Murakkab: Resource-Efficient Agentic Workflow Orchestration in Cloud Platforms

Gohar Irfan Chaudhry[1], Esha Choukse[2], Haoran Qiu[2], Íñigo Goiri[2], Rodrigo Fonseca[2], Adam Belay[1], Ricardo Bianchini[3]
[1]MIT CSAIL [2]Microsoft Azure Research – Systems [3]Microsoft Azure

arXiv:2508.18298v2 [cs.MA] 3 Sep 2025

Abstract

Agentic workflows commonly coordinate multiple models and tools with complex control logic. They are quickly becoming the dominant paradigm for AI applications. However, serving them remains inefficient with today's frameworks. The key problem is that they expose workflows as opaque sequences of model and tool calls that tightly couple agent logic with model and hardware choices. Often, these workflow components are fragmented across different entities, preventing systems from reasoning about trade-offs across accuracy, latency, energy, and cost. This leads to resource waste and degraded service-level objectives (SLOs).

We present Murakkab, a resource-efficient serving system for agentic workflows. Murakkab introduces a declarative abstraction that decouples workflow specification from execution configuration. A profile-guided optimizer and adaptive runtime jointly manage the full stack: orchestrating workflow components, mapping them to models and hardware, and dynamically reconfiguring execution to satisfy user-defined SLOs. By exposing the internal structure of agentic workflows, Murakkab enables cross-layer optimization that existing frameworks and cloud schedulers cannot achieve.

Our evaluation on diverse workflows shows that Murakkab reduces GPU usage by up to 2.8×, energy consumption by 3.7×, and cost by 4.3× while maintaining SLOs.

1 Introduction

Recent advancements in Large Language Models (LLMs) and multi-modal LLMs are rapidly reshaping fields like education [18, 22], software engineering [34, 36], and healthcare [14, 30]. Recent efforts have focused on extending these models with the ability to invoke external tools (*e.g.,* web browser or code execution) at inference time [11, 64, 74, 80], enabling them to take actions beyond static text generation. This has paved the way for ***agentic workflows***, where multiple models and external tools collaborate to complete complex tasks, marking a shift from single-model inference to more sophisticated, compound AI applications.

Despite their promise, today's agentic workflows are fragmented across agent frameworks, agent providers, and cloud platforms, often within different organizational boundaries (Figure 1): developers use frameworks such as LangChain [40] or LlamaIndex [47] to stitch together models and tools, invoke provider APIs like OpenAI [62] or Databricks [23], and rely on cloud providers for compute infrastructure. Each layer pursues different goals (latency, quality, cost, or utilization), but coordination is minimal, leading to inefficiencies:

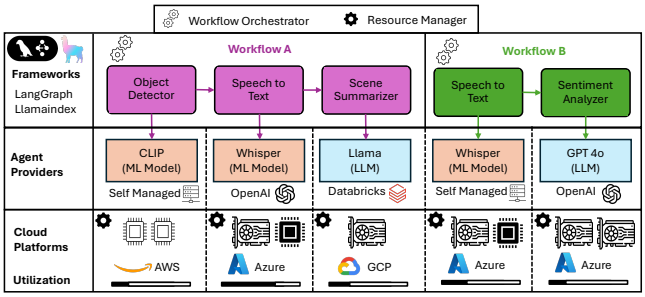

Figure 1. Today workflow developers use frameworks to call agents from *different* providers hosted on *multiple* cloud platforms. This fragmentation results in inefficiencies.

1. **Tight coupling:** Hard-coded parameters (*e.g.,* models) and hardware choices hinder automated optimization.
2. **Disjoint orchestration:** Frameworks (construct workflows) and resource managers (deploy and serve workflows) operate in silos, causing suboptimal scheduling.
3. **Difficult trade-offs:** Accuracy, latency, energy, and cost objectives require navigating a large configuration space that grows with workflow depth and model/tool choices.

These limitations drive up costs, degrade service-level objectives (SLOs), and waste resources. Static, imperative definitions further make workflows brittle: updating models or hardware requires manual refactoring and redeployment.

Our Work. We present Murakkab, a resource-efficient system for serving multi-tenant agentic workflows, built on two principles: (1) declarative workflow specification and, (2) adaptive, SLO-aware runtime system design.

Developers describe workflows with a declarative specification as *logical* tasks and dependencies, decoupled from model, tool, or hardware choices. This separation enables Murakkab to integrate workflow orchestration with resource management, dynamically reconfiguring workflow parameters (*e.g.,* models or tools) and hardware configurations as needed, based on offline profiles and online monitoring to optimize quality, latency, energy, and cost objectives.

This design enables Murakkab to efficiently explore the configuration space of workflow- and hardware-level knobs while meeting user-defined SLOs. Unlike existing deployments that treat agentic workflows as black boxes, Murakkab leverages end-to-end visibility to coordinate scheduling, colocation, and multiplexing across multi-tenant workflows.

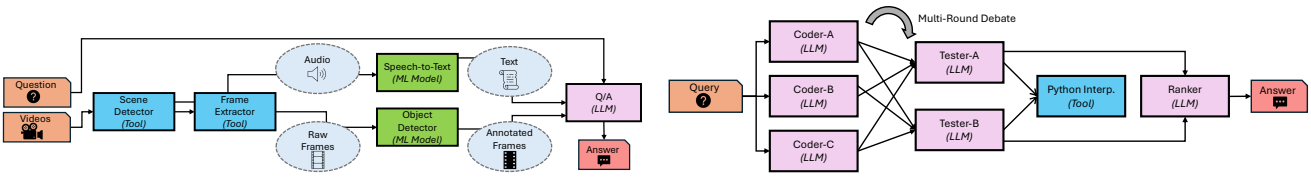

(a) Video Q/A workflow: multi-modal with tools handling different modalities before feeding into an LLM for answering the query.

(b) Code generation workflow: text-only with an LLM Debate structure to write, test and execute code.

Figure 2. Two agentic workflows with different characteristics and components.

We evaluate Murakkab on representative workflows (video question answering, code generation and mathematical problem solving) using production-scale traces. Murakkab achieves up to 2.8× lower GPU usage, 3.7× less energy, and 4.3× lower cost than state-of-the-art baselines like LangGraph while preserving quality and latency SLOs.

Summary. This paper makes the following contributions:

- Characterization of inefficiencies from workflow-, agent-, and hardware-level configurations.
- A declarative programming model and adaptive runtime unifying orchestration with resource management.
- Profile-guided optimization and workflow-aware scheduling for multi-tenant workloads.
- Evaluation showing significant efficiency gains without workflow quality or execution latency SLO violations.

2 Background and Motivation

2.1 Agentic Workflows

We define an ***agent*** as a composable unit that autonomously makes decisions, invokes tools, and participates in workflows, either independently or in collaboration with other agents [13, 21]. Each agent is typically powered by: (1) a *model* (*e.g.,* an LLM) for reasoning and language understanding, (2) a set of *instructions* specifying its goals, behavior, and constraints, and (3) *tools* for retrieving knowledge or taking actions. When multiple agents (each with specialized roles) interact and collaborate to achieve complex goals, we call the resulting process an ***agentic workflow***.

Serving agentic workflows differs fundamentally from single-model inference. Such workflows require coordinated interaction and data exchange among multiple agents to produce a final response. These agents often span diverse use cases, modalities (*e.g.,* text, audio, images etc.), execution times, and software/hardware requirements. Composing effective workflows thus demands a nuanced understanding of the trade-offs between task accuracy, request-serving latency, energy consumption, and hardware cost [38].

2.2 Example Workflows

We consider three widely used, representative workflows [60]: *Video Q/A*, *Code Generation*, and *Math Q/A* which exhibit distinct characteristics. For the rest of the paper, we focus on the first two, as shown in Figure 2, and include evaluation of the third one in Appendix A.1. A wide range of other workflows can be orchestrated using a similar approach.

Video Q/A. Video Q/A workflows are deployed for interactive query processing [67], security footage analysis [4], and many other use-cases [68, 72]. We construct a comprehensive agentic workflow to capture these diverse use-cases [7, 79]. This multi-modal workflow answers textual queries about input videos (Figure 2a). Several agents in this workflow collaborate to produce the final result:

1. *Scene Detector* to separate the input videos into distinct scenes that ensure better capture of key-frames and extract the audio clip corresponding to each scene.
2. *Frame Extractor* to sample frames from each scene.
3. *Speech-to-Text* for audio transcription.
4. *Object Detector* to annotate frames with objects of interest that are related to the query.
5. *Multi-modal LLM* (or LMM) to answer the user query given the processed frames and audio transcript.

The resulting agentic workflow for video question/answering (Q/A) can be modeled as a directed acyclic graph (DAG), where nodes represent individual agents (*e.g.,* Speech-to-Text transcription agent) and edges capture the data flow and execution dependencies among the agents (*e.g.,* passing frames from Frame Extractor to Object Detector).

Code Generation. An important use case for agentic workflows is coding and software engineering [2, 3, 9, 52] or generating code on-the-fly for accomplishing tasks in a workflow [75]. We present a representative text-only workflow that translates natural language descriptions into executable Python code (Figure 2b). It adopts the *LLM Debate* framework [49], where coder agents propose candidate solutions and tester agents generate tests and execute them using a Python interpreter. Each agent plays a unique role (*e.g.,* algorithm developer, unit tester) and may employ the same or different LLMs. The agents engage in iterative rounds of debate, aiming to reach consensus or terminating after a predefined number of rounds. The final output is selected as the highest-voted solution, determined by an LLM based on both the proposed candidates and the original query.

2.3 Rigid and Imperative Definitions

Current state-of-the-practice for agentic workflow development are frameworks such as LangGraph [41], LangChain [40], LlamaIndex [47], AutoGen [10] and CrewAI [1]. They follow

```
# ===== Workflow nodes (coupled app logic + configs) =====
scene_detection = Tool(
    fn=SceneDetector(),    # Implemented earlier by the developer.
    key=AWS_KEY, resources={"CPUs": 32}
)
frame_extractor = Tool(
    fn=FrameExtractor(),   # Implemented earlier by the developer.
    key=AWS_KEY, params={"num_frames": 15}, resources={"CPUs": 32}
)
speech_to_text = MLModel(
    name="Whisper", key=OPENAI_API_KEY, resources={"PTUs": 50}
)
object_detection = MLModel(
    name="CLIP", key=AZURE_KEY, resources={"CPUs": 128}
)
question_answer = LLM(
    name="Llama-3.2", key=DATABRICKS_API_KEY,
    params={"batch": 256}, resources={"GPUs": 8, "Type": "H100"},
    system_prompt="You are an agent that can understand videos.",
    user_prompt="Answer the given question about the video."
)
# ===== Query and data =====
ques    = "What is the name of the person wearing the red dress?"
videos  = ["road_trip.mp4"]
# ===== Workflow (ordering of nodes/data flow) =====
scenes, audio  = scene_detection(videos)
frames         = frame_extractor(scenes)
transcript     = speech_to_text(audio)
obj_frames     = object_detection(frames)
answer         = question_answer(ques, transcript, obj_frames)
```

Listing 1. Simplified video Q/A workflow definition. It shows tightly coupled application logic and execution details typical of existing frameworks.

an ***imperative*** paradigm: developers define workflow components (*i.e.,* agents), their execution logic, interactions, and numerous ***configuration parameters***.

For example, in a video Q/A workflow (Listing 1), the developer specifies parameters for `FrameExtractor` (number of frames per scene), enables options like speech-to-text or object detection, and defines the ***execution mechanism***, or model, for each component (*e.g.,* `SceneDetector` in OpenCV, Whisper [59] for transcription, Llama-3.2 [31] for reasoning).

Developers must also make ***resource allocation*** decisions, such as assigning CPUs/GPUs for self-hosted or rented VMs, or choosing service tiers (*e.g.,* provisioned throughput units or PTU [12]) for managed services.

This tightly couples configuration choices with the workflow DAG, making deployed workflows *rigid*. Any change requires manual updates and redeployment across the stack.

2.4 Fragmented Stack and Goals

Agent developers develop workflows that expose endpoints callable by end users. The workflows run on hardware offerings from cloud providers. Entities involved in developing, using, and executing agentic workflow (*i.e.,* developers, end users, and providers) prioritize different goals. *Developers* value response quality, latency, or both. *End-users* additionally care about execution cost. *Providers* prioritize resource efficiency and cost. However, the decision-making burden for workflow configuration today largely falls on developers, who must navigate an overwhelming design space and deployment knobs without holistic system visibility. Ultimately, these decisions impact all involved entities in the agentic workflow life-cycle. These configurations include:

- ***Workflow-level configurations***: *e.g.,* whether to include a Speech-to-Text Transcript agent.
- ***Agent/Node-level configurations***: *e.g.,* how many frames to extract in Frame Extractor, which LLM to use for Q/A.
- ***Hardware-level configurations***: *e.g.,* CPU *vs.* GPU and parallelism degree for each model.

2.5 Characterization and Motivation

Poor Resource Efficiency. Frameworks for agentic workflows (*e.g.,* LangGraph [41] and LlamaIndex [47]) place the burden of workflow configuration on developers. Most developers are neither systems nor ML experts, and cannot reason about accelerator choice, model parallelism, or model selection. Configurations are therefore often arbitrary and inefficient, leading to poor performance and high cost. Even experienced developers, lacking insight into user priorities, default to maximizing accuracy at the expense of efficiency.

From the cloud provider's perspective, these workflows are opaque. The cloud platform has little visibility into workflow components (*e.g.,* models) or interactions (*e.g.,* task sequences and data flow). Tightly coupled application logic and execution details prevent the platform from reconfiguring workflows to improve resource efficiency while meeting SLOs on accuracy or latency. The result is resource underutilization, increased costs, and a larger energy footprint (often passed on to users through higher prices).

> **Insight 1:** *Cloud platforms lack visibility into workflow internals (*e.g., *tasks, requirements), preventing end-to-end optimization. Meanwhile, developers lack control or insight into system-level resource behavior.*

Multi-level Trade-offs. Configurations spanning across workflow level, agent level, and hardware layer introduce complex trade-offs over accuracy, latency, energy, and cost.

Workflow- and Agent-Level knobs. Figure 3 shows a subset of configurations for the video Q/A and code generation workflows and their response accuracy. For video Q/A, Figure 3a shows that both workflow-level knobs (*e.g.,* whether to enable speech-to-text) and agent-level knobs (*e.g.,* number of frames extracted) significantly affect accuracy. Figure 3b shows the tokens generated for each configuration. Higher token generation is associated with higher latency, cost, and GPU load. For example, using Gemma-3-27B [73] with 10 extracted frames and STT enabled achieves the highest video Q/A accuracy (66.2%) but also generates the most tokens compared to other configurations of the same model. Token counts vary widely across requests, from 250 to nearly 1000, often with a heavy tail.

For the code generation workflow, accuracy is sensitive to the number of debaters and rounds in LLM Debate (Figure 3c).

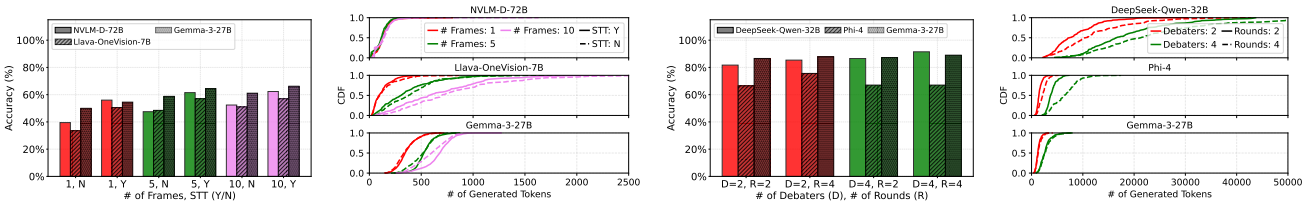

(a) Video Q/A accuracy. Varying frames, enable STT, and LMM. **(b)** Video Q/A token generation by varying the knobs. **(c)** Code Gen. accuracy. Varying debaters, rounds, and LLM. **(d)** Code Gen. tokens generated by varying the knobs.

Figure 3. Workflow accuracy under different configurations and the token generation load on the respective models.

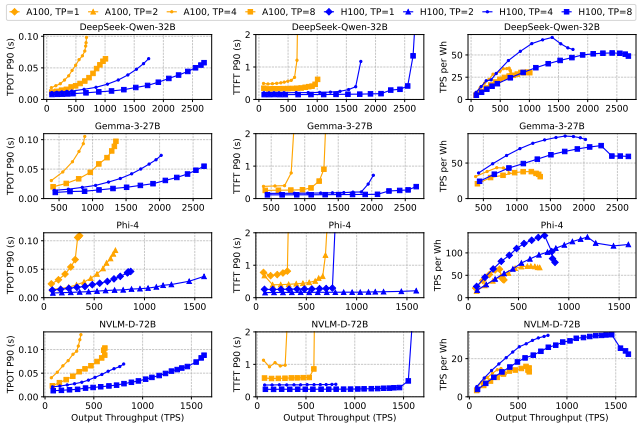

Figure 4. Model performance under different hardware and parallelism configurations.

Some configurations, such as DeepSeek-Qwen-32B [24] (a reasoning model), achieve the highest accuracy but at a much higher token generation cost (Figure 3d). At the median, it generates ≈20,000 tokens versus ≈2,500 for Gemma-3-27B under the same workflow configuration. Yet token counts alone do not capture the full end-to-end workflow characteristics, such as latency, cost, or energy.

Hardware-level knobs. Different models in a workflow have varying compute/memory requirements. Their interaction with the hardware determines whether a workflow can meet user SLOs while satisfying cloud platform constraints, like the number of GPUs allocated. Figure 4 shows models running on different hardware configurations (*e.g.,* GPU types and parallelism). We evaluate time to first token (TTFT) and time per output token (TPOT) under varying load, and measure energy efficiency as throughput per Wh. The results reveal a rich trade-off space, where configurations offer different balances of throughput, latency, and energy efficiency.

Insight 2: *Configuration knobs introduce fundamental trade-offs: workflow- and agent-level knobs affect accuracy* vs. *latency/cost, while additionally, hardware-level knobs drive cost* vs. *performance trade-offs, making it challenging to optimize workflows across multiple objectives.*

High-Dimensional Configuration Space. The configuration space for the end-to-end agentic workflow can grow exponentially with the configuration choices of the individual agents that make up the workflow. Figures 5a and 5b show that even modest agentic workflows like video Q/A or code generation, with only a small set of knobs, produce explosive configuration spaces to choose from. Each point in these Pareto frontiers represents a valid end-to-end configuration derived from varying model selection, hardware accelerator assignments, and resource scaling (with model parallelism) strategies. Navigating this configuration space manually is infeasible given that: (1) objectives vary across users (*e.g.,* low-latency *vs.* low-cost *vs.* high accuracy), (2) the optimal configuration depends on dynamic runtime context (*e.g.,* user traffic, model availability, cloud resource availability), and (3) model and hardware choices are rapidly evolving.

Insight 3: *Configuration complexity of each workflow grows combinatorially with the number of agents and their parameters at each level:* $(O(\#WorkflowKnobs \times \#AgentKnobs \times \#HardwareKnobs))$.

3 Murakkab Design

In light of this characterization, we propose Murakkab, a system that: (1) has visibility into workflow requirements, (2) holistically controls the entire workflow development and execution stack, and (3) enables automated configuration.

Instead of rigid, imperative workflows hand-tuned by developers, Murakkab supports *declarative* specifications that *decouple* high-level intent from low-level configuration. Developers focus on application logic, while the system *dynamically configures* workflows to adapt to evolving runtime conditions. *Developers* define workflow structure and task dependencies, *end-users* specify SLOs (*e.g.,* quality, latency, cost), and the *platform* dynamically optimizes execution for resource efficiency (*e.g.,* energy, cost) in a multi-tenant environment. Thus, workflows can automatically adapt to new models, varying load, or changing resources (without code rewrites or redeployments).

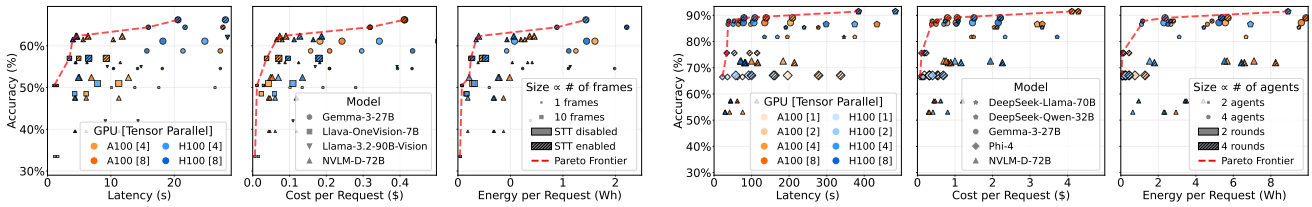

(a) Video Q/A workflow. **(b)** Code generation workflow.

Figure 5. Large space of workflow configurations along a subset of knobs and metrics.

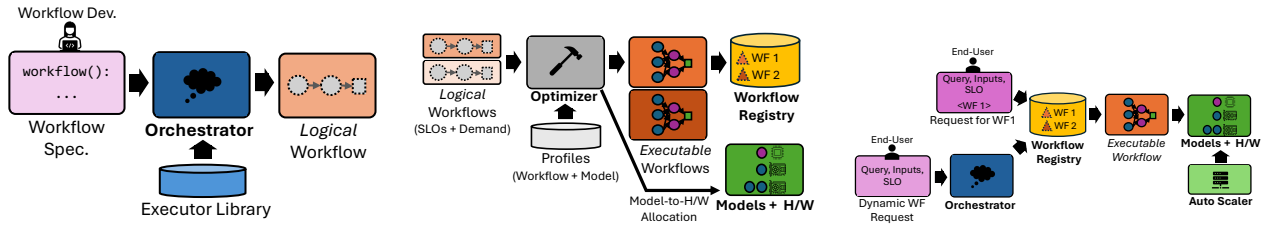

(a) Development phase. **(b)** Optimization and deployment phase. **(c)** Execution phase.

Figure 6. Murakkab manages end-to-end workflow life-cycle: from development to optimized deployment and execution.

```
# ===== Sub-tasks in the workflow =====
scene_detect   = "Given a list of videos, identify scenes in each."
frame_extract  = "Given a list of scenes, extract frames."
stt            = "Given a list of scenes, convert audio to text."
q_a            = "Answer the query given some context."
# ===== Workflow description using the sub-tasks and data flow =====
def workflow(query, videos):
    scenes      = scene_detect(videos)
    frames      = frame_extract(scenes)
    transcript  = stt(scenes)
    answer      = q_a(query, [frames, transcript])
    return answer
# ===== Execution with example request =====
query   = "What is the name of the person wearing the red dress?"
videos  = ["road_trip.mp4"]
result  = run(workflow(query, videos), slo=LOW_LATENCY)
```

Listing 2. Murakkab's declarative workflow specification of the video Q/A abstracts away configuration details, letting developers focus on application logic.

3.1 Workflow Life-Cycle

Murakkab cohesively manages the end-to-end life-cycle, unlike existing systems that fragment these among different entities. There are three main phases in an agentic workflow's life-cycle, as shown in Figure 6: (1) Workflow Development (Figure 6a), (2) Workflow Optimization (Figure 6b), and (3) Workflow Execution (Figure 6c).

3.2 Development

Murakkab adopts a declarative paradigm for workflow specification that decouples *application logic* from low-level *execution details*. Workflow developers may specify high-level tasks, and optionally, the data flow between them, without the need to provide any workflow configurations (*e.g.,* resource allocation or model selection details).

Executor Library. Murakkab is designed to interoperate with existing model and tool ecosystems. It supports LLMs and traditional machine learning (ML) models from repositories such as Hugging Face [25], as well as tools from open-source libraries and platforms, including OpenAI Agents SDK [61], Google Vertex AI Agent Garden [29], NVIDIA NeMo Agent Toolkits [55], and Microsoft Azure AI Foundry Tools [11]. In Murakkab, an *executor* is a functional unit within a workflow that can take one of three forms:

(1) ***LLM***: specialized LLM configurations (fine-tuning, few-shot learning, or even just domain-specific prompting);

(2) ***Structured compositions***: aggregations of models, *e.g.,* a self-reflection [66] or an LLM-Debate pattern [49] built from multiple LLMs;

(3) ***Tool***: utility modules for AI workflows to take actions with (*e.g.,* OpenCV frame extractor for video processing, web-search, file-search, computer-use, or any third-party tools that follow the MCP specification [63]). Traditional ML models are also included as tools (*e.g.,* CNN-based image classifiers or Word2Vec sentiment analyzers).

Attributes. Each model or tool in the library exposes three attributes: (1) a textual description, (2) an interface specification, and (3) a key-value list of configurable parameters. For example, the frame extraction tool exposes the knobs: `F` (number of frames to extract) and `cores` (number of CPU cores to run on). The LLM Debate composition exposes the knobs: `D` (number of debaters), `R` (number of rounds), and `model` (which LLM to use). The *orchestrator* uses these descriptions and interfaces to rank and assign executors (models or tools) for workflow tasks, deferring parameter configuration to a later optimization phase (Section 3.3).

A key design choice in Murakkab is mapping a broad range of unknown tasks to executors that are built from a finite, known set of models and tools in the library. If none is found, Murakkab prompts the developer to onboard a suitable one.

Workflow Specification. We allow tasks to be expressed in *natural language* with varying levels of control over how the task may be executed, in contrast to imperative paradigms discussed earlier. Recent agentic-workflow development frameworks (*e.g.,* DSPy [69]) are moving in this direction. Murakkab takes this a step further and completely decouples configuration details from the workflow specification.

Declarative Specification. Developers can manually decompose tasks into sub-tasks and define data flow between them, providing precise control over workflow steps. Listing 2 shows the developer specifies sub-tasks (*e.g.,* scene detection, speech-to-text transcription, frame extraction) and their dependencies, ensuring sequential execution to achieve the desired outcome, such as answering a question about a video. Configuration details (*e.g.,* which LLM to use, number of frames to extract, resource allocation) are omitted from the specification. However, Murakkab does not restrict developers from specifying any execution preferences (*e.g.,* particular LLM choice or hardware constraint), which are then incorporated into the optimization process as constraints.

Interface. Once deployed, workflows are exposed to end users via dedicated REST endpoints. Each request may include an SLO, such as accuracy, latency, or cost tier. Murakkab uses these SLOs to configure workflows per request and optimize resource efficiency across tasks.

Workflow Orchestrator. The *workflow orchestrator* transforms a *declarative workflow* specification into a *logical workflow*. It interprets the specification, parses tasks and sub-tasks, and maps each to an appropriate executor from Murakkab's library. At the core is an LLM with tool-calling capabilities [64], which receives a list of available executors and their interfaces, along with task descriptions, and selects the best executor for each sub-task. Recent work [43, 53, 78] has focused on automating the discovery and refinement of effective agentic workflows to maximize task performance and minimize human intervention, which Murakkab can benefit from. A feedback loop allows developers to inspect and refine the generated specification, supporting hybrid workflows with both manual and system-generated tasks.

Logical Workflow. This is an abstract execution plan that captures the functional intent of each task without binding to specific models, resources, or hardware. It is represented as a directed acyclic graph (DAG), where nodes correspond to executors and edges indicate data flow. This representation remains request-agnostic, containing no per-request details such as query text, input payloads, or SLOs. Execution specifics (*e.g.,* model selection or hardware allocation) are deferred to later stages.

The orchestrator performs type-checking on the DAG to ensure output types from source nodes match input types of destination nodes. In case of mismatches, the workflow is regenerated with error feedback to the LLM. Persistent errors prompt the developer to revise the specification.

3.3 Optimization and Deployment

Profiles. Accurate, fine-grained performance characterization is essential for optimizing multi-tenant agentic workflows with dynamic execution patterns. Inspired by Profile-Guided Optimization (PGO) [54, 76], Murakkab builds offline profiles across diverse configurations to inform runtime decisions. Each profile captures three key metrics per workflow configuration: response quality, end-to-end latency, and resource usage. To support this, Murakkab maintains two profiling layers: workflow profiles and model profiles.

Workflow Profiles. Murakkab profiles representative agentic workflows (*e.g.,* video Q/A) under various combinations of workflow-level knobs (*e.g.,* enabling/disabling STT) and executor-level knobs (*e.g.,* varying the number of frames extracted by the extractor tool). *Quality* is assessed using open-source benchmarks and datasets (*e.g.,* VideoMME [27], HumanEval [17], and Math [35]) with ground-truth results. When public datasets fall short, Murakkab can onboard curated evaluation datasets from developers or collect request-response pairs and user feedback (*e.g.,* rating, approvals), similar to deployed systems like ChatGPT [58]. These inputs are periodically incorporated into profile updates to refine configuration choices for active workflows.

Configuration choices also affect per-request resource demands. Workflow profiles quantify executor-level load, including prompt and completion tokens for LLM-based executors, serving as a proxy for resource usage. These profiles capture metrics similar to those in our earlier characterization (*e.g.,* Figures 3a and 3b).

Model Profiles. To assess workflow resource demands, executor-level load must be mapped to model and hardware metrics. Murakkab maintains profiles for each model, capturing performance across software configurations (*e.g.,* tensor parallelism, prefill/decode separation) and hardware setups (*e.g.,* GPU type, clock frequency). Each profile reports: (1) latency (TTFT and TPOT for LLMs), (2) energy consumption across hardware, and (3) cost per configuration. Profiles span load levels to expose trade-offs and guide the optimizer in allocating load and instances.

Profiling is lightweight; performed once per configuration and reused across workflows. New models and accelerators are profiled upon integration. Decoupling workflow and model profiles enables workflows to benefit immediately from model updates, with selective re-profiling as needed.

This separation enables independent evolution, profile reuse, rapid model onboarding, and efficient hardware adaptation. The optimizer uses profiles as structured priors to evaluate cost–accuracy–latency trade-offs and drive reconfiguration under dynamic conditions.

Workflow Optimizer. The optimizer transforms the *logical workflow* from the orchestrator into an *executable workflow* by selecting: (1) workflow parameters, (2) models/tools per executor, and (3) hardware and parallelism strategies.

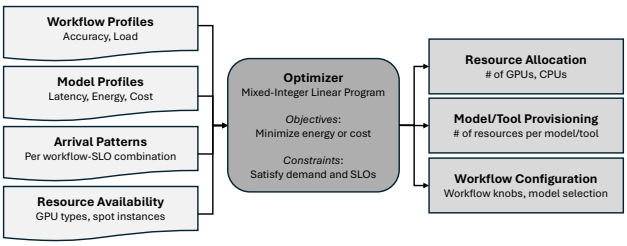

Figure 7. Summary of Murakkab's optimization process.

It combines workflow and model profiles to estimate latency, cost, and energy for candidate configurations. Workflow profiles capture task accuracy and executor-level load; model profiles quantify latency, energy, and cost under varying load and system setups. For requests with SLOs on accuracy, latency, or cost, the optimizer selects configurations that meet constraints while maximizing efficiency.

For example, to estimate latency for a video Q/A request, Murakkab: (1) filters out configurations that do not meet the quality SLO, (2) uses token distributions and model profiles to estimate latency and cost, (3) selects a configuration satisfying all SLOs (or returns an infeasibility error).

Beyond individual workflows, Murakkab optimizes system-wide resource usage. It leverages global visibility to align configurations across workflows, enabling colocation and executor multiplexing for diverse SLOs.

During each *optimization epoch*, the optimizer considers workflow load and resource availability (*e.g.,* spot instances) to allocate model instances. It balances peak and average load to maximize sharing while respecting SLOs. Figure 7 summarizes the optimization process.

Formulation. The optimizer uses a Mixed Integer Linear Programming (MILP) formulation to allocate resources for agentic workflows under varying SLOs. Inputs include: (1) *workflow profiles* defining accuracy, token needs, and SLO targets (latency, cost, accuracy), (2) *model profiles* specifying throughput, latency, energy, GPU requirements, (3) *arrival patterns* showing request distributions per workflow-SLO pair, and (4) *resource constraints* including cost limits.

The MILP matches workflows to feasible model profiles, then determines instance counts and workload distribution. A key insight is separating peak provisioning from average utilization: GPUs are allocated for peak demand, while average loads guide cost-efficient usage.

The optimizer supports multiple objectives: minimizing energy, minimizing cost, or maximizing accuracy given a cost budget. The output is a deployment plan: GPU instance counts per model profile and workflow-to-instance assignments, ensuring SLO compliance and resource efficiency. We provide the detailed formulation in Appendix A.5.

Workflows that can benefit from *external* calls to proprietary models can also be supported by the optimizer by assigning appropriate cost, latency, and accuracy attributes to those model profiles.

Once an executable workflow is generated for all valid SLO tiers of an onboarded workflow, it is added to the Murakkab workflow registry and is ready to serve requests.

3.4 Execution

At runtime, Murakkab receives incoming requests from end-users with a payload that contains the identifier of the agentic workflow being invoked, any input query/data, and the SLOs. Murakkab looks up the registry to obtain the corresponding executable workflow and submits it for execution.

Dynamic Workflow Requests. End-users can either invoke a *particular* agentic workflow or send a request with a natural language query (Figure 6c), any input data to operate on, and SLOs, *without* specifying an agentic workflow to use. The workflow orchestrator parses the query into one or more sub-tasks, mapping each to an appropriate executor or an existing workflow. Thus, Murakkab dynamically composes workflows from existing building blocks available to it.

SLOs. Each request has the option to specify a quality, latency and cost SLOs. We assign four SLO tiers for quality and end-to-end latency: *best*, *good*, *fair*, and *basic*. The SLO tiers correspond to the best, 95th, 80th, and 50th percentile values of accuracy and latency available among the set of all workflow, model, and hardware configurations.

Runtime Optimization. While request dispatch is deterministic given the selected workflow, achieving efficient execution under dynamic, multi-tenant conditions requires continuous adaptation. The optimizer runs in the background after every optimization epoch, in our case every 60 minutes, to adapt to the most up-to-date load and resource availability in the system. The state of the previous epochs is used to project the load for each workflow in the next epoch. Using this information, the optimizer reconfigures the workflows and updates their executable workflows in the registry.

Auto-Scaler. Predicting end-to-end resource usage in agentic workflows is difficult, as input-dependent control flow and intermediate outputs propagating along data flows determine actual demand (*e.g.,* Figures 3b and 3d). For example, a video Q/A workflow with 10 frames and STT on Llava-OneVision-7B [42] produces 600 and 1200 tokens in the 50th and 99th percentile, respectively, highlighting high variance.

To handle such variability, Murakkab includes an *auto-scaler* that monitors per-model instance load over short windows (seconds to minutes) and rapidly scales out when needed.

We set thresholds for autoscaling based on the performance-throughput characteristic in executor profiles. This mechanism prioritizes avoiding SLO violations over short-term allocation optimality. Murakkab also maintains spare resources to absorb demand spikes and, when it detects significant

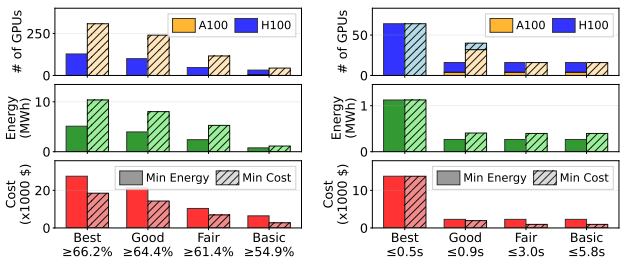

(a) Accuracy SLOs. **(b)** Latency SLOs.

Figure 8. Video Q/A workflow configured for different SLOs and optimization objectives.

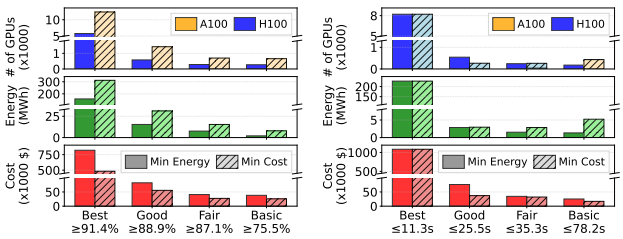

(a) Accuracy SLOs. **(b)** Latency SLOs.

Figure 9. Code generation workflow configured for different SLOs and optimization objectives.

deviations in workload or resource usage, triggers early re-optimization to adapt quickly.

4 Evaluation

4.1 Experimental Setup

Hardware. We run our experiments on A100 and H100 VMs from Microsoft Azure. Each A100 VM has 8×NVIDIA A100 (80GB) GPUs and an AMD EPYC 7V12 64-Core processor, while each H100 VM has 8×NVIDIA H100 (80GB) GPUs with an Intel Xeon (Sapphire Rapids) processor. We use vLLM (v0.9) [39] as the LLM inference engine, speaches-ai (v0.7) [5] as the speech-to-text model serving engine, and OmDet [81] for object detection model serving.

Production Traces. Since no publicly available traces exist for production agentic workflow serving, we approximate workload arrivals using LLM serving traces collected over a 24-hour period in May 2024 from Azure's LLM inference service for *chat* and *coding* applications [70] (Appendix A.4).

Agentic Workflows. Our evaluation focuses on the video Q/A and code generation workflows. We map the *chat* requests from the trace to the *video Q/A* workflow and *coding* requests to the *code generation* workflow.

Policies. We consider three scheduling policies:

(1) *Static* is a hand-crafted baseline that balances cost and accuracy for both workflows, using Gemma-3-27B and A100 GPUs. However, it lacks visibility into the workflows and has a disconnect between orchestration and resource management, resulting in lack of adaptability to shifting demand. This represents existing systems like LangGraph [41].

(2) *Murakkab Optimized* (Mrkb Opt) optimizes each workflow–SLO combination for a specific objective (either minimizing energy or cost).

(3) *Murakkab Optimized + Multiplexing* (Mrkb Opt+Mult) *jointly* optimizes requests across all workflow–SLO combinations to maximize colocation and sharing of model instances for even better efficiency.

When evaluating Murakkab over the production traces, we set the optimization epoch to 60 minutes, and also conduct a sensitivity analysis for this choice in Section 4.6.

4.2 Single-Workflow Optimization

To understand how Murakkab optimizes a single workflow under varying request SLOs and optimization objectives, we run Murakkab with two optimization objectives: *minimize energy consumption* and *minimize execution cost*. We use the workload traces [70] for each workflow and assume that all requests have the same SLO for each experiment. We report the type and number of GPUs allocated, energy consumption in MWh, and the cost in $.

Accuracy SLOs. Figure 8a shows the requests across SLO accuracies. When minimizing energy, Murakkab lowers energy consumption from 5.1 MWh at 66.2% accuracy (*best*) to 3.9 MWh at 64.4% accuracy (*good*) – an energy reduction of 23.5% with negligible accuracy impact. Tolerating a drop to 61.4% (*fair*) yields substantial savings of 2.6 MWh from the peak consumption. On the other hand, when optimizing for cost, Murakkab reduces from $18.5k to $14.3k while marginally dropping accuracy between *best* to *good*. Allowing accuracy to drop further to 61.4% reduces expenses to $6.9k, a ≈ 4× decrease compared to the most expensive configuration. Murakkab prefers using H100 GPUs when minimizing energy, trading off cost for energy savings.

Figure 9a shows results for the code generation workflow. Energy use ranges from 312 MWh to 2 MWh, and cost from $820k to $25k, across varying accuracy levels. Relaxing the accuracy SLO from *best* to *good* yields a sharp drop in energy (≈ 10.5×) and cost (≈ 8.7×), mainly due to Murakkab switching from DeepSeek-Qwen-32B to Gemma-3-27B, which has lower token and operation costs.

Latency SLOs. Murakkab serves latency-sensitive requests with varying SLO tiers, while minimizing energy or cost. A similar pattern in resource and cost savings emerges for both workflows as shown in Figures 8b and 9b. We guarantee a basic accuracy tier even for latency SLO requests (*e.g.*, 50% for video Q/A). For example, Murakkab can reduce energy consumption from 1.1 MWh to 266 kWh at a slight increase in end-to-end latency, from 0.5 s to 0.9 s for the video Q/A workflow. On the other hand, it can reduce energy consumption from 227 MWh to 2.8 MWh (as a result of changing to a different model with different tensor parallelism).

We detail the most common workflow configurations chosen by Murakkab's optimizer for each of the experiments

Policy	# of GPUs	Energy (MWh)	Cost (×1000 $)
Static	2560	80.4	201.5
Murakkab Opt	1151	27.1	56.2
Murakkab Opt+Mult	908	21.6	46.5

Table 1. Comparison of resource usage, energy, and cost. Murakkab substantially reduces GPUs, energy, and cost compared to static allocation, with further gains from multiplexing across all video Q/A and code generation workflow requests.

over the entire arrival trace in Appendices A.2 and A.3 (Tables 4 and 5). We observe similar trends for the Math Q/A workflow and report the results in Appendix A.1 (Figures 16a and 16b).

4.3 Multi-Workflow Optimization

We measure the improvements in efficiency as a result of joint optimization and multiplexing across different workflow–SLO combinations. For this experiment, we run video Q/A and code generation requests together and assign 70% requests to be high-accuracy and 30% requests to low-latency, both with *good tier* (Section 3.4).

Figure 10a shows the *static baseline* has a fixed allocation with a fixed cost. Its energy consumption varies depending on resource usage, but is significantly higher than the other configurations due to higher idle-power consumption of its GPUs. Table 1 summarizes the results for the 24 hour trace. This configuration uses ≈ 2,560 A100 GPUs consistently resulting in a flat hourly cost. A non-hand-crafted baseline could be *significantly* more inefficient than our current baseline from requiring bigger models and more resources or result in more SLO violations if under-provisioned.

Mrkb Opt. This policy optimizes each workflow–SLO combination to minimize cost, although *without* considering the demand across other combinations running in the system. It can adapt to changing load (leveraging the difference between peak and average utilization) to change model and resource allocation over time. It requires 1,151 GPUs, 27.1 MWh energy, and $56,246 cost.

Mrkb Opt+Mult. This policy additionally leverages its *holistic* view of all workflows and their demand to maximize multiplexing opportunities. It requires 908 GPUs (21.1% reduction), 21.6 MWh energy (20.2% reduction), and $46,494 cost (17.3% reduction) respectively.

Accuracy and Latency. The static baseline cannot distinguish between per-request SLO variations to configure workflows appropriately. Therefore, high-accuracy and low-latency SLO requests have the same accuracy and latency, resulting in latency SLO violations, as shown in Figure 10b. Murakkab leverages per-request SLO information to provision enough resources for each request, as can be seen in the difference in accuracy and latency of the two categories of requests, allowing for higher resource-efficiency.

Available A100s	Available H100s	Allocated A100s	Allocated H100s	Energy (MWh)	Cost (×1000 $)
2000	0	1292	0	24.7	55.7
2000	100	780	100	17.8	52.4
2000	200	536	200	14.5	60.3
2000	300	288	300	12.4	64.8
2000	400	50	400	11.1	70.6
2000	500	0	495	11.0	75.4

Table 2. Murakkab serving multiple workflows varying resource constraints minimizing energy. As more H100 GPUs are available, Murakkab trades off cost for energy (same SLO).

4.4 Adapting to Dynamic Resource Availability

To evaluate how Murakkab adapts to changes in resources (*e.g.,* spot-instances, cloud provider prioritizing resources for other services), we run the same arrival pattern and workflow-SLO distribution as in Section 4.3. However, we constrain the type and amount of resources available to Murakkab. We assume a fixed number of A100 GPUs available in the cluster (2000) and vary the number of H100 GPUs available from 0 to 500, increasing by 100 each time. We let Murakkab optimize to minimize energy consumption when serving the incoming workflow requests.

Table 2 shows that Murakkab leverages the H100 GPUs as they become available, while reducing the number of A100 GPUs. Starting with an allocation of 1292 A100 GPUs and no H100 GPUs, which consumes 24.7 MWh energy, Murakkab adapts to the changing H100 availability and leverage most of them (up to 495 H100s) to bring down the energy consumption to 11 MWh. This shows that Murakkab adjusts its resource allocation under resource-constrained settings to satisfy the SLOs while striving for higher resource efficiency.

Execution Analysis. Figure 11 shows Murakkab adapting to changes in load. We consider the *400×H100 GPUs* configuration (without loss of generality) and observe that:

1. Murakkab *scales model instances up/down* after every optimization epoch to adapt to the shift in load dynamics.
2. Murakkab *changes GPU allocation* with system load. It prefers H100 GPUs due to their energy efficiency and uses A100 GPUs for the rest, depending on load fluctuations (*e.g.,* at T=9 hours).
3. Murakkab *reconfigures workflows to maximize colocation.* For example, Gemma serves code generation requests with high-accuracy SLOs and multiplexes load from video Q/A requests with high-accuracy SLOs when there is surplus capacity (*e.g.,* from T=5–9 hours and T=15–19).

4.5 Workflow/DAG-Aware Scheduling

Consider the request: *verify the student's coding solution from a video*, with a *30-second latency SLO.* The orchestrator constructs a DAG with a fan-out for the two sub-tasks that can execute in parallel: (1) video Q/A (Figure 2a) to extract the

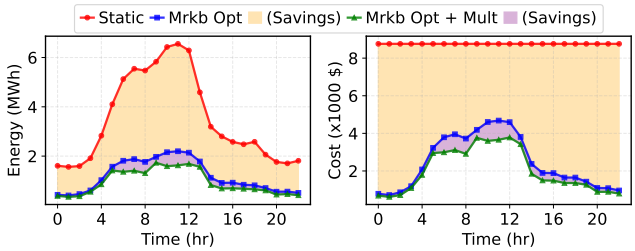

(a) Energy and cost over time. Static policy is agnostic to load shifts while Murakkab reconfigures workflows for higher efficiency.

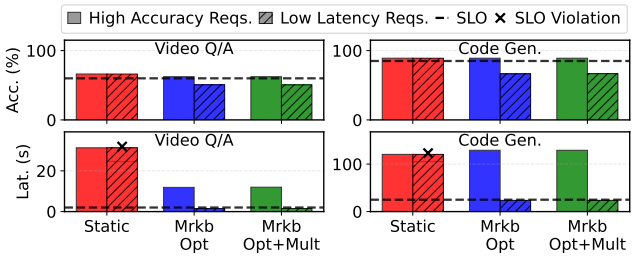

(b) Accuracy and latency for all requests. Static policy does not distinguish between request SLO categories.

Figure 10. Comparing three policies: (1) a hand-crafted static configuration, (2) Murakkab optimizing individual workflows (*Opt*), and (3) Murakkab jointly-optimizing across workflows + multiplexing resources (*Opt+Mult*).

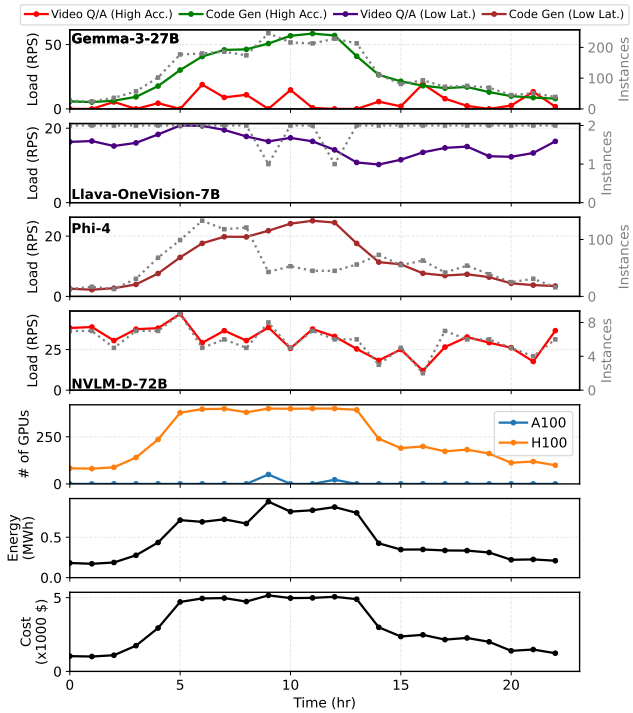

Figure 11. Murakkab adjusts resource allocation and model instances across all workflow-SLO combinations with changing load (under a constraint of 400×H100 GPUs.)

student's solution, and (2) code generation (Figure 2b) to produce a reference solution for comparison.

Scheduling. Leveraging workflow and model profiles, Murakkab identifies viable scheduling options based on cluster resource availability (Figure 12). All configurations use Gemma-3-27B on 4×A100 GPUs. We highlight setups that leverage both GPUs and CPUs for improved efficiency.

OmDet and Whisper on GPUs. Figure 12a shows OmDet and Whisper running on dedicated A100 GPUs, with a total of 6×A100 GPUs. The two sub-tasks run in near-perfect parallel, with full overlap in execution. However, despite meeting the latency goal, dedicated GPUs are underutilized, and CPUs are mostly idle.

OmDet and Whisper on CPUs. Figure 12b shows Whisper and OmDet running on CPUs, reducing GPU usage to 4×A100 and utilizing idle CPU resources. Whisper runs efficiently without saturating CPUs, making it a good candidate for offloading. In contrast, OmDet fully loads all cores and significantly increases latency, resulting in an SLO violation.

OmDet on GPU, Whisper on CPU. Figure 12c shows OmDet running on 1×A100 GPU and Whisper on CPUs, with a total of 5×A100 GPUs provisioned. Both sub-tasks complete nearly simultaneously, and Whisper's added CPU latency has minimal impact on end-to-end time. This configuration meets the latency SLO while reducing GPU usage compared to the all-GPU setup.

Murakkab in Action. Leveraging workflow stage visibility and pre-generated execution profiles, Murakkab selects the configuration with OmDet on GPU and Whisper on CPU (Figure 12c) as the most resource-efficient option (balancing overall utilization and latency). In contrast, traditional systems treat agentic workflows as opaque, lacking coordination between orchestration and scheduling. This leads to inefficient execution, higher costs, and SLO violations.

4.6 Optimization Frequency Sensitivity Analysis

The choice of optimization epoch presents a fundamental trade-off between resource efficiency, cost, and system responsiveness. We evaluate intervals from 20 minutes to 6 hours, measuring cost, utilization, and worst-case dropped requests. Provisioning new instances (*i.e.,* VM allocation, software setup, and model transfer to GPUs) is assumed to take 20 minutes [28, 33, 56]. We use an exponentially weighted moving average (EWMA) [20], with $\alpha = 0.5$, to predict workload demand at every epoch.

Three-Zone Cost Structure. Figure 13a shows three distinct zones that guide optimal selection.

Zone 1 (10–60 minutes): Buffer-dominated. Frequent reoptimization induces high transition overhead. Excessive GPU provisioning during transitions leads to lower utilization,

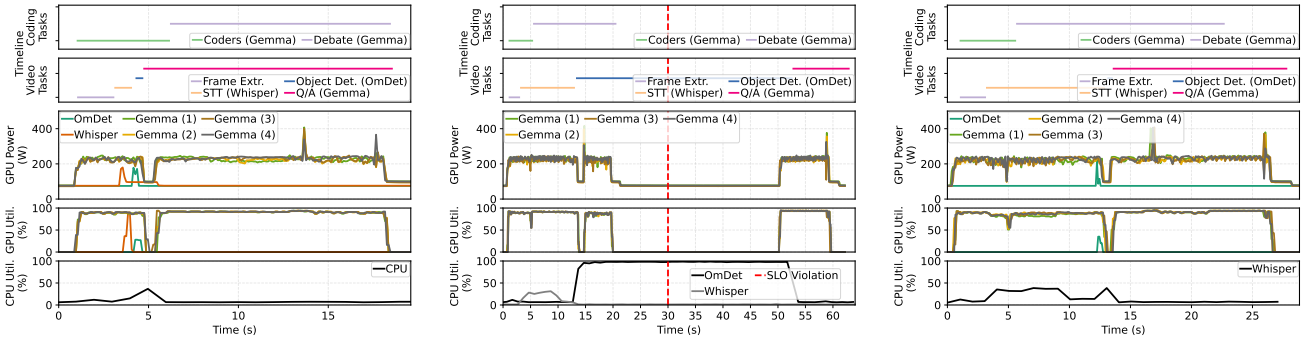

(a) OmDet on 1×A100, Whisper on 1×A100, Gemma on 4×A100s. End-to-end execution completes *within* latency SLO *(using 6 GPUs)*.

(b) OmDet on CPUs, Whisper on CPUs, Gemma on 4×A100s. End-to-end execution ***violates*** latency SLO *(using 4 GPUs)*.

(c) OmDet on 1×A100, Whisper on CPUs, Gemma on 4×A100s. End-to-end execution completes *within* latency SLO *(using 5 GPUs)*.

Figure 12. Executing a user request involving parallel video Q/A and code generation. Murakkab selects the configuration in Figure 12c to minimize energy use while meeting the 30-second latency SLO and maintaining response quality.

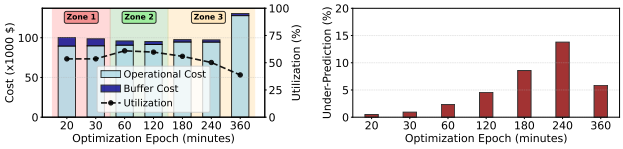

(a) Cost *vs.* resource utilization. **(b)** Demand under-prediction.

Figure 13. Murakkab sensitivity to optimization epoch.

despite responsive demand adaptation. Frequent model and tool changes can also reduce KV cache efficiency for LLMs.

Zone 2 (60–180 minutes): Balanced. Transition costs and prediction uncertainty offset, yielding peak cost efficiency. Utilization peaks at an epoch of around 60 minutes, reflecting the best balance between adaptation frequency and stability.

Zone 3 (180+ minutes): Uncertainty-dominated. At long intervals, buffer cost from transition overhead is minimal, but coarse-grained provisioning increasingly diverges from fine-grained, dynamic demand. This reduces prediction accuracy and drives over-provisioning, lowering efficiency.

Worst-case System Responsiveness. We measure responsiveness by demand under-prediction (requests that would be dropped without auto-scaling). This reflects the *worst-case*, as auto-scalers can mitigate mismatches post-provisioning. Figure 13b shows under-prediction is minimal at short intervals but rises steadily, peaking at ~15% around 240 minutes. Beyond that, longer intervals lead to over-provisioning, lowering under-prediction but hurting utilization.

Operational Implications. This analysis reveals a key trade-off: short intervals improve responsiveness but increase transition costs, while long intervals reduce overhead but hurt prediction accuracy and utilization. Providers can tune this balance based on priorities (*e.g.*, cost, utilization, responsiveness). In our setup, a 60-minute interval offers a middle ground (high utilization, low cost, and manageable shortage risk). While we use EWMA for forecasting, the trade-off holds with more advanced predictors as well.

5 Related Work

Agentic Workflow Development. Frameworks like LangGraph [41], LangChain [40], and AutoGen [10] build agentic workflows *imperatively*, composing model and tool calls. DSPy [69] and Palimpzest [45, 46] take a *declarative* approach, focusing on prompt and query optimization. However, these frameworks still blur configuration and logic, burden developers with resource management, and struggle to scale efficiently across large configuration spaces.

Automated Workflow Generation. There has been extensive work from the ML community on automating workflow generation to improve response quality [43, 46, 48, 53, 77, 78]. This line of work is complementary to Murakkab, which can integrate these workflow generation techniques into the Murakkab workflow orchestrator (Section 3.2).

Systems Optimization. Prior systems have focused on accelerating workflow execution with improved data flow management [65, 71] and scheduling of LLM inference calls in the context of agentic applications [44, 50]. Others target response quality by exploring model selection [16, 26] and scaling properties [15, 57] of test-time compute. Systems like SpotServe [51] and Loki [8] address resource and load dynamism, but remain confined to single-model serving. Recent works [19, 37, 38] have examined the energy costs of test-time compute and their impact on response quality. In contrast, Murakkab introduces a fully automated system to navigate these complex tradeoffs for agentic workflows.

6 Conclusion

Murakkab brings a declarative programming model and adaptive runtime to the serving of multi-tenant agentic workflows. By decoupling workflow logic from execution configurations and tightly integrating orchestration with cluster management, Murakkab enables dynamic reconfiguration and resource-aware scheduling while honoring diverse SLOs. Our evaluation shows that Murakkab substantially improves

efficiency (reducing GPU usage, energy consumption, and cost) without compromising workflow quality or latency. We see opportunities in extending adaptive co-optimization across larger clusters, supporting more heterogeneous accelerators, and exploring workload specialization for emerging classes of agentic applications.

References

[1] 2025. CrewAI — The Leading Multi-Agent Platform. https://www.crewai.com/. Accessed: 19 August 2025.

[2] 2025. Features | Cursor – The AI Code Editor. https://cursor.com/features. Accessed: 19 August 2025.

[3] 2025. GitHub Copilot · Your AI pair programmer. https://github.com/features/copilot. Accessed: 19 August 2025.

[4] 2025. Plix AI. https://plix.ai/. Accessed: 19 August 2025.

[5] 2025. Speaches – OpenAI-API compatible server for speech-to-text, speech-to-speech, and translation. https://speaches.ai/. Accessed: 20 August 2025.

[6] Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Xin Wang, Rachel Ward, Yue Wu, Dingli Yu, Cyril Zhang, and Yi Zhang. 2024. Phi-4 Technical Report. arXiv:2412.08905 [cs.CL] https://arxiv.org/abs/2412.08905

[7] Shubham Agrawal, Adeola Adesoba, Dhruv Nandakumar, Katherine Huang, and Vignesh Srinivasakumar. 2024. Build an Agentic Video Workflow with Video Search and Summarization. NVIDIA Developer Blog. Available at: https://developer.nvidia.com/blog/build-an-agentic-video-workflow-with-video-search-and-summarization/.

[8] Sohaib Ahmad, Hui Guan, and Ramesh K. Sitaraman. 2024. Loki: A System for Serving ML Inference Pipelines with Hardware and Accuracy Scaling. In *Proceedings of the 33rd International Symposium on High-Performance Parallel and Distributed Computing* (Pisa, Italy) *(HPDC '24)*. Association for Computing Machinery, New York, NY, USA, 267–280. https://doi.org/10.1145/3625549.3658688

[9] Anthropic. 2025. Claude Code: Deep coding at terminal velocity. https://www.anthropic.com/claude-code. Accessed: 19 August 2025.

[10] AutoGen. 2024. AutoGen. https://microsoft.github.io/autogen/stable//index.html.

[11] Microsoft Azure. 2025. Azure AI Foundry Tool Library. https://learn.microsoft.com/en-us/azure/ai-foundry/agents/how-to/tools/overview.

[12] Microsoft Azure. 2025. Understanding costs associated with provisioned throughput units (PTU). https://learn.microsoft.com/en-us/azure/ai-foundry/openai/how-to/provisioned-throughput-onboarding.

[13] Microsoft Azure. 2025. What is an AI Agent? https://learn.microsoft.com/en-us/azure/ai-foundry/agents/overview.

[14] Marco Cascella, Jonathan Montomoli, Valentina Bellini, and Elena Bignami. 2023. Evaluating the feasibility of ChatGPT in healthcare: an analysis of multiple clinical and research scenarios. *Journal of Medical Systems* 47, 1 (2023), 33.

[15] Lingjiao Chen, Jared Quincy Davis, Boris Hanin, Peter Bailis, Ion Stoica, Matei Zaharia, and James Zou. 2024. Are More LLM Calls All You Need? Towards Scaling Laws of Compound Inference Systems. arXiv:2403.02419 [cs.LG] https://arxiv.org/abs/2403.02419

[16] Lingjiao Chen, Jared Quincy Davis, Boris Hanin, Peter Bailis, Matei Zaharia, James Zou, and Ion Stoica. 2025. Optimizing Model Selection for Compound AI Systems. arXiv:2502.14815 [cs.AI] https://arxiv.org/abs/2502.14815

[17] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating Large Language Models Trained on Code. arXiv:2107.03374 [cs.LG] https://arxiv.org/abs/2107.03374

[18] Zhendong Chu, Shen Wang, Jian Xie, Tinghui Zhu, Yibo Yan, Jinheng Ye, Aoxiao Zhong, Xuming Hu, Jing Liang, Philip S Yu, et al. 2025. LLM agents for education: Advances and applications. *arXiv preprint arXiv:2503.11733* (2025).

[19] Jae-Won Chung, Jiachen Liu, Jeff J. Ma, Ruofan Wu, Oh Jun Kweon, Yuxuan Xia, Zhiyu Wu, and Mosharaf Chowdhury. 2025. The ML.ENERGY Benchmark: Toward Automated Inference Energy Measurement and Optimization. arXiv:2505.06371 [cs.LG] https://arxiv.org/abs/2505.06371

[20] Petar Cisar, Saša Bošnjak, and Sanja Maravic Cisar. 2010. EWMA algorithm in network practice. *International Journal of Computers Communications & Control* 5, 2 (2010), 160–170.

[21] Google Cloud. 2025. What are AI Agents? https://cloud.google.com/discover/what-are-ai-agents.

[22] Ling Dai, Yuan-Hao Jiang, Yuanyuan Chen, Zinuo Guo, Tian-Yi Liu, and Xiaobao Shao. 2024. Agent4EDU: Advancing AI for Education with Agentic Workflows. In *Proceedings of the 2024 3rd International Conference on Artificial Intelligence and Education*. 180–185.

[23] Databricks. 2025. Databricks Large Language Model Serving. https://docs.databricks.com/en/large-language-models/index.html.

[24] DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue

Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. arXiv:2501.12948 [cs.CL] https://arxiv.org/abs/2501.12948

[25] Hugging Face. 2025. Hugging Face Models. https://huggingface.co/.

[26] Tao Feng, Yanzhen Shen, and Jiaxuan You. 2025. GraphRouter: A Graph-based Router for LLM Selections. arXiv:2410.03834 [cs.AI] https://arxiv.org/abs/2410.03834

[27] Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, Peixian Chen, Yanwei Li, Shaohui Lin, Sirui Zhao, Ke Li, Tong Xu, Xiawu Zheng, Enhong Chen, Caifeng Shan, Ran He, and Xing Sun. 2025. Video-MME: The First-Ever Comprehensive Evaluation Benchmark of Multi-modal LLMs in Video Analysis. arXiv:2405.21075 [cs.CV] https://arxiv.org/abs/2405.21075

[28] Yao Fu, Leyang Xue, Yeqi Huang, Andrei-Octavian Brabete, Dmitrii Ustiugov, Yuvraj Patel, and Luo Mai. 2024. ServerlessLLM: Low-Latency serverless inference for large language models. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*. 135–153.

[29] Google. 2025. Google Vertex AI Agent Garden. https://console.cloud.google.com/vertex-ai/agents/agent-garden.

[30] Sagar Goyal, Eti Rastogi, Sree Prasanna Rajagopal, Dong Yuan, Fen Zhao, Jai Chintagunta, Gautam Naik, and Jeff Ward. 2024. HealAI: A healthcare LLM for effective medical documentation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*. 1167–1168.

[31] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, et al. 2024. The Llama 3 Herd of Models. arXiv:2407.21783 [cs.CL] https://arxiv.org/abs/2407.21783

[32] Gurobi Optimization, LLC. 2024. Gurobi Optimizer Reference Manual. https://www.gurobi.com

[33] Jianwei Hao, Ting Jiang, Wei Wang, and In Kee Kim. 2021. An empirical analysis of VM startup times in public IaaS clouds. In *2021 IEEE 14th International Conference on Cloud Computing (CLOUD)*. IEEE, 398–403.

[34] Junda He, Christoph Treude, and David Lo. 2025. LLM-Based Multi-Agent Systems for Software Engineering: Literature Review, Vision, and the Road Ahead. *ACM Transactions on Software Engineering and Methodology* 34, 5 (2025), 1–30.

[35] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring Mathematical Problem Solving With the MATH Dataset. *NeurIPS* (2021).

[36] Haolin Jin, Linghan Huang, Haipeng Cai, Jun Yan, Bo Li, and Huaming Chen. 2024. From LLMs to LLM-based agents for software engineering: A survey of current, challenges and future. *arXiv preprint arXiv:2408.02479* (2024).

[37] Junho Jin, Gu-Yeon Wei, and David Brooks. 2025. The Energy Cost of Reasoning: Analyzing Energy Usage in LLMs with Test-time Compute. arXiv:2505.14733 [cs.LG] https://arxiv.org/abs/2505.14733

[38] Jiin Kim, Byeongjun Shin, Jinha Chung, and Minsoo Rhu. 2025. The Cost of Dynamic Reasoning: Demystifying AI Agents and Test-Time Scaling from an AI Infrastructure Perspective. arXiv:2506.04301 [cs.LG] https://arxiv.org/abs/2506.04301

[39] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient Memory Management for Large Language Model Serving with PagedAttention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.

[40] LangChain. 2024. LangChain. https://github.com/langchain-ai/langchain.

[41] LangGraph. 2024. LangGraph. https://www.langchain.com/langgraph.

[42] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024. LLaVA-OneVision: Easy Visual Task Transfer. arXiv:2408.03326 [cs.CV] https://arxiv.org/abs/2408.03326

[43] Zelong Li, Shuyuan Xu, Kai Mei, Wenyue Hua, Balaji Rama, Om Raheja, Hao Wang, He Zhu, and Yongfeng Zhang. 2024. Autoflow: Automated workflow generation for large language model agents. *arXiv preprint arXiv:2407.12821* (2024).

[44] Chaofan Lin, Zhenhua Han, Chengruidong Zhang, Yuqing Yang, Fan Yang, Chen Chen, and Lili Qiu. 2024. Parrot: efficient serving of LLM-based applications with semantic variable. In *Proceedings of the 18th USENIX Conference on Operating Systems Design and Implementation* (Santa Clara, CA, USA) *(OSDI'24)*. USENIX Association, USA, Article 50, 17 pages.

[45] Chunwei Liu, Matthew Russo, Michael Cafarella, Lei Cao, Peter Baile Chen, Zui Chen, Michael Franklin, Tim Kraska, Samuel Madden, Rana Shahout, and Gerardo Vitagliano. [n.d.]. Palimpzest: Optimizing AI-Powered Analytics with Declarative Query Processing. In *Proceedings of the Conference on Innovative Database Research (CIDR)* (2025).

[46] Chunwei Liu, Matthew Russo, Michael Cafarella, Lei Cao, Peter Baille Chen, Zui Chen, Michael Franklin, Tim Kraska, Samuel Madden, and Gerardo Vitagliano. 2024. A Declarative System for Optimizing AI Workloads. arXiv:2405.14696 [cs.CL]

[47] Jerry Liu. 2022. LlamaIndex. https://github.com/jerryjliu/llama_index.

[48] Jiale Liu, Yifan Zeng, Shaokun Zhang, Chi Zhang, Malte Højmark-Bertelsen, Marie Normann Gadeberg, Huazheng Wang, and Qingyun Wu. 2025. Divide, Optimize, Merge: Fine-Grained LLM Agent Optimization at Scale. arXiv:2505.03973 [cs.CL] https://arxiv.org/abs/2505.03973

[49] Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. 2024. A dynamic LLM-powered agent network for task-oriented agent collaboration. In *First Conference on Language Modeling*.

[50] Michael Luo, Xiaoxiang Shi, Colin Cai, Tianjun Zhang, Justin Wong, Yichuan Wang, Chi Wang, Yanping Huang, Zhifeng Chen, Joseph E. Gonzalez, and Ion Stoica. 2025. Autellix: An Efficient Serving Engine for LLM Agents as General Programs. arXiv:2502.13965 [cs.LG] https://arxiv.org/abs/2502.13965

[51] Xupeng Miao, Chunan Shi, Jiangfei Duan, Xiaoli Xi, Dahua Lin, Bin Cui, and Zhihao Jia. 2023. SpotServe: Serving Generative Large Language Models on Preemptible Instances. arXiv:2311.15566 [cs.DC] https://arxiv.org/abs/2311.15566

[52] Daye Nam, Andrew Macvean, Vincent Hellendoorn, Bogdan Vasilescu, and Brad Myers. 2024. Using an LLM to Help With Code Understanding. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering* (Lisbon, Portugal) *(ICSE '24)*. Association for Computing Machinery, New York, NY, USA, Article 97, 13 pages. https://doi.org/10.1145/3597503.3639187

[53] Boye Niu, Yiliao Song, Kai Lian, Yifan Shen, Yu Yao, Kun Zhang, and Tongliang Liu. 2025. Flow: Modularized agentic workflow automation. *arXiv preprint arXiv:2501.07834* (2025).

[54] Diego Novillo. 2014. SamplePGO - The Power of Profile Guided Optimizations without the Usability Burden. In *2014 LLVM Compiler Infrastructure in HPC*. 22–28. https://doi.org/10.1109/LLVM-HPC.2014.8

[55] NVIDIA. 2025. NeMo Agent Toolkit. https://developer.nvidia.com/nemo-agent-toolkit.

[56] NVIDIA. 2025. NVIDIA DOCA Overview. https://docs.nvidia.com/doca/archive/2-9-0/nvidia+doca+overview/index.html.

[57] Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E. Gonzalez, M Waleed Kadous, and Ion Stoica. 2025. RouteLLM: Learning to Route LLMs with Preference Data. arXiv:2406.18665 [cs.LG] https://arxiv.org/abs/2406.18665
[58] OpenAI. 2023. ChatGPT (Mar 14 version) [Large language model]. https://chat.openai.com/chat.
[59] OpenAI. 2024. Whisper Large V3 Model. https://huggingface.co/openai/whisper-large-v3.
[60] OpenAI. 2025. Coding, Math, and Multimodal capabilities are the widely used benchmarks for agent and model evaluation. https://openai.com/index/introducing-o3-and-o4-mini/
[61] OpenAI. 2025. OpenAI Agents SDK. https://openai.github.io/openai-agents-python/tools/.
[62] OpenAI. 2025. OpenAI Large Language Models and API. https://platform.openai.com/docs/.
[63] Model Context Protocol. 2025. Model Context Protocol (MCP). https://modelcontextprotocol.io/docs/getting-started/intro.
[64] Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. 2023. ToolLLM: Facilitating large language models to master 16000+ real-world APIs. *arXiv preprint arXiv:2307.16789* (2023).
[65] Deepti Raghavan, Keshav Santhanam, Muhammad Shahir Rahman, Nayani Modugula, Luis Gaspar Schroeder, Maximilien Cura, Houjun Liu, Pratiksha Thaker, Philip Levis, and Matei Zaharia. 2025. Alto: Orchestrating Distributed Compound AI Systems with Nested Ancestry. arXiv:2403.04311 [cs.AI] https://arxiv.org/abs/2403.04311
[66] Matthew Renze and Erhan Guven. 2024. Self-reflection in llm agents: Effects on problem-solving performance. *arXiv preprint arXiv:2405.06682* (2024).
[67] Francisco Romero, Johann Hauswald, Aditi Partap, Daniel Kang, Matei Zaharia, and Christos Kozyrakis. 2022. Optimizing Video Analytics with Declarative Model Relationships. *Proc. VLDB Endow.* 16, 3 (Nov. 2022), 447–460. https://doi.org/10.14778/3570690.3570695
[68] Francisco Romero, Mark Zhao, Neeraja J. Yadwadkar, and Christos Kozyrakis. 2021. Llama: A Heterogeneous & Serverless Framework for Auto-Tuning Video Analytics Pipelines. arXiv:2102.01887 [cs.DC] https://arxiv.org/abs/2102.01887
[69] Stanford NLP Group. 2023. DSPy: The Framework for Programming—Not Prompting—Language Models. https://github.com/stanfordnlp/dspy.
[70] Jovan Stojkovic, Chaojie Zhang, Íñigo Goiri, Josep Torrellas, and Esha Choukse. 2025. DynamoLLM: Designing LLM inference clusters for performance and energy efficiency. In *2025 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 1348–1362.
[71] Xin Tan, Yimin Jiang, Yitao Yang, and Hong Xu. 2025. Teola: Towards End-to-End Optimization of LLM-based Applications. arXiv:2407.00326 [cs.DC] https://arxiv.org/abs/2407.00326
[72] Yunlong Tang, Jing Bi, Siting Xu, Luchuan Song, Susan Liang, Teng Wang, Daoan Zhang, Jie An, Jingyang Lin, Rongyi Zhu, Ali Vosoughi, Chao Huang, Zeliang Zhang, Pinxin Liu, Mingqian Feng, Feng Zheng, Jianguo Zhang, Ping Luo, Jiebo Luo, and Chenliang Xu. 2025. Video Understanding with Large Language Models: A Survey. arXiv:2312.17432 [cs.CV] https://arxiv.org/abs/2312.17432
[73] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, Gaël Liu, Francesco Visin, Kathleen Kenealy, Lucas Beyer, Xiaohai Zhai, Anton Tsitsulin, Robert Busa-Fekete, Alex Feng, Noveen Sachdeva, Benjamin Coleman, Yi Gao, Basil Mustafa, Iain Barr, Emilio Parisotto, David Tian, Matan Eyal, Colin Cherry, Jan-Thorsten Peter, Danila Sinopalnikov, Surya Bhupatiraju, Rishabh Agarwal, Mehran Kazemi, Dan Malkin, Ravin Kumar, David Vilar, Idan Brusilovsky, Jiaming Luo, Andreas Steiner, Abe Friesen, Abhanshu Sharma, Abheesht Sharma, Adi Mayrav Gilady, Adrian Goedeckemeyer, Alaa Saade, Alex Feng, Alexander Kolesnikov, Alexei Bendebury, Alvin Abdagic, Amit Vadi, András György, André Susano Pinto, Anil Das, Ankur Bapna, Antoine Miech, Antoine Yang, Antonia Paterson, Ashish Shenoy, Ayan Chakrabarti, Bilal Piot, Bo Wu, Bobak Shahriari, Bryce Petrini, Charlie Chen, Charline Le Lan, Christopher A. Choquette-Choo, CJ Carey, Cormac Brick, Daniel Deutsch, Danielle Eisenbud, Dee Cattle, Derek Cheng, Dimitris Paparas, Divyashree Shivakumar Sreepathihalli, Doug Reid, Dustin Tran, Dustin Zelle, Eric Noland, Erwin Huizenga, Eugene Kharitonov, Frederick Liu, Gagik Amirkhanyan, Glenn Cameron, Hadi Hashemi, Hanna Klimczak-Plucińska, Harman Singh, Harsh Mehta, Harshal Tushar Lehri, Hussein Hazimeh, Ian Ballantyne, Idan Szpektor, Ivan Nardini, Jean Pouget-Abadie, Jetha Chan, Joe Stanton, John Wieting, Jonathan Lai, Jordi Orbay, Joseph Fernandez, Josh Newlan, Ju yeong Ji, Jyotinder Singh, Kat Black, Kathy Yu, Kevin Hui, Kiran Vodrahalli, Klaus Greff, Linhai Qiu, Marcella Valentine, Marina Coelho, Marvin Ritter, Matt Hoffman, Matthew Watson, Mayank Chaturvedi, Michael Moynihan, Min Ma, Nabila Babar, Natasha Noy, Nathan Byrd, Nick Roy, Nikola Momchev, Nilay Chauhan, Noveen Sachdeva, Oskar Bunyan, Pankil Botarda, Paul Caron, Paul Kishan Rubenstein, Phil Culliton, Philipp Schmid, Pier Giuseppe Sessa, Pingmei Xu, Piotr Stanczyk, Pouya Tafti, Rakesh Shivanna, Renjie Wu, Renke Pan, Reza Rokni, Rob Willoughby, Rohith Vallu, Ryan Mullins, Sammy Jerome, Sara Smoot, Sertan Girgin, Shariq Iqbal, Shashir Reddy, Shruti Sheth, Siim Põder, Sijal Bhatnagar, Sindhu Raghuram Panyam, Sivan Eiger, Susan Zhang, Tianqi Liu, Trevor Yacovone, Tyler Liechty, Uday Kalra, Utku Evci, Vedant Misra, Vincent Roseberry, Vlad Feinberg, Vlad Kolesnikov, Woohyun Han, Woosuk Kwon, Xi Chen, Yinlam Chow, Yuvein Zhu, Zichuan Wei, Zoltan Egyed, Victor Cotruta, Minh Giang, Phoebe Kirk, Anand Rao, Kat Black, Nabila Babar, Jessica Lo, Erica Moreira, Luiz Gustavo Martins, Omar Sanseviero, Lucas Gonzalez, Zach Gleicher, Tris Warkentin, Vahab Mirrokni, Evan Senter, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, Yossi Matias, D. Sculley, Slav Petrov, Noah Fiedel, Noam Shazeer, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Jean-Baptiste Alayrac, Rohan Anil, Dmitry, Lepikhin, Sebastian Borgeaud, Olivier Bachem, Armand Joulin, Alek Andreev, Cassidy Hardin, Robert Dadashi, and Léonard Hussenot. 2025. Gemma 3 Technical Report. arXiv:2503.19786 [cs.CL] https://arxiv.org/abs/2503.19786
[74] Jize Wang, Ma Zerun, Yining Li, Songyang Zhang, Cailian Chen, Kai Chen, and Xinyi Le. 2024. GTA: a benchmark for general tool agents. *Advances in Neural Information Processing Systems* 37 (2024), 75749–75790.
[75] Xingyao Wang, Yangyi Chen, Lifan Yuan, Yizhe Zhang, Yunzhu Li, Hao Peng, and Heng Ji. 2024. Executable Code Actions Elicit Better LLM Agents. In *Forty-first International Conference on Machine Learning*. https://openreview.net/forum?id=jJ9BoXAfFa
[76] Patrick Wintermeyer, Maria Apostolaki, Alexander Dietmüller, and Laurent Vanbever. 2020. P2GO: P4 profile-guided optimizations. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks*. 146–152.
[77] Shirley Wu, Parth Sarthi, Shiyu Zhao, Aaron Lee, Herumb Shandilya, Adrian Mladenic Grobelnik, Nurendra Choudhary, Eddie Huang, Karthik Subbian, Linjun Zhang, et al. 2025. Optimas: Optimizing Compound AI Systems with Globally Aligned Local Rewards. *arXiv preprint arXiv:2507.03041* (2025).
[78] Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xionghui Chen, Jiaqi Chen, Mingchen Zhuge, Xin Cheng, Sirui Hong, Jinlin Wang, et al. 2024. Aflow: Automating agentic workflow generation. *arXiv preprint arXiv:2410.10762* (2024).
[79] Lu Zhang, Tiancheng Zhao, Heting Ying, Yibo Ma, and Kyusong Lee. 2024. OmAgent: A Multi-modal Agent Framework

for Complex Video Understanding with Task Divide-and-Conquer. arXiv:2406.16620 [cs.CL] https://arxiv.org/abs/2406.16620

[80] Yaolun Zhang, Yinxu Pan, Yudong Wang, and Jie Cai. 2024. PyBench: Evaluating LLM agent on various real-world coding tasks. *arXiv preprint arXiv:2407.16732* (2024).

[81] Tiancheng Zhao, Peng Liu, and Kyusong Lee. 2024. OmDet: Large-scale vision-language multi-dataset pre-training with multimodal detection network. *IET Computer Vision* 18, 5 (Jan. 2024), 626–639. https://doi.org/10.1049/cvi2.12268

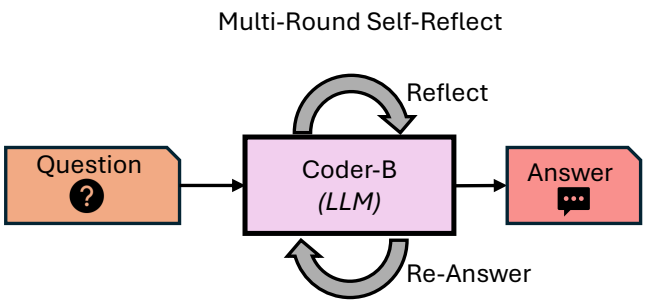

Figure 14. Math Q/A workflow (self-reflect [66] structure).

Policy	# of GPUs	Energy (MWh)	Cost (×1000 $)
Static	4448	169.87	367.2
Mrkb Opt	1875	62.88	123.0
Mrkb Opt+Mult	1660	52.66	104.6

Table 3. Comparison of resource usage, energy, and cost. Murakkab substantially reduces GPUs, energy, and cost compared to static allocation, with further gains from multiplexing across all math Q/A and code generation workflow requests.

SLO	Objective	Tier	Model	Frames	GPU	TP	TPOT (s)	TPS
Acc.	Cost	Best	Gemma-3-27B	10	A100	4	0.0624	699
		Fair	NVLM-D-72B	5	A100	4	0.0966	325
		Basic	Llava-OneVision-7B	5	A100	4	0.0224	2244
		Good	Gemma-3-27B	5	A100	4	0.0624	700
	Energy	Best	Gemma-3-27B	10	H100	4	0.0484	1688
		Fair	NVLM-D-72B	5	H100	4	0.0650	766
		Basic	Llava-OneVision-7B	5	H100	4	0.0079	3271
		Good	Gemma-3-27B	5	H100	4	0.0472	1668
Lat.	Cost	Best	Llava-OneVision-7B	1	H100	4	0.0044	479
		Fair	Llava-OneVision-7B	1	A100	4	0.0224	2244
		Basic	Llava-OneVision-7B	1	A100	4	0.0224	2244
		Good	Llava-OneVision-7B	1	A100	4	0.0085	926
	Energy	Best	Llava-OneVision-7B	1	H100	4	0.0044	479
		Fair	Llava-OneVision-7B	1	H100	4	0.0070	2836
		Basic	Llava-OneVision-7B	1	H100	4	0.0070	2836
		Good	Llava-OneVision-7B	1	H100	4	0.0070	2836

Table 4. Video Q/A workflow configurations chosen by Murakkab corresponding to the experiment in Section 4.2

A Appendix

A.1 Math Q/A

Math Q/A, as shown in Figure 14, is a workflow to answer mathematical problems. It uses a self-reflect [66] structure where a *mathematician* agent generates a response to the question, self-reflects on the response, and iteratively continues this until it is confident in the answer or a limit on the maximum number of rounds is reached.

A.1.1 Characterization. We present the results (Figure 15) for different models and different number of self-refine rounds (R) (Figure 15a) and the token generation load in terms of prompt (Figure 15c) and completion tokens (Figure 15b) for different configurations.

A.1.2 Single-Workflow Optimization. Similar to Section 4.2, we optimize the math Q/A workflow using Murakkab for different accuracy and latency SLO tiers and optimization objectives. The results are presented in Figure 16.

A.1.3 Multi-Workflow Optimization. Similar to Section 4.3, we evaluate end-to-end execution of running 24 hours of Azure traces [70]. This time, the two categories of requests from the traces map to two agentic workflows: *chat* requests maps to *math Q/A* workflow requests and *coding* requests map to the *code generation* workflow (Figure 2b). The results are shown in Figure 17. The hand-crafted *static* baseline policy (similar to LangGraph [41]) has a fixed allocation of DeepSeek-Qwen-32B for the math Q/A workflow and Gemma-3-27B for code generation workflow – balancing quality and resource usage. As explained in Section 4, the static baseline has no visibility into the workflow, load or resource availability and suffers from under-utilization. Murakkab on the other hand is adaptive to load for each workflow–SLO combination and appropriately configures workflows at each optimization epoch, resulting in massive resource savings. When allowed to multiplex across resources, it yields further savings by colocating the requests from all workflows onto shared model instances whenever appropriate. Table 3 summarizes the results for the 24 hour duration. We observe a reduction of ≈ 2.7× in GPUs, ≈ 3.2× in energy, and ≈ 3.5× in cost from the static baseline to Murakkab's best optimization policy.

A.2 Video

We present the most commonly chosen configuration by Murakkab for each workflow–SLO combination of the video Q/A workflow under different optimization objectives in Table 4 (corresponding to the experiment presented in Section 4.2). It shows various knobs (*e.g.,* model, number of frames, GPU type etc.) and the offered load per model instance in tokens-per-second (TPS) with the observed latency in time-per-output-token (TPOT) For example, we can observe the changing model and number of frames processed as the accuracy SLO is relaxed. For the latency SLO requests, Murakkab keeps the same workflow- and model-level knobs but changes GPU type and increases the allowed load per model instance to increase batching as the SLO is relaxed.

A.3 Code Generation

We present the most commonly chosen configuration by Murakkab for each workflow–SLO combination of the code generation workflow under different optimization objectives in Table 5 (corresponding to the experiment presented in Section 4.2). It shows various knobs (*e.g.,* model, number of debaters, number of debate rounds etc.) and the offered load per model instance in tokens-per-second (TPS) with the observed latency in time-per-output-token (TPOT) For

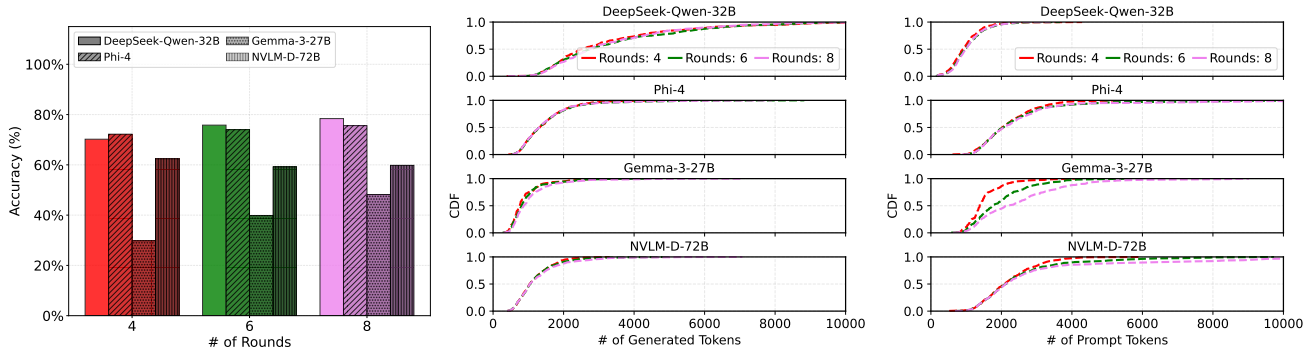

(a) Math Q/A accuracy. **(b)** Math Q/A model load (generated tokens). **(c)** Math Q/A model load (prompt tokens).

Figure 15. Workflow accuracy under different configurations and the token processing load on the respective models.

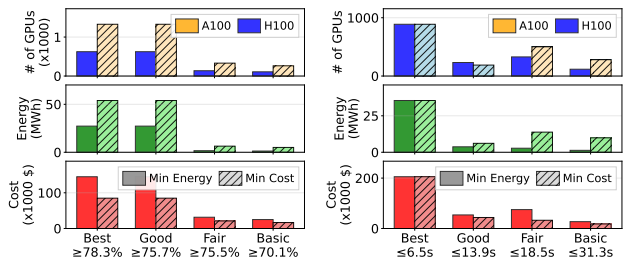

(a) Accuracy SLOs. **(b)** Latency SLOs.

Figure 16. Math Q/A workflow configured for different SLO and optimization objectives.

SLO	Objective	Tier	Model	Debaters	Rounds	GPU	TP	TPOT (s)	TPS
Acc.	Cost	Best	DeepSeek-Qwen-32B	4	4	A100	4	0.0767	653
		Good	Gemma-3-27B	4	4	A100	4	0.0624	700
		Fair	Gemma-3-27B	2	4	A100	4	0.0624	700
		Basic	Phi-4	2	4	A100	2	0.0609	623
	Energy	Best	DeepSeek-Qwen-32B	4	4	H100	4	0.0387	1390
		Good	Gemma-3-27B	4	4	H100	4	0.0496	1709
		Fair	Gemma-3-27B	2	4	H100	4	0.0487	1693
		Basic	Phi-4	2	4	H100	1	0.0373	757
Lat.	Cost	Best	NVLM-D-72B	2	2	H100	8	0.0129	84
		Good	Phi-4	2	2	H100	2	0.0169	1036
		Fair	Phi-4	2	2	H100	2	0.0218	1185
		Basic	Phi-4	2	2	A100	2	0.0609	623
	Energy	Best	NVLM-D-72B	2	2	H100	8	0.0129	84
		Good	Phi-4	2	2	H100	1	0.0165	248
		Fair	Phi-4	2	2	H100	1	0.0253	552
		Basic	Phi-4	2	2	H100	1	0.0372	755

Table 5. Code generation workflow configurations chosen by Murakkab corresponding to the experiment in Section 4.2

example, we can observe the changing model and number of debaters as the accuracy SLO is relaxed. For the latency SLO requests, Murakkab mostly keeps the same workflow- and model-level knobs but changes GPU type and increases the allowed load per model instance to increase batching as the SLO is relaxed. Notably, it also changes the tensor parallelism for the same model (Phi-4 [6]) between optimizing for cost and energy.

A.4 Azure Traces

We use a subset of LLM serving traces released by Azure [70] from 08:00 05/15/2024 to 08:00 05/16/2024 shown in Figure 18.

A.5 Optimization Formulation

We formulate our optimization problem (Section 3.3) as a temporal resource allocation task for serving multiple workflows with heterogeneous SLOs across diverse model profiles.

Sets and Indices.

- $\mathcal{W}$: workflows
- $\mathcal{S}$: SLO types
- $\mathcal{M}$: model profiles
- $\mathcal{C}_w$: workflow configurations for w
- $\mathcal{G}$: resource types

Parameters.

- $\lambda^{\text{peak}}_{w,s}$: Peak request rate for workflow w with SLO s
- $\lambda^{\text{avg}}_{w,s}$: Average request rate for workflow w with SLO s
- α: Unified buffer factor (default 1.15)
- $\tau_{w,s}$: SLO threshold for workflow w and SLO type s
- a_c: Accuracy of workflow configuration $c \in \mathcal{C}_w$
- t_c: Tokens per request for workflow configuration c
- θ_m: Token throughput (tokens/sec) for model profile m
- t^{TTFT}_m: Time to first token for model profile m
- t^{TPOT}_m: Time per output token for model profile m
- g_m: Parallelism for model m
- e_m: Energy consumption (kWh) for model profile m
- c_g: Cost per instance per second for resource type $g \in \mathcal{G}$
- B_g: Maximum available resource instances of type g

Decision Variables.

- $n_m \in \mathbb{Z}^+$: Number of instances of model profile m
- $x^{\text{peak}}_{w,s,c,m} \in \mathbb{R}^+$: Peak load allocation from (w, s, c) to model m
- $x^{\text{avg}}_{w,s,c,m} \in \mathbb{R}^+$: Average load allocation from (w, s, c) to model m

Constraints.

Demand Satisfaction (Peak):

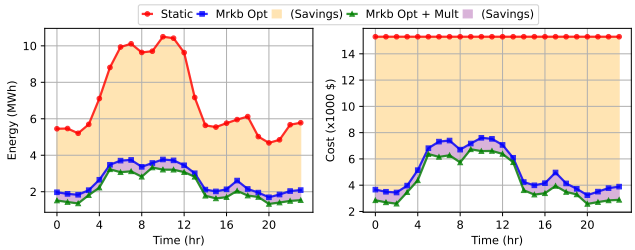

(a) Energy and cost over time. Static policy is agnostic to load shifts while Murakkab reconfigures workflows for higher efficiency.

(b) Accuracy and latency for all requests. Static policy does not distinguish between request SLO categories.

Figure 17. Comparing a hand-crafted static configuration (only A100s) to: (1) Murakkab optimizing individual workflows, and (2) Murakkab jointly-optimizing across workflows + multiplexing resources (Math Q/A + Code Gen).

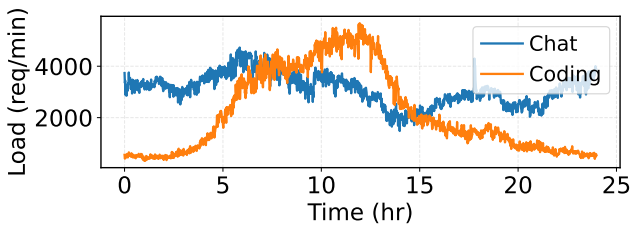

Figure 18. Azure LLM serving traces.

Ensure peak demand is met with buffer:

$$\lambda^{\text{peak}}_{w,s} \leq \sum_{c \in C_w, m \in \mathcal{M}} x^{\text{peak}}_{w,s,c,m} \leq \alpha \cdot \lambda^{\text{peak}}_{w,s}, \quad \forall w \in \mathcal{W}, s \in \mathcal{S} \quad (1)$$

Demand Satisfaction (Average):
Similar bounds for average demand:

$$\lambda^{\text{avg}}_{w,s} \leq \sum_{c \in C_w, m \in \mathcal{M}} x^{\text{avg}}_{w,s,c,m} \leq \alpha \cdot \lambda^{\text{avg}}_{w,s}, \quad \forall w \in \mathcal{W}, s \in \mathcal{S} \quad (2)$$

Capacity Constraint with Multiplexing:
Account for statistical multiplexing:

$$\mu_m \cdot \sum_{w,s,c} x^{\text{peak}}_{w,s,c,m} \cdot t_c \leq n_m \cdot \theta_m, \quad \forall m \in \mathcal{M} \quad (3)$$

where μ_m is the model-specific multiplexing factor.

SLO Filtering:
For accuracy SLO (s = max_accuracy):

$$x^{\text{peak}}_{w,s,c,m} = 0 \quad \text{if } a_c < \tau_{w,s} \quad (4)$$

For latency SLO (s = min_latency):

$$x^{\text{peak}}_{w,s,c,m} = 0 \quad \text{if } \ell^{\text{TTFT}}_m + t_c \cdot \ell^{\text{TPOT}}_m > \tau_{w,s} \quad (5)$$

Cost Budget Constraint:
If cost SLO is specified:

$$\sum_{w,s,c,m} x^{\text{avg}}_{w,s,c,m} \cdot \frac{t_c}{\theta_m} \cdot g_m \cdot c_{g(m)} \leq \text{Cost}_{\text{budget}} \quad (6)$$

where $\text{Cost}_{\text{budget}} = \sum_{w \in \mathcal{W}} \tau_{w,\text{cost}} \cdot \sum_s \lambda^{\text{avg}}_{w,s}$.

Resource Budget Constraint:

$$\sum_{m:\text{GPU}(m)=g} n_m \cdot g_m \leq B_g, \quad \forall g \in \mathcal{G} \quad (7)$$

SLO Filtering Constraints:

$$x^{\text{peak}}_{w,s,c,m} = 0 \quad \text{if } a_c < \tau_{w,s} \quad \text{(accuracy)} \quad (8)$$

$$x^{\text{peak}}_{w,s,c,m} = 0 \quad \text{if } \ell^{\text{TTFT}}_m + t_c \ell^{\text{TPOT}}_m > \tau_{w,s} \quad \text{(latency)} \quad (9)$$

Cost Budget Constraint:

$$\sum_{w,s,c,m} x^{\text{avg}}_{w,s,c,m} \frac{t_c}{\theta_m} g_m c_{g(m)} \leq \sum_w \tau_{w,\text{cost}} \sum_s \lambda^{\text{avg}}_{w,s} \quad (10)$$

Objective Functions.

Minimize Energy:

$$\min \sum_m n_m e_m g_m \quad (11)$$

Minimize Cost:

$$\min \sum_m n_m g_m c_{g(m)} \quad (12)$$

Maximize Accuracy Under a Cost Budget:

$$\max \frac{\sum_{w,s,c,m} x^{\text{avg}}_{w,s,c,m} a_c}{\sum_{w,s} \lambda^{\text{avg}}_{w,s}} - \epsilon \cdot \text{Cost}_{\text{total}} \quad (13)$$

where $\epsilon = 0.001$.

Solution Method.

The formulated Mixed Integer Linear Program (MILP) is solved using Gurobi [32] with a time limit of 300 seconds. The solution yields an allocation of model instance counts n^*_m and load distributions across model instances for all workflows.