompt; (ii) an Action-Generation Module (AGM) that converts LLM outputs into executable actions for the interface; and (iii) a Reinforcement-Learning and Fine-Tuning module (RLFT) that implements a two-stage training procedure.

1) *Prompt Aggregator*: The PAM consolidates system objectives, temporal constraints, and user policies into prompts suitable for LLM reasoning. The prompts should (a) encode

| Symbol | Explanation |
|------------------|---|
| (C) (D) (M) (N) | CPU / Disk / Memory / Network-intensive, downstream (→) |
| (c) (d) (m) (n) | CPU / Disk / Memory / Network-intensive, upstream (←) |
| ① – ⑳ | Microservice identifiers |
| <think> | generate hypotheses |
| <Fault> | list possible causes |
| <Counterfactual> | effect if resources change |
| <root> | root cause: IDs and type |
| <Carrier> | effectiveness of resources change |

TABLE II: Resource and CoT symbol definitions

guidance policy and QoS/SLO constraints, (b) carry explicit microservice context, and (c) specify a machine-readable output schema to enable auditable and reproducible decisions.

- 1) *Initial Reasoning Guidance*: The PAM provides a reasoning template that directs the LLM to assess interdependent system states systematically before proposing scheduling actions. The template requires explicit intermediate considerations and specifies which aspects to cover during CoT reasoning. This structure may improve decision accuracy, stability, and interpretability by making each reasoning step traceable, especially in edge cases.
- 2) *Deployment Description*: For each deployment, the PAM accepts a natural-language descriptor that specifies its functional role, resource-demand profile, runtime behavior, and *readiness time* (the time to become operational). Readiness time guides orchestration decisions; services with long readiness times *should* be scaled out earlier to preserve responsiveness and efficiency.
- 3) *Call Graph Structure*: The call graph is central to reasoning about dependency structure and fan-in/fan-out. It exposes critical paths, bottlenecks, latency-sensitive paths, and failure propagation to inform resilience and resource allocation. We model static topology and dynamic traces in natural language with an attention mechanism and use

¹Prometheus: <https://prometheus.io>; Jaeger: <https://www.jaegertracing.io>

compact special tokens (Table II) to preserve structure and reduce space.

- 4) **Expected Generation Results:** The PAM defines a structured output schema; all annotations and pipeline outputs conform to it (Fig. 4). We treat RCA and resource allocation as distinct yet tightly coupled tasks that share evidence (metrics, traces, and faults) and reasoning. We use a four-stage pipeline per service per analysis window: (1) **Analysis:** summarize salient signals; (2) **Fault Identification:** list anomalies and candidate symptoms; (3) **Counterfactual Reasoning:** assess whether resource changes would mitigate degradation; (4) **Root-Cause Resolution:** record validated causes and remediations. Symbols are given in Table II.
- 5) **Cluster Running State:** As outlined in Section IV-A, this field is paired with a `Carrier` that encodes system-level sensitivity and specifies how the discretized state responds to load and configuration changes.

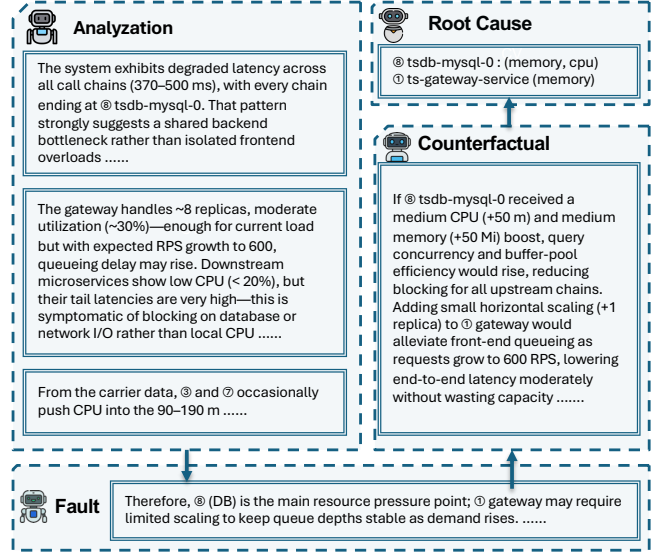


Fig. 4: An Expected Output From LLM

C. Action Generation

The action-generation module converts LLM outputs into actions that the interface can execute. For RCA, it parses root causes enclosed in `<root>` tags. For resource allocation, it uses the counterfactual reasoning segment to extract configuration changes, selects those it predicts will improve outcomes, and emits an executable action specification.

D. Reinforcement Learning and Fine-Tuning Module

Our framework optimizes the reasoning process rather than low-level control actions, enabling the model to learn high-level semantics and generalizable allocation rules. Supervised fine-tuning [39] trains on annotated prompt-output pairs to enforce the structured schema, encode practical engineers’ heuristics, and reasoning patterns. The reinforcement learning component operates in two distinct stages: (i) root-cause identification, optimized with causal analysis and validity checks; and (ii) prediction of whether proposed counterfactual changes achieve the intended system behavior.

V. SYSTEM MODELING

This section defines the problem formulation and optimization objective, describes system-wide operation, and presents the algorithmic design used to achieve the objective.

A. Problem Formulation

In ORACL, we formalize resource management for microservices running on a cluster of physical machines. Let $M = \{m_1, m_2, \dots, m_i\}$ denote the set of machines. For each machine $m_i \in M$, the allocatable resources are represented as $R_{m,j} = \{r_{m,1}, r_{m,2}, \dots, r_{m,j}\}$, where each $r_{m,j}$ denotes a resource type j such as CPU or memory. The set of feasible actions is $A_{m,i} = \{a_{m,1}, a_{m,2}, \dots, a_{m,i}\}$, where each $a_{m,i}$ encodes a scaling operation i (horizontal scaling, CPU vertical scaling, or memory vertical scaling) on machine m . We denote the textual input provided to the system as *Prompt*, and the generated response sequence of the LLM as (g_1, g_2, \dots, g_T) .

Our methodology consists of two interconnected sub-problems. The first is **root cause identification**, which we treat as a classification problem. Given input signals *Prompt*, the model predicts $\hat{y} = f_\theta(\text{Prompt})$, aiming to maximize the probability $\mathbb{P}(\hat{y} = y^* \mid \text{Prompt})$, where y^* is the true underlying cause.

The second sub-problem is **resource allocation**, which we formulate as a Markov Decision Process (MDP). At each step t , the system observes a state $s_t \in \mathcal{S}$, takes an action $a_t \in \mathcal{A}$, and transitions to a new state s_{t+1} . The objective is to maximize the QoS, which can be naturally achieved by letting the LLM perform proactive counterfactual reasoning: given a hypothetical intervention \tilde{a}_t , which normally corresponds to the action suggested after identifying a potential problem, the model predicts a future state $\hat{s}_{t+1} = f_\theta(s_t, \tilde{a}_t)$ and seeks to maximize the alignment $\mathbb{P}(\hat{s}_{t+1} \approx s_{t+1} \mid s_t, \tilde{a}_t)$. This design ensures that “what-if” predictions guide scaling actions while maintaining consistency with observed system performance.

In summary, ORACL decomposes the learning problem into two stages that share the same chain-of-thought policy π , whose conceptual form is expressed in Eq. (1):

$$\pi^* \propto \underbrace{\mathbb{P}(\hat{s}_{t+1} = s_{t+1} \mid \text{Prompt})}_{\text{counterfactual consistency}} \underbrace{\mathbb{P}(\hat{y} = y^* \mid \hat{s}_{t+1}, \text{Prompt})}_{\text{root cause classification}}. \quad (1)$$

B. System-wide Actions

Operation procedure: In ORACL, the dual-objective optimization problem is formulated as a unified system that operates in two distinct modes. As shown in Algorithm 1, we first initialize the base LLM (lines 1–2). When offline training is required, the procedure follows Algorithm 2 (lines 3–5). After the offline training phase, if the task is labeled as an RCA problem, the current system state—either retrieved from historical records or collected from an online cluster—is transformed by the prompt aggregator described in the previous

Algorithm 1: System-wide Operation Procedure

Input: Flags: *Required_offline_Training*, *Required_Online_Training*;
 Data instance: *data_instance*; Task type: *Task_type*;
 Base model: *base_LLM*;
Output: System-level output (e.g., root cause or scheduling actions)

```

1 Initialization:
2 Load base_LLM;
3 if Required_offline_Training then
4   Offline_Training(data_instance, base_LLM);
5 end;

6 // Root-Cause Identification
7 if Task_type = "Root_cause_Identification" then
8   system_state ← get_cluster_state();
9   if system_state is None then
10    system_state ← data_instance;
11  end;
12  prompt ← system_state;
13  Root_cause ←
14    base_LLM.generate(prompt).extract_between_tags();
15  return Root_cause;
16 end;

17 // Dynamic Allocation Loop
18 if Task_type = "Allocation" then
19   while True do
20     Online_Training(base_LLM,
21       Required_Online_Training);
22     sleep(interval_time);
23   end;

```

section. The aggregated prompt is then passed to the LLM (lines 6–12). The LLM produces outputs structured into four segments corresponding to the CoT reasoning process, from which the root causes are extracted by parsing the contents enclosed in the <root> tags (lines 13–15). In contrast, when the task is specified as scheduling (lines 16–18), the LLM manager enters a continuous loop in which it proactively allocates resources according to Algorithm 3, followed by a sleep interval (lines 20–21). In practical deployments, multiple LLM managers may run in parallel to provision resources for pods associated with the same call graph, thereby improving system-level efficiency. The two algorithms referenced above are elaborated in subsequent sections.




Fig. 5: System-wide operation with a Dual Objective Procedure

Optimization targets: As illustrated in Fig. 5, the RCA objective is to accurately identify all faulty components (blue), with precision as the primary evaluation metric. For resource allocation, only counterfactual configurations (red) with positively predicted outcomes are retained to ensure consistency with the CoT-based reasoning process. By aligning counterfac-

Algorithm 2: ORACL: Offline Training

Input: Annotated dataset: \mathcal{D} ; Base model: *base_LLM*;
 Number of supervised steps: *SFT_steps*; Number of GSPO steps: *GSPO_steps*
Output: Updated model parameters θ

```

1 Initialization:
2 Split  $\mathcal{D}$  into (data_partition_0, data_partition_1) using
   Annotated_Data.split();
3 Initialize optimizer and learning scheduler for base_LLM;
4 // Supervised Fine-Tuning (SFT)
5 for step ← 1 to SFT_steps do
6   logits ← base_LLM(data_partition_0.inputs);
7   loss ←
8     CrossEntropy(logits, data_partition_0.response);
9   update(base_LLM, loss);
10 // GSPO Fine-Tuning
11 for step ← 1 to GSPO_steps do
12   out ←
13     base_LLM.generate(data_partition_1.prompts);
14   [Observation, Fault, Counterfactual, Root_cause] ←
15     out.extract_between_tags();
16   loss ←
17     Overall_Loss(Root_cause, data_partition_1.root_cause);
18   update(base_LLM, loss);

```

tual selection with accurate system-response predictions (yellow), the resulting decisions naturally drive improvements in overall system performance. Discrepancies between predicted and observed outcomes are further fed back to refine the reasoning process and enhance robustness over time [40].

C. Root-Cause Identification

Algorithm 2 illustrates the offline training procedure for RCA. The annotated dataset, formatted as described in the previous section, is first divided into two partitions (lines 1–2), and the optimizer, along with the base LLM parameters, is initialized (line 3). The first partition is used to enable the model not only to retain its pretrained knowledge but also to adapt to the reasoning patterns commonly employed by cloud engineers when diagnosing faults and proposing solutions.

$$\mathcal{L}(\theta) = - \sum_{t=1}^T \log P_{\theta}(y_t | y_{<t}) \quad (2)$$

Subsequently, the entropy loss is calculated as defined in (2), and the model parameters are updated accordingly (lines 5–8). The model then proceeds to the critical stage, where reinforcement learning is applied to optimize the CoT process itself. In this stage, the LLM generates four reasoning steps, and the outputs are evaluated by a reward function in conjunction with the annotated data. The design and objective of this reward function are discussed in detail in Section V-C1.

1) *Objective:* The **primary objective** is to maximize the probability of correctly identifying the root cause in Eq. (3):

$$\pi^* = \arg \max_{\pi} \mathbb{E} \left[P_{\pi}(t_{n+1}, t_{n+2}, \dots, t_m \mid t_1, t_2, \dots, t_n) \right]. \quad (3)$$

where

$$t_1, \dots, t_n \in \mathcal{T}_{\text{reason}}, \quad t_{n+1}, \dots, t_m \in \mathcal{T}_{\text{rc}},$$

denote the reasoning token sequence and the root cause token sequence, respectively.

The **secondary objective** is to maximize the probability of uncovering false positives during reasoning as shown in Eq. (4):

$$\max_{\pi} \mathbb{P}_{\pi} \left(t_1^{(\text{cf})}, \dots, t_n^{(\text{cf})} \mid t_1^{(\text{cand})}, \dots, t_m^{(\text{cand})} \right), \quad (4)$$

where $t_{1:m}^{(\text{cand})} \in \mathcal{T}_{\text{cand}}$ denote the candidate root-cause tokens already generated, and $t_{1:n}^{(\text{cf})} \in \mathcal{T}_{\text{cf}}$ denote the subsequent counterfactual reasoning tokens that may revise or reject earlier candidates.

This step ensures that the model not only identifies faulty components with high accuracy but also validates its own reasoning by aligning counterfactual predictions with observed system behavior. In practice, this coupling reduces hallucinations in LLM outputs, improves the robustness of root cause identification, and enhances the reliability of subsequent resource allocation decisions, thereby aligning learning objectives with the overarching goals of SLO compliance and efficient cloud resource use.

2) *Reward Functions*: The reward functions address correctness, structural properties, length, and the model’s over-shifting behavior, as described below.

a) *Correctness*: The correctness score comprises three components: pod-level accuracy, root-cause correctness, and a thresholded penalty that discourages excessive guessing.

Pod-level match (guess reward). The pod-level match score S_{pod} is defined as

$$S_{\text{pod}} = \frac{\text{POD}_{\text{match}}}{\text{POD}_{\text{match}} + \text{POD}_{\text{not}} + \epsilon}, \quad (5)$$

where $\text{POD}_{\text{match}}$ and POD_{not} denote the numbers of correctly and incorrectly predicted pods, respectively, and $\epsilon > 0$ is a small constant for numerical stability. As shown in Eq. (5), this term rewards correctly guessed faulty pods irrespective of the specific CPU/Memory/Disk dimension.

Root-cause correctness. Given the numbers of true positives, false positives, and false negatives, denoted by TP, FP, and FN, respectively, the root-cause correctness is measured by the F_{β} score:

$$F_{\beta} = \frac{(1 + \beta^2) \text{TP}}{(1 + \beta^2) \text{TP} + \beta^2 \text{FN} + \text{FP} + \epsilon}, \quad \beta > 1, \quad (6)$$

where $\beta > 1$ biases the metric toward recall to reduce missed root causes. As in Eq. (6), false negatives are weighted more heavily than false positives.

Base reward. The base reward R_{base} combines the pod-level match score in Eq. (5) and the root-cause correctness in Eq. (6):

$$R_{\text{base}} = \alpha S_{\text{pod}} + (1 - \alpha) F_{\beta} + \delta \cdot \mathbb{I}[\text{TP} > 0], \quad 0 < \alpha \ll 1, \delta > 0. \quad (7)$$

where α is a small weight on the pod-level term and δ provides an additional bonus whenever at least one true positive is identified. Equation (7) thus encourages the model to correctly localize the faulty pod while emphasizing root-cause correctness.

Threshold penalty. To discourage missing too many true root causes, a thresholded penalty is applied to the base reward in Eq. (7). Let $\tau_{\text{fn}} \in \mathbb{N}$ be the tolerance threshold for false negatives. The final correctness reward R_{Result} is given by

$$R_{\text{Result}} = R_{\text{base}} - d \cdot \max(0, \text{FN} - \tau_{\text{fn}}), \quad (8)$$

where $d > 0$ controls the penalty strength. When $\text{FN} \leq \tau_{\text{fn}}$, Eq. (8) reduces to the base reward in Eq. (7). In contrast, false positives FP are only softly penalized through the F_{β} term in Eq. (6), and no explicit threshold is imposed on them.

b) *Format Compliance*: The objective is to ensure that the LLM correctly outputs the required tags. If all tag pairs are valid, no penalty is applied. Otherwise, each violation reduces the representative score. Formally, let C_{total} denote the total number of checks and C_{invalid} the number of checks not passed. We define:

$$R_{\text{format}} = -\frac{C_{\text{invalid}}}{C_{\text{total}}}, \quad (9)$$

where each of the following counts as a failed check:

- 1) Empty reasoning content between a start and end tag,
- 2) Incorrect ordering of start and end tags,
- 3) Nested tags inside another tag’s scope.

Thus, the final score decreases proportionally to the number of violations.

c) *Length Constraint with Smooth Decay*: To balance correctness and inference efficiency, we constrain the model’s reasoning length within a statistically determined range $[L_{\text{min}}, L_{\text{max}}]$. Although these bounds are derived from empirical measurements, the model typically requires more explicit reasoning steps and exploratory attempts than an expert human engineer. Consequently, L_{max} is intentionally set to exceed the typical annotation length produced by human experts. As shown in Eq. (10), answers shorter than L_{min} receive the maximum reward; those with lengths between L_{min} and L_{max} are linearly down-weighted; and any answer exceeding L_{max} obtains zero reward.

$$s(L) = \begin{cases} 0, & L \leq L_{\text{min}}, \\ -\frac{L - L_{\text{min}}}{L_{\text{max}} - L_{\text{min}}}, & L_{\text{min}} < L < L_{\text{max}}, \\ -1, & L \geq L_{\text{max}}. \end{cases} \quad (10)$$

Formally, let $\text{correct}(c) \in \{0, 1\}$ be the correctness indicator, the final reward is depicted in the (11):

$$R_{\text{length}}(c) = \text{correct}(c) \cdot s(L(c)). \quad (11)$$

d) *KL Regularization*: To prevent the new policy from diverging excessively from the reference model and thereby losing the embedded expertise of cloud engineers, a KL penalty term is introduced, as defined in Eq. (12).

$$R_{\text{KL}} = -\beta D_{\text{KL}}(\pi_{\theta} \parallel \pi_{\text{old}}), \quad (12)$$

where π_θ denotes the new policy, π_{old} represents the reference policy from the previous step, and $\beta > 0$ controls the strength of the regularization.

e) *Final Reward*: The overall reward R_{total} is defined as the sum of all the aforementioned components, as shown in Eq. (13):

$$R_{\text{total}} = R_{\text{result}} + R_{\text{format}} + R_{\text{length}} + R_{\text{KL}}. \quad (13)$$

3) *With GSPO*: As the main issue is the retrieval of datapoints, we allow our LLM to generate multiple instances each round by deploying Group Sequence Policy Optimization (GSPO) [41] with the objective function shown in Eq. (14):

$$J_{\text{GSPO}}(\theta) = \mathbb{E} \left[\frac{1}{G} \sum_{i=1}^G \min \left(s_i(\theta) \hat{A}_i, \text{clip}(s_i(\theta)) \hat{A}_i \right) \right], \quad (14)$$

Here, each prompt x is sampled from the dataset \mathcal{D} , and $\{y_i\}_{i=1}^G$ denotes the G candidate outputs generated by the model. The advantage \hat{A}_i for the i -th output is computed by normalizing its total reward $R_{\text{total}}(x, y_i)$ with the group mean and standard deviation.

$$s_i(\theta) = \left(\frac{\pi_\theta(y_i | x)}{\pi_{\theta_{\text{old}}}(y_i | x)} \right)^{\frac{1}{|y_i|}} \quad (15)$$

where the ratio compares sequence-level likelihoods [42] under the new policy π_θ and the old policy $\pi_{\theta_{\text{old}}}$, and the geometric mean (the $|y_i|$ -th root) helps prevent instability on long sequences, as shown in Eq. (15). Finally, the clipping function $\text{clip}(\cdot)$ bounds $s_i(\theta)$ within $[1 - \epsilon, 1 + \epsilon]$ to avoid overly large policy updates.

D. Resource Allocation

Algorithm 3: ORACL: Scheduling with Optional Online Training

Input: Base model *base_LLM*;
 Online training flag *Required_Online_Training*
Output: Updated model *base_LLM* (if online training enabled)

```

1 Initialization:
2 carrier ← state_gathering_phase(previous_system_state);
3 previous_system_state ← get_cluster_state();
4 Prompt Construction and Inference:
5 prompt ←
   compose_prompt(previous_system_state, carrier);
6 output ← base_LLM.generate(prompt);
7 [Observation, Fault, Counterfactual, Root_cause] ←
   output.extract_between_tags();
8 Counterfactual Action Execution:
9 apply_action(Counterfactual);
10 Online Adaptation (Optional):
11 if Required_Online_Training then
12   new_system_state ← get_cluster_state();
13   Loss ←
     is_Expected(Counterfactual, new_system_state);
14   base_LLM.step(Loss);

```

The orchestration procedure consists of an inference phase and an optional online training phase. As shown in Algorithm 3, because the base LLM does not have an explicit notion

of a “good” system state, we prepend a brief statistical state-gathering phase [26] that characterizes how sensitive the system is to configuration changes. Using a lightweight clustering technique [43], historical system states are summarized into a compact representation, referred to as the *carrier* (lines 1–2). The carrier is then prepended to the current cluster state, and the combined information is encoded into a prompt (lines 3–5).

The LLM prediction is parsed by an interpreter, defined in Eq. (16), to extract the counterfactual action that is expected to lead to a positive change in the system:

$$\hat{o}, a = \pi_{\text{Interpreter}}(\text{Counterfactual}), \quad (16)$$

where $\pi_{\text{Interpreter}}(\cdot)$ can be instantiated as either another LLM or a rule-based procedure. The resulting action a is then applied to the managed application (lines 6–9).

If online training is enabled, the true outcome o is observed from the cluster after executing a scaling action a (lines 10–12). We then compute a simple binary loss, analogous to that used in DeepSeek-Math (lines 13), as given in Eq. (17). Specifically, the loss is set to 1 if the predicted outcome \hat{o} matches the observed outcome o , and 0 otherwise:

$$\mathcal{L}(\hat{o}, o) = \begin{cases} 1, & \text{if } \hat{o} = o, \\ 0, & \text{otherwise.} \end{cases} \quad (17)$$

Finally, the base LLM is updated using this loss signal via a single optimization step (line 14).

VI. PERFORMANCE EVALUATION

To assess the efficacy of ORACL for both RCA and autoscaling in microservice environments, we conduct experiments on a realistic Kubernetes-based testbed and on open-source datasets. We detail the experimental settings, benchmarks, and provide a comprehensive analysis of the results. The primary objective is to demonstrate that these tasks can be unified within a single framework and that the CoT-based method is highly adaptable without extensive training, thereby suggesting promising directions for future research.

A. Experimental Setup

The prototype is deployed on a five-node Kubernetes cluster consisting of one master and four worker nodes, each equipped with 8 vCPUs and 8 GB of memory, with cgroup v2 enabled. Two additional GPU nodes are used for training and inference: one with four NVIDIA A100 GPUs and another with eight NVIDIA RTX 6000 Ada GPUs. We use Qwen3-32B as the model backend, and implement model training with PyTorch, Unsloth², and DeepSpeed ZeRO-3 to enable efficient distributed training with quantization-based optimization. RCA experiments are conducted using the Murphy dataset [16] on the Hotel Reservation benchmark [44], while autoscaling evaluations are performed on Sock-Shop [45] and Train-Ticket [44]. Workloads are generated using a distributed Locust setup³ following traffic patterns derived from the Alibaba trace [38]. Performance metrics and traces are collected through Prometheus, cAdvisor, and Jaeger.

²<https://github.com/unslothai/unsloth>

³<https://locust.io>

B. Baselines

Given the two tasks in this work, we evaluate our model against two groups of baselines. **Root-cause detection.** We compare with four representative systems: ExplainIT [46] (correlating observed-problem metrics with candidate-root-cause metrics), NetMedic [47] and Sage [48] (dependency-graph diagnosis with correlation-based edge weights and heuristics), and Murphy [16] (MRF-based diagnosis over loosely associated metrics). We also evaluate our constructed prompts on three prompt-only LLM APIs: ChatGPT-4o [11], Deepseek-V3 [37], and Claude-3.5-Haiku [49].

Resource allocation. For autoscaling, we consider three baselines: KuScal [50], a threshold-based Kubernetes scaler with a 75% CPU target as suggested in prior work [51]; CoScale [52], a correlation-driven Q-learning autoscaler with GRU-based workload prediction and support for horizontal/vertical scaling and brownout; and FIRM [22], a model-based controller that allocates resources along critical microservice paths to reduce end-to-end latency. In addition to these external baselines, we include ORACL (Raw), an ablated version of our framework without RCA training, to isolate the contribution of the RCA-enhanced reasoning process.

C. Root Cause Identification

To evaluate the effectiveness of our CoT design and demonstrate the benefits of our training strategy, we first compare ORACL against several state-of-the-art RCA systems. We then apply the prompts generated by our interface to a set of commercial large-scale models, each with at least an order of magnitude more parameters, to assess the robustness and generalization of our training method.




Fig. 6: RCA Comparison

1) *Comparison with State-of-the-Art RCA Models:* In Fig. 6, we evaluate ORACL against four representative RCA frameworks, including ExplainIT, NetMedic, Sage, and Murphy on the original datasets released with Murphy, using Recall, Precision, and Accuracy. For a fair comparison, we adopt a strict protocol in which only the single most confident prediction of each method is counted.

Traditional statistical and rule-based baselines (ExplainIT, NetMedic) exhibit limited diagnostic capability and fail to capture cyclic failure conditions [16], while Sage shows

unstable precision and incomplete recall. Murphy, a recent learning-based method, attains slightly higher recall than ORACL but substantially lower precision and accuracy under topology churn. In contrast, ORACL achieves 0.83 recall, 0.89 precision, and 0.92 accuracy, combining high anomaly sensitivity with a low false-positive rate and demonstrating robust generalization across diverse microservice conditions.

2) *Comparison with Commercial LLMs:* Table III summarizes the comparative performance of ORACL and three state-of-the-art commercial LLMs under the default ORACL configuration. Specifically, we compare ORACL against DeepSeek-V3, ChatGPT-4o, and Claude 3.5 Haiku. Public reports indicate that DeepSeek-V3 comprises approximately 671 billion parameters, while ChatGPT-4o is widely believed to operate in a similar trillion-scale regime. In contrast, Claude 3.5 Haiku is explicitly positioned as a lightweight model with a substantially smaller parameter budget, although exact figures are not publicly disclosed. Despite using significantly fewer parameters than all three commercial systems, ORACL consistently achieves superior performance in our setting. Quantitatively, ORACL improves precision by approximately 3%–28%, recall by 6%–38%, and accuracy by 4%–15% over the commercial baselines. For example, compared to ChatGPT-4o, ORACL attains gains of +6% in precision, +6% in recall, and +5% in accuracy. Overall, ORACL achieves the highest precision, recall, and accuracy among all evaluated models.

TABLE III: Performance comparison across commercial models

| Model | Precision | Recall | Accuracy |
|------------------|-------------|-------------|-------------|
| Claude-3.5-Haiku | 0.61 | 0.45 | 0.77 |
| ChatGPT-4o | 0.83 | 0.77 | 0.87 |
| DeepSeek-V3 | 0.86 | 0.72 | 0.88 |
| ORACL | 0.89 | 0.83 | 0.92 |

D. Scheduling Evaluations

We evaluate ORACL under two research questions: (1) its performance relative to existing approaches under limited training, and (2) its performance when competing models are fully converged. Experiments are conducted in two scenarios, where the 500-minute and 5000-minute test traces are each divided into five equal intervals, and the reported results are averaged per interval. ORACL performs a single data-collection phase of approximately 1–2 hours (Section V-D) and requires no further online retraining.

The first scenario evaluates Sock-shop with only 24 hours of training data. The second considers the larger Train-ticket system, where all models are allowed to fully converge. Due to its substantially larger call graph, Train-ticket is handled by multiple LLM-agent replicas operating on different subgraphs; When conflicting add/remove decisions arise for the same resource, the system conservatively prioritizes the add action. With inference model batching and a simple implementation, the overhead is negligible.

1) *Requests-per-Second (RPS) Analysis:* To evaluate the system’s throughput and the scheduler’s ability to effectively




Fig. 7: Comparison of Sock-Shop performance metrics.

utilize available resources, we analyze requests per second (RPS), which measures how many requests the system can process per second.

Sock-shop (limited training). Fig. 7a reports the RPS of the four methods on Sock-shop over five time periods. ORACL consistently achieves the highest throughput, between 3500 and 3670 requests/s, delivering about 5% more requests/s than CoScal and over 20% more than FIRM, while KuScal lags far behind. In addition, ORACL’s throughput is also the most stable over time, whereas FIRM shows noticeable oscillations under previously unseen states, and KuScal is consistently under-provisioned. With only one day of training data, the LLM-based ORACL scheduler already achieves both higher and more stable throughput than all baselines. The ablation result of ORACL(Raw) further supports the findings from prior works in Section II and highlights the benefit of RCA training: without RCA-guided training, the model may fail to identify the correct resource-allocation actions, leading to random performance fluctuations across the five time periods; therefore, the remaining experiments focus on the full ORACL framework.

Train-ticket (fully converged). Fig. 8a presents the RPS results on Train-ticket when all baseline methods are trained to convergence. FIRM now achieves the highest average throughput of around 750 requests/s, with ORACL and CoScal closely following at around 730 requests/s. Compared with FIRM, ORACL incurs only about 2% lower throughput, while still improving throughput over KuScal by more than 30%. The small gap between ORACL and the best-converged deep learning model suggests that ORACL generalizes well to a more complex application, even though it does not rely on long-term online training.

2) **Average Response Time:** To evaluate the system’s stability in handling real workloads with minimal failures, we further compare the average end-to-end response time across the same five periods.

Sock-shop (limited training). Fig. 7b shows that, on Sock-shop, ORACL consistently achieves the lowest latency across all five periods. Excluding the initial warm-up interval, where FIRM spikes to nearly 600ms, ORACL attains a steady-state average of 162.3ms, compared with 234.0ms for CoScal, 205.5ms for FIRM, and 268.5ms for KuScal. This corresponds to a latency reduction of approximately 31%, 21%, and 40%,

respectively. The pronounced first-period spike indicates that FIRM is sensitive to previously unseen workload patterns with limited training, whereas ORACL maintains low, stable latency after only a short offline data-collection phase.

Train-ticket (fully converged). The response-time results for Train-ticket are summarized in Fig. 8b. All three advanced methods (FIRM, CoScal, and ORACL) achieve low latencies on the order of tens of milliseconds. Compared with the best deep-learning baseline (FIRM), ORACL incurs about 10% higher latency and only a few percent overhead over CoScal, but still reduces latency by roughly 65% over KuScal. Looking at the per-period behavior, ORACL achieves the lowest or comparable latency in the first three periods, where the workload is relatively stable. In the last two periods, its latency increases more noticeably, while FIRM and CoScal remain slightly lower. This suggests that, in the presence of long-running shifts in user behavior or carrier composition, the converged numerical models can capture fine-grained patterns more accurately, whereas ORACL trades some optimality for its lightweight training.

3) **Cumulative Distribution Function (CDF) of Response Time:** To evaluate the effectiveness of resource usage and better characterize system behavior, we further analyze the response-time CDFs to reveal how each scheduler behaves across the full latency distribution.

Sock-shop (limited training). Fig. 7c shows the response-time CDF on Sock-shop. ORACL achieves the tightest distribution, reducing latency by about 30–50% at the 60th–80th percentiles and by 15–60% at the tail (95th–99th percentiles) relative to the baselines. Among the non-ORACL baselines, FIRM performs best at lower quantiles but exhibits the heaviest tail, indicating instability under unseen workloads, while CoScal and KuScal display smoother yet consistently higher latencies. These results indicate that ORACL effectively stabilizes tail behavior under limited training by prioritizing QoS-critical services.

Train-ticket (fully converged). Fig. 8c presents the response-time CDFs on Train-ticket. ORACL clearly dominates at low percentiles, with median and 66th-percentile latencies reduced by roughly 50–67% relative to FIRM and by over 80% relative to CoScal, indicating its strong effectiveness in optimizing the critical path for short requests. ORACL remains competitive around the 75th percentile and, although FIRM and CoScal




Fig. 8: Comparison of Train-Ticket performance metrics.

TABLE IV: CDF Comparison in Train-ticket

| System | 50th | 66th | 75th | 80th | 90th | 95th | 99th | 99.99th |
|--------|------|------|------|------|------|------|------|---------|
| FIRM | 9 | 10 | 12 | 15 | 45 | 120 | 360 | 2900 |
| CoScal | 22 | 36 | 51 | 63 | 95 | 180 | 510 | 1800 |
| KuScal | 240 | 660 | 960 | 1100 | 1700 | 2300 | 3900 | 7900 |
| ORACL | 3 | 5 | 11 | 26 | 140 | 260 | 460 | 2700 |

slightly outperform it in the mid-to-high range (80th–95th percentiles) after full convergence, ORACL matches FIRM near the tail and significantly outperforms KuScal, reducing worst-case latency by about two-thirds. This behavior stems from ORACL’s design: the LLM-based scheduler prioritizes reasoning about the root causes of QoS violations over aggressively reclaiming resources from all services, thereby preserving capacity for non-critical microservices. While this conservative strategy can yield marginally higher mid-percentile latency, it consistently delivers strong median performance and substantially improved tail behavior.

4) **Summary:** Overall, the threshold-based CoScal provides reasonable performance and convergence speed, but its coarse state discretization and limited action space prevent it from fully exploiting the available resources, making it the most “unremarkable” method across all ten periods. The LLM-based, chain-of-thought ORACL scheduler is the most adaptive when the application or workload changes: in both Sockshop and Train-ticket, it maintains stable behavior without per-application hyperparameter tuning and achieves strong throughput and latency with only a short data-collection phase.

Deep-learning-based FIRM, on the other hand, benefits from large amounts of training data. Under limited training (Sock-shop), it suffers from instability and occasional severe latency spikes, but once fully converged on a complex workload (Train-ticket), it can slightly surpass ORACL in both throughput and latency. In summary, ORACL offers a favorable trade-off: it avoids the long training time and instability of deep-learning approaches while achieving near state-of-the-art performance on a complex microservice system and significantly outperforming Kubernetes-based baselines.

E. Attention Mechanism Analysis

To examine the model’s internal reasoning and interpretability, we analyze token-wise attention distributions across early,

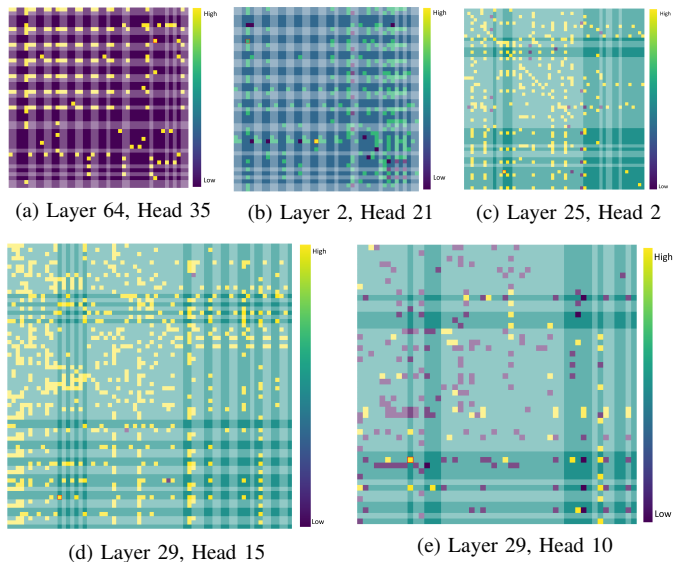


Fig. 9: Attention Head Visualization Across Different Layers.

middle, and late layers under different system-state inputs. Tokens that contribute more to the final prediction are highlighted with higher color intensity. Although exhaustive visualization of all reasoning tokens is infeasible due to the large context size of LLM inference, representative attention patterns indicate that the CoT-based reasoning process captures both optimization-relevant signals and root-cause-related evidence. Instead of exhaustively listing all tokens, we qualitatively summarize the observed irregular attention patterns and correlated token groups to provide a general understanding of the model’s reasoning behavior.

In the early layers, as depicted in Fig. 9b, the model primarily focuses on capturing local contextual and syntactic relations. At this stage, the attention patterns mainly highlight tokens that represent resource-related words, reflecting how the model learns word-level and adjacent dependencies, similar to conventional NLP models. Since the special tokens introduced in our setup carry limited semantic meaning, it is reasonable to observe lower attention scores on the left side of the figure.

The middle layers suggest dependency-aware representations, in our case, that is, the causal patterns between microservices, assigning closer attention to token pairs repre-

senting pod-level interactions on the same call graph. Metric-behavior alignment is also evident, as shown by the red-highlighted regions corresponding to Front-end-Memory (Fig. 9e), Rate-CPU (Fig. 9d), and Rate-Memory (Fig. 9c).

In the later stages, particularly in the final layer (as shown in Fig. 9a), the model focuses more on the structural patterns of the predefined templates to generate plans interpretable by our framework. Moreover, as the LLM is optimized for next-token prediction, some token pairs exhibit elevated attention beyond structural alignment. While such peaks may appear to reflect causal reasoning, they more likely indicate output alignment toward correct root-cause predictions rather than explicit causal analysis.

VII. CONCLUSIONS AND FUTURE WORK

Our proposed ORACL system unifies root-cause detection and resource allocation tasks within a single framework. Compared with existing techniques, it requires significantly shorter training time while still achieving competitive performance in practical settings. Furthermore, most of the workflow is specified at a high level, which reduces the overall coding effort. While our results demonstrate the feasibility of optimizing chain-of-thought reasoning for resource allocation and root cause finding, there exist several opportunities for further improvement as discussed below:

- **Inference Efficiency.** Although inference is generally fast, the computational cost and variable latency of LLMs can affect responsiveness under frequent requests or workload bursts. In some experiments, inference exceeded the allocated time window, causing transient stalls. Future work will explore domain-adapted models and optimized serving infrastructures to reduce cost, stabilize latency, and improve peak-load robustness.
- **Multi-Agent Negotiation.** In the current implementation, conflicts among multiple LLM-agent replicas are resolved using a simple rule-based strategy. Future work could investigate more fine-grained negotiation and coordination mechanisms to support scalable LLM-agent deployment and more effective multi-agent resource management.
- **Granularity of Analysis.** The current data set covers only a limited set of microservice architectures and deployment environments, and reasoning is conducted mainly at the service and resource levels. Future work could use RAG [53] to integrate time-aware telemetry, node-level scheduling, and code-level efficiency signals for finer-grained diagnosis and more precise allocation decisions.

Despite these limitations, our experiments suggest that direct CoT optimization is a promising direction. With richer data, tighter integration with cloud-native tooling, and improved model architectures, LLM-based agents could evolve into autonomous controllers for cloud resource management.

REFERENCES

- [1] R. Mahmud, S. Pallewatta, M. Goudarzi, and R. Buyya, "Ifogsim2: An extended ifogsim simulator for mobility, clustering, and microservice management in edge and fog computing environments," *Journal of Systems and Software*, vol. 190, p. 111351, 2022.
- [2] Alibaba Cloud, "Alibaba cloud documentation," <https://www.alibabacloud.com/help>, 2023, accessed: 2025-02-01.
- [3] M. Xu, L. Yang, Y. Wang, C. Gao, L. Wen, G. Xu, L. Zhang, K. Ye, and C. Xu, "Practice of alibaba cloud on elastic resource provisioning for large-scale microservices cluster," *Software: Practice and Experience*, vol. 54, no. 1, pp. 39–57, 2024.
- [4] Amazon Web Services, "Overview of amazon web services," <https://docs.aws.amazon.com/whitepapers/latest/aws-overview/aws-overview.pdf>, 2023, accessed: 2025-02-01.
- [5] X. Yang, S. M. Blackburn, and K. S. McKinley, "Elfen scheduling: Fine-Grain principled borrowing from Latency-Critical workloads using simultaneous multithreading," in *2016 USENIX Annual Technical Conference (USENIX ATC 16)*. Denver, CO: USENIX Association, Jun. 2016, pp. 309–322. [Online]. Available: <https://www.usenix.org/conference/atc16/technical-sessions/presentation/yang>
- [6] T. Zhang, H. Qiu, G. Castellano, M. Rifai, C. S. Chen, and F. Pianese, "System log parsing: A survey," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 8, pp. 8596–8614, 2023.
- [7] J. Jiang, W. Lu, J. Chen, Q. Lin, P. Zhao, Y. Kang, H. Zhang, Y. Xiong, F. Gao, Z. Xu *et al.*, "How to mitigate the incident? an effective troubleshooting guide recommendation technique for online service systems," in *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2020, pp. 1410–1420.
- [8] M. Xu, L. Wen, J. Liao, H. Wu, K. Ye, and C. Xu, "Auto-scaling approaches for cloud-native applications: A survey and taxonomy," *arXiv preprint arXiv:2507.17128*, 2025.
- [9] F. Y. Yan, H. Ayers, C. Zhu, S. Fouladi, J. Hong, K. Zhang, P. Levis, and K. Winstein, "Learning in situ: a randomized experiment in video streaming," in *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, 2020, pp. 495–511.
- [10] S. Shen, Z. Yao, L. Shi, L. Wang, L. Lai, Q. Tao, L. Su, R. Chen, W. Yu, H. Chen *et al.*, "Bridging the gap between relational {OLTP} and graph-based {OLAP}," in *2023 USENIX Annual Technical Conference (USENIX ATC 23)*, 2023, pp. 181–196.
- [11] OpenAI, "Gpt-5 system card," <https://cdn.openai.com/pdf/8124a3ce-ab78-4f06-96eb-49ea29ffb52f/gpt5-system-card-aug7.pdf>, 2025, model release announcement and system card for GPT-5.
- [12] Q. Team, "Qwen3 technical report," <https://arxiv.org/abs/2505.09388>, 2025, large language model series, Alibaba.
- [13] Anthropic, "Claude 4 (opus & sonnet) model family," <https://www.anthropic.com>, 2025.
- [14] F. He, T. Zhu, D. Ye, B. Liu, W. Zhou, and P. Yu, "The emerged security and privacy of llm agent: A survey with case studies," *ACM Computing Surveys*, 2024.
- [15] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," *Advances in neural information processing systems*, vol. 35, pp. 24824–24837, 2022.
- [16] V. Harsh, W. Zhou, S. Ashok, R. N. Mysore, B. Godfrey, and S. Banerjee, "Murphy: Performance diagnosis of distributed cloud applications," in *Proceedings of the ACM SIGCOMM 2023 Conference*, 2023, pp. 438–451.
- [17] N. Pham *et al.*, "Baro: Robust root cause analysis for microservice failures," in *Proceedings of the ACM FSE*, 2024.
- [18] X. Yao *et al.*, "Chain-of-event: Interpretable event-level root cause analysis for microservices," in *Proceedings of the ACM FSE*, 2024.
- [19] Y. Chen *et al.*, "Rcacopilot: Large language model assisted root cause analysis in cloud incidents," in *Proceedings of EuroSys*, 2024.
- [20] Y. Ren *et al.*, "Slim: Scalable lightweight interpretable model for root cause analysis in microservice failures," in *Proceedings of the ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (FSE)*, 2024.
- [21] H. Bai, M. Xu, K. Ye, R. Buyya, and C. Xu, "Drpc: Distributed reinforcement learning approach for scalable resource provisioning in container-based clusters," *IEEE Transactions on Services Computing*, 2024.
- [22] H. Qiu, S. S. Banerjee, S. Jha, Z. T. Kalbarczyk, and R. K. Iyer, "{FIRM}: An intelligent fine-grained resource management framework for {SLO-Oriented} microservices," in *14th USENIX symposium on operating systems design and implementation (OSDI 20)*, 2020, pp. 805–825.
- [23] H. Qiu, W. Mao, C. Wang, H. Franke, A. Youssef, Z. T. Kalbarczyk, T. Başar, and R. K. Iyer, "{AWARE}: Automate workload autoscaling with reinforcement learning in production cloud systems," in *2023 USENIX Annual Technical Conference (USENIX ATC 23)*, 2023, pp. 387–402.

- [24] N. Akbari, J. Grundy, A. Cheema, and A. N. Toosi, "Intentcontinuum: Using llms to support intent-based computing across the compute continuum," *arXiv preprint arXiv:2504.04429*, 2025.
- [25] M. Seyedkazemi Ardebili and A. Bartolini, "Kubeintellect: A modular llm-orchestrated agent framework for end-to-end kubernetes management," 2025.
- [26] Y. Zhang, Z. Zhou, S. Elnikety, and C. Delimitrou, "Analytically-driven resource management for cloud-native microservices," *arXiv preprint arXiv:2401.02920*, 2024.
- [27] Z. Zhang, F. Yang, X. Qin, J. Zhang, Q. Lin, G. Cheng, D. Zhang, S. Rajmohan, and Q. Zhang, "The vision of autonomic computing: Can llms make it a reality?" *arXiv preprint arXiv:2407.14402*, 2024.
- [28] Amazon Web Services, "What is amazon ec2 auto scaling?" <https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2-auto-scaling.html>, 2025, accessed: 2025-10.
- [29] K. Cheng, S. Zhang, C. Tu, X. Shi, Z. Yin, S. Lu, Y. Liang, and Q. Gu, "Proscale: Proactive autoscaling for microservice with time-varying workload at the edge," *IEEE Transactions on Parallel and Distributed Systems*, vol. 34, no. 4, pp. 1294–1312, 2023.
- [30] K. Hu, M. Xu, K. Ye, and C. Xu, "Lsram: A lightweight autoscaling and slo resource allocation framework for microservices based on gradient descent," *Software: Practice and Experience*, vol. 55, no. 4, pp. 714–730, 2025.
- [31] V. Sachidananda and A. Sivaraman, "Erlang: Application-aware autoscaling for cloud microservices," in *EuroSys*, 2024.
- [32] J. Park, B. Choi, C. Lee, and D. Han, "Graph neural network-based slo-aware proactive resource autoscaling framework for microservices," *IEEE/ACM Transactions on Networking*, vol. 32, no. 4, pp. 3331–3346, 2024.
- [33] B. Krishnamurthy and S. G. Shiva, "Large language model-guided sarsa algorithm for dynamic task scheduling in cloud computing," *Mathematics*, vol. 13, no. 6, p. 926, 2025.
- [34] L. Sui, K. Zhang, F. Wu, and X. Huang, "Large models for resource allocation in edge computing power networks," *IEEE Network*, 2024.
- [35] F. Yu, F. Yang, X. Qin, Z. Zhang, J. Zhang, Q. Lin, H. Zhang, Y. Dang, S. Rajmohan, D. Zhang *et al.*, "Enabling autonomic microservice management through self-learning agents," *arXiv preprint arXiv:2501.19056*, 2025.
- [36] Y. Gan, Y. Zhang, D. Cheng, A. Shetty, P. Rathi, N. Katarki, A. Bruno, J. Hu, B. Ritchken, B. Jackson *et al.*, "An open-source benchmark suite for microservices and their hardware-software implications for cloud & edge systems," in *Proceedings of the twenty-fourth international conference on architectural support for programming languages and operating systems*, 2019, pp. 3–18.
- [37] DeepSeek-AI, "Deepseek-r1," 2025, deepSeek Technical Report. [Online]. Available: <https://www.deepseek.com>
- [38] S. Luo, H. Xu, K. Ye, G. Xu, L. Zhang, G. Yang, and C. Xu, "The power of prediction: microservice auto scaling via workload learning," in *Proceedings of the 13th Symposium on Cloud Computing*, 2022, pp. 355–369.
- [39] G. Dong, H. Yuan, K. Lu, C. Li, M. Xue, D. Liu, W. Wang, Z. Yuan, C. Zhou, and J. Zhou, "How abilities in large language models are affected by supervised fine-tuning data composition," in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2024, pp. 177–198.
- [40] D. Guo, D. Yang, H. Zhang, J. Song, R. Zhang, R. Xu, Q. Zhu, S. Ma, P. Wang, X. Bi *et al.*, "Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning," *arXiv preprint arXiv:2501.12948*, 2025.
- [41] C. Zheng, S. Liu, M. Li, X.-H. Chen, B. Yu, C. Gao, K. Dang, Y. Liu, R. Men, A. Yang *et al.*, "Group sequence policy optimization," *arXiv preprint arXiv:2507.18071*, 2025.
- [42] C. Zheng, P. Ke, Z. Zhang, and M. Huang, "Click: Controllable text generation with sequence likelihood contrastive learning," in *Findings of the Association for Computational Linguistics: ACL 2023*, A. Rogers, J. Boyd-Graber, and N. Okazaki, Eds. Toronto, Canada: Association for Computational Linguistics, Jul. 2023, pp. 1022–1040.
- [43] R. Blundell, S. Bond, and F. Windmeijer, "Estimation in dynamic panel data models: improving on the performance of the standard gmm estimator," *Nonstationary panels, panel cointegration, and dynamic panels*, pp. 53–91, 2001.
- [44] Y. Gan, Y. Zhang, D. Cheng, K. He, P. Pancholi, T. Rosing, P. Gratz, N. E. Jerger, and D. Wentzlaff, "An open-source benchmark suite for microservices and their hardware-software implications for cloud and edge systems," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2019.
- [45] Weaveworks, "Sock-shop: A microservices demo application," <https://github.com/microservices-demo/microservices-demo>, 2016, accessed: 2025.
- [46] V. Jeyakumar, O. Madani, A. Parandeh, A. Kulshreshtha, W. Zeng, and N. Yadav, "Explaint!—a declarative root-cause analysis engine for time series data," in *Proceedings of the 2019 International Conference on Management of Data*, 2019, pp. 333–348.
- [47] S. Kandula, R. Mahajan, P. Verkaik, S. Agarwal, J. Padhye, and P. Bahl, "Detailed diagnosis in enterprise networks," in *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, 2009, pp. 243–254.
- [48] Y. Gan, M. Liang, S. Dev, D. Lo, and C. Delimitrou, "Sage: practical and scalable ml-driven performance debugging in microservices," in *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2021, pp. 135–151.
- [49] Anthropic, "Claude 3.5 haiku," 2025, anthropic Model Card. [Online]. Available: <https://www.anthropic.com>
- [50] The Kubernetes Authors, "Kubernetes: Production-grade container orchestration," <https://kubernetes.io/>, 2014, accessed: 2025-11-24.
- [51] C. Carrión, "Kubernetes scheduling: Taxonomy, ongoing issues and challenges," *ACM Computing Surveys*, vol. 55, no. 7, pp. 1–37, 2022.
- [52] M. Xu, C. Song, S. Ilager, S. S. Gill, J. Zhao, K. Ye, and C. Xu, "Coscal: Multifaceted scaling of microservices with reinforcement learning," *IEEE Transactions on Network and Service Management*, vol. 19, no. 4, pp. 3995–4009, 2022.
- [53] M. He, R. Gao, M. Conway, and B. E. Chapman, "Query pipeline optimization for cancer patient question answering systems," *arXiv preprint arXiv:2412.14751*, 2024.



Haoyu Bai received his B.Sc. and Master's degree from the University of Melbourne, Australia. He is currently pursuing a Ph.D. degree with the School of Computing and Information Systems, University of Melbourne. His research interests include efficient microservice-based applications, resource management, and cloud system optimization.



Muhammed Tawfiqul Islam is a Lecturer in the School of Computing and Information Systems at the University of Melbourne. He received his Ph.D. in Engineering from the University of Melbourne in 2021, working in the CLOUDS Laboratory. His doctoral research on scalable cloud resource management received the IEEE TCSC Outstanding Ph.D. Dissertation Award in 2021. His research interests span distributed systems, cloud and edge computing, microservice scheduling, and performance and cost-efficient resource management.



Minxian Xu is an Associate Professor with the Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences. He received his Ph.D. degree from the University of Melbourne in 2019. His research interests include resource scheduling, performance optimization, infrastructure for large language models, and cloud computing systems. He has co-authored more than 80 peer-reviewed papers in leading international conferences and journals. His Ph.D. thesis received the 2019 IEEE TCSC Outstanding Ph.D. Dissertation Award, and he was

the recipient of the 2023 IEEE TCSC Award for Excellence (Early Career Award).



Rajkumar Buyya is a Redmond Barry Distinguished Professor and Director of the Quantum Cloud Computing and Distributed Systems (qCLOUDS) Laboratory at the University of Melbourne, Australia. Recognized as one of the world's most highly cited researchers in computer science and software engineering (h-index: 181; g-index: 394; 170,000+ citations), Dr. Buyya is a Fellow of IEEE, a Foreign Fellow of Academia Europaea, and a Fellow of ACM. He co-founded five major IEEE/ACM international conferences—CCGrid, Cluster, Grid, e-Science, and UCC—and served as the Chair of their inaugural meetings. He served as the founding Editor-in-Chief of the IEEE Transactions on Cloud Computing. He is currently serving as Co-Editor-in-Chief of the Journal of Software: Practice and Experience, which was established 55+

years ago.