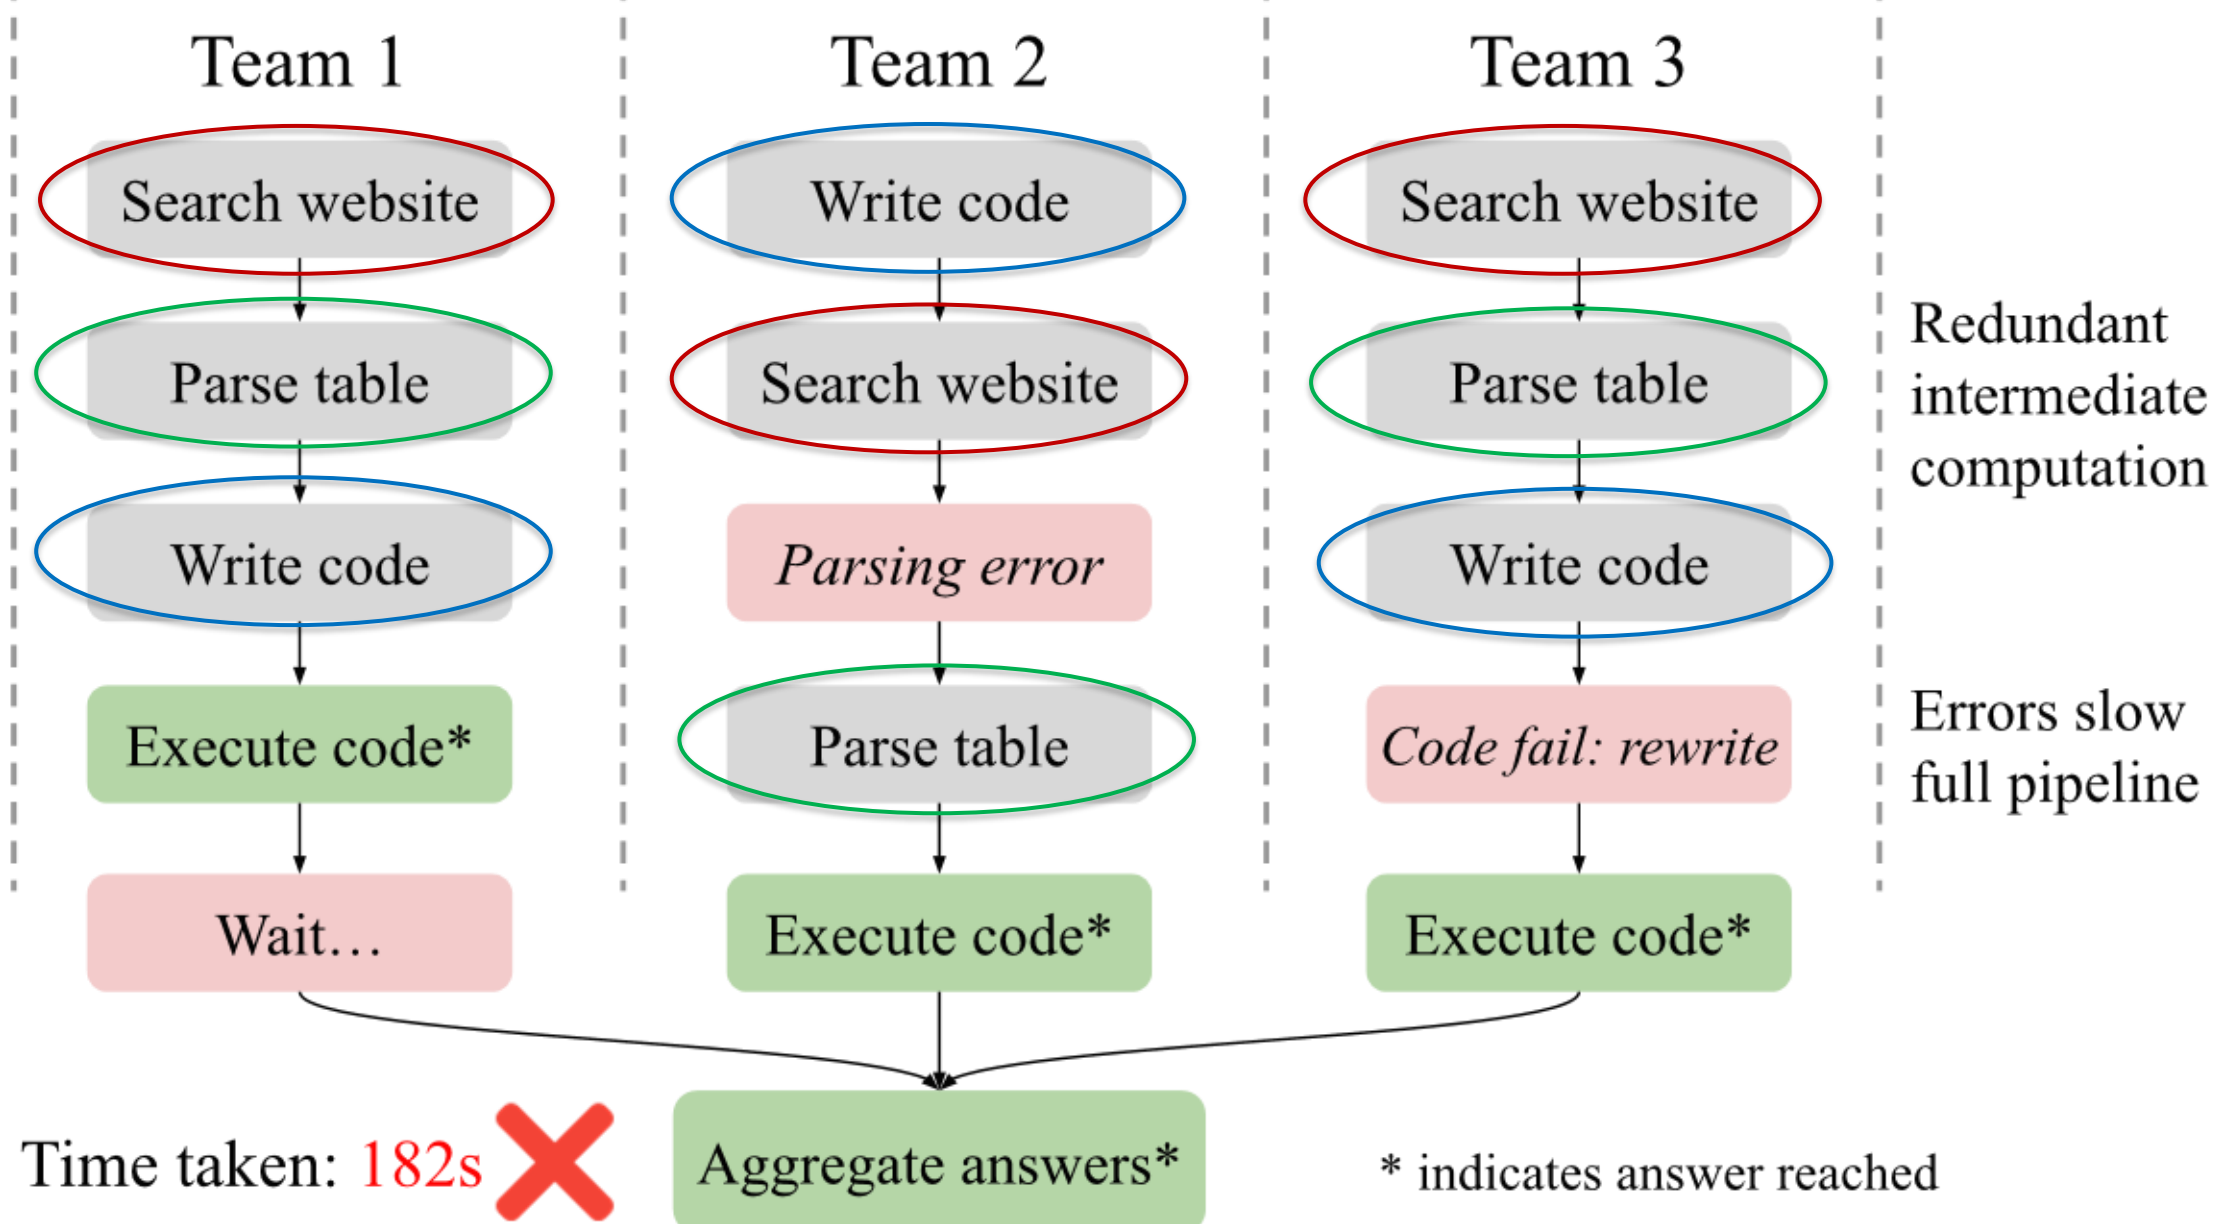


### Motivation – Parallel Execution

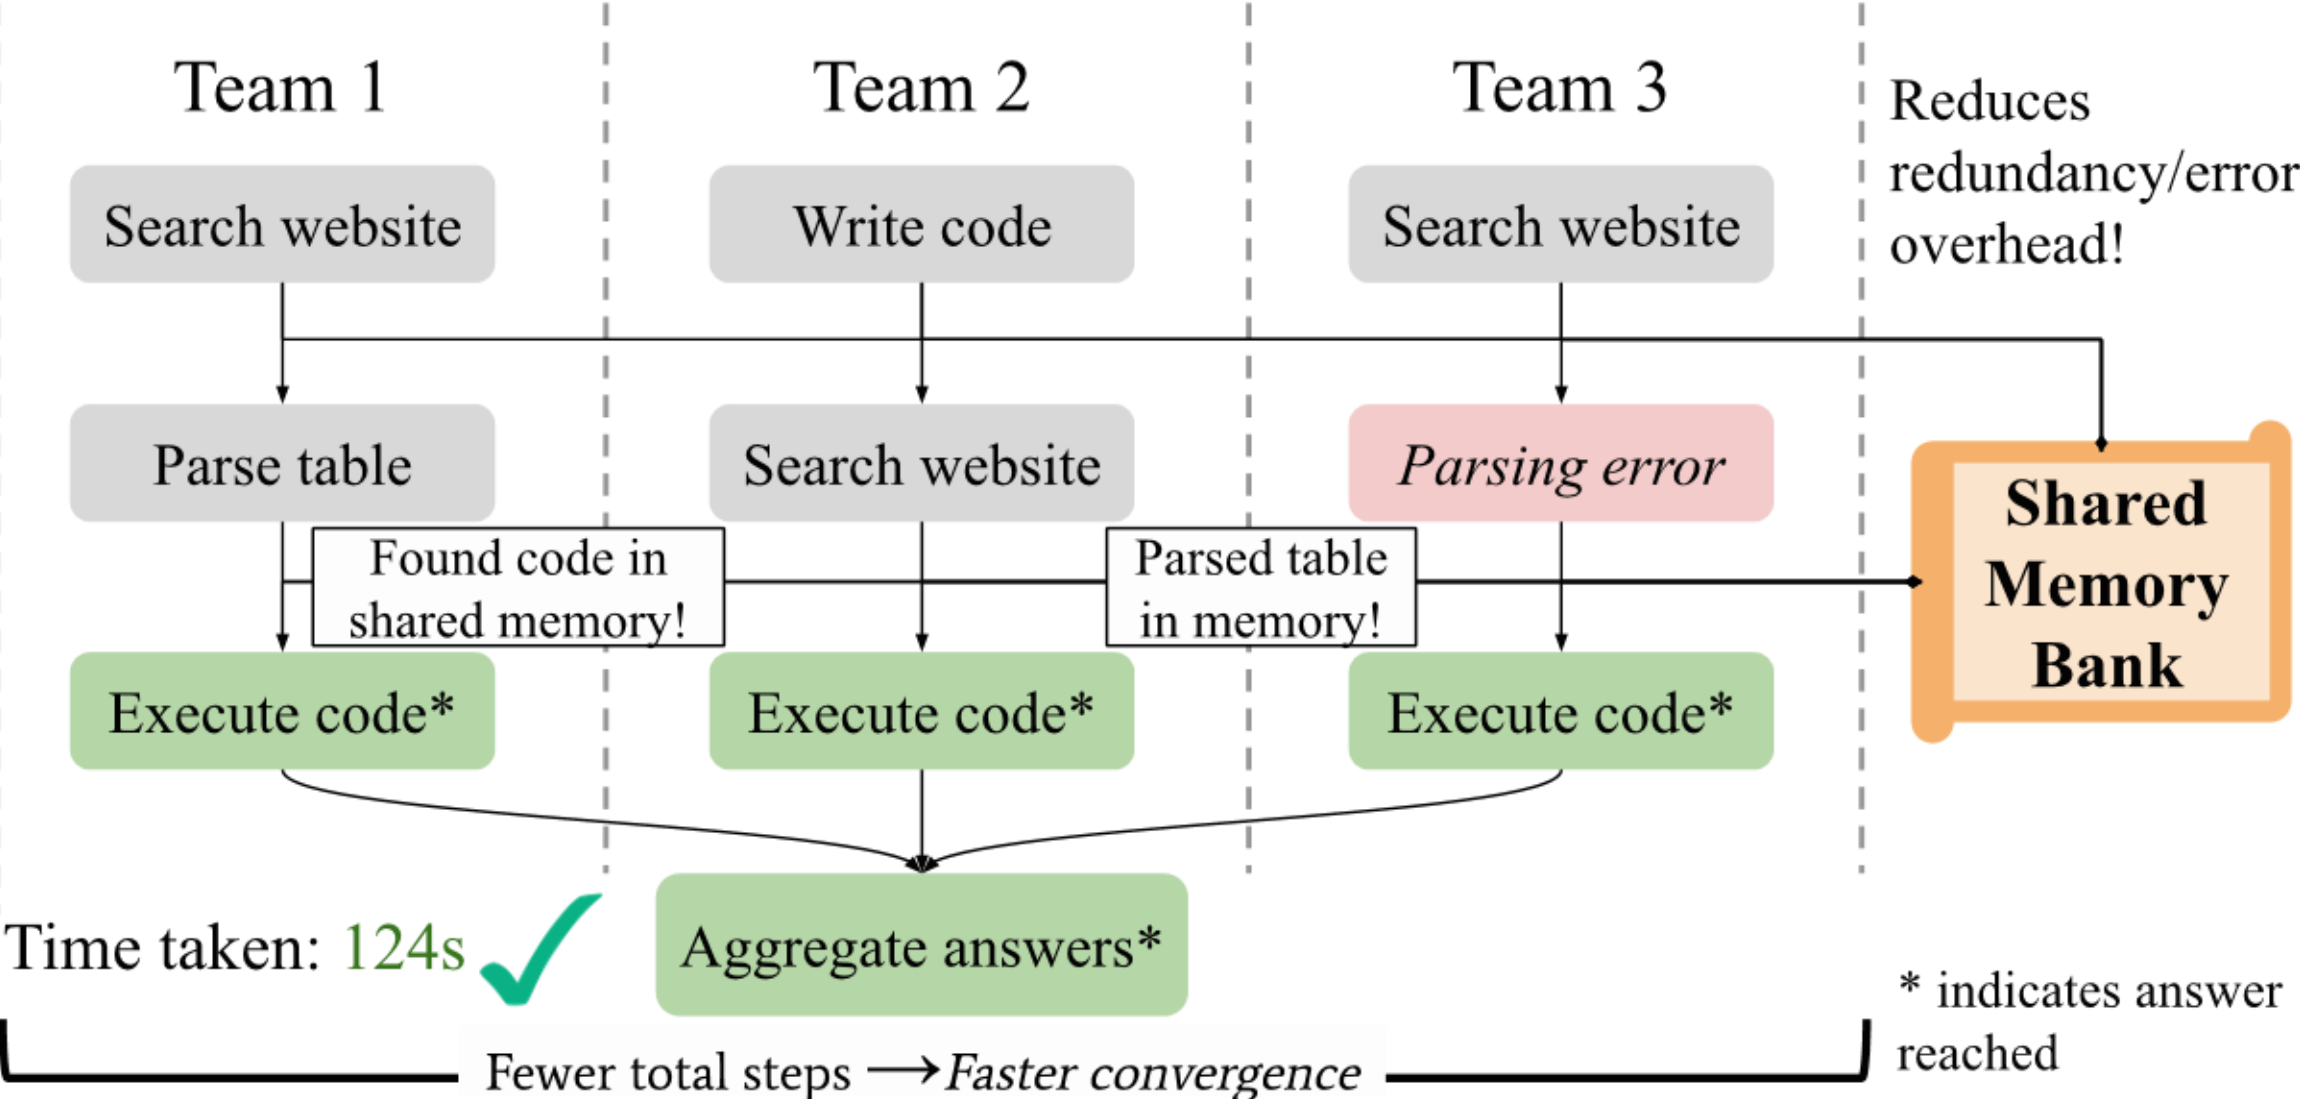
- ❖ Self-consistency improves performance in agentic tasks
  - ❖ However, teams are blind to each other; no sharing progress
  - ❖ Overlapping computation, isolated errors slow the full system
- ### Contributions
- ❖ Propose memory to reduce redundant parallel computation
  - ❖ Design RL-based strategy with usage-aware reward shaping
  - ❖ Show selective sharing improves efficiency AND performance

### Parallel Agentic Framework

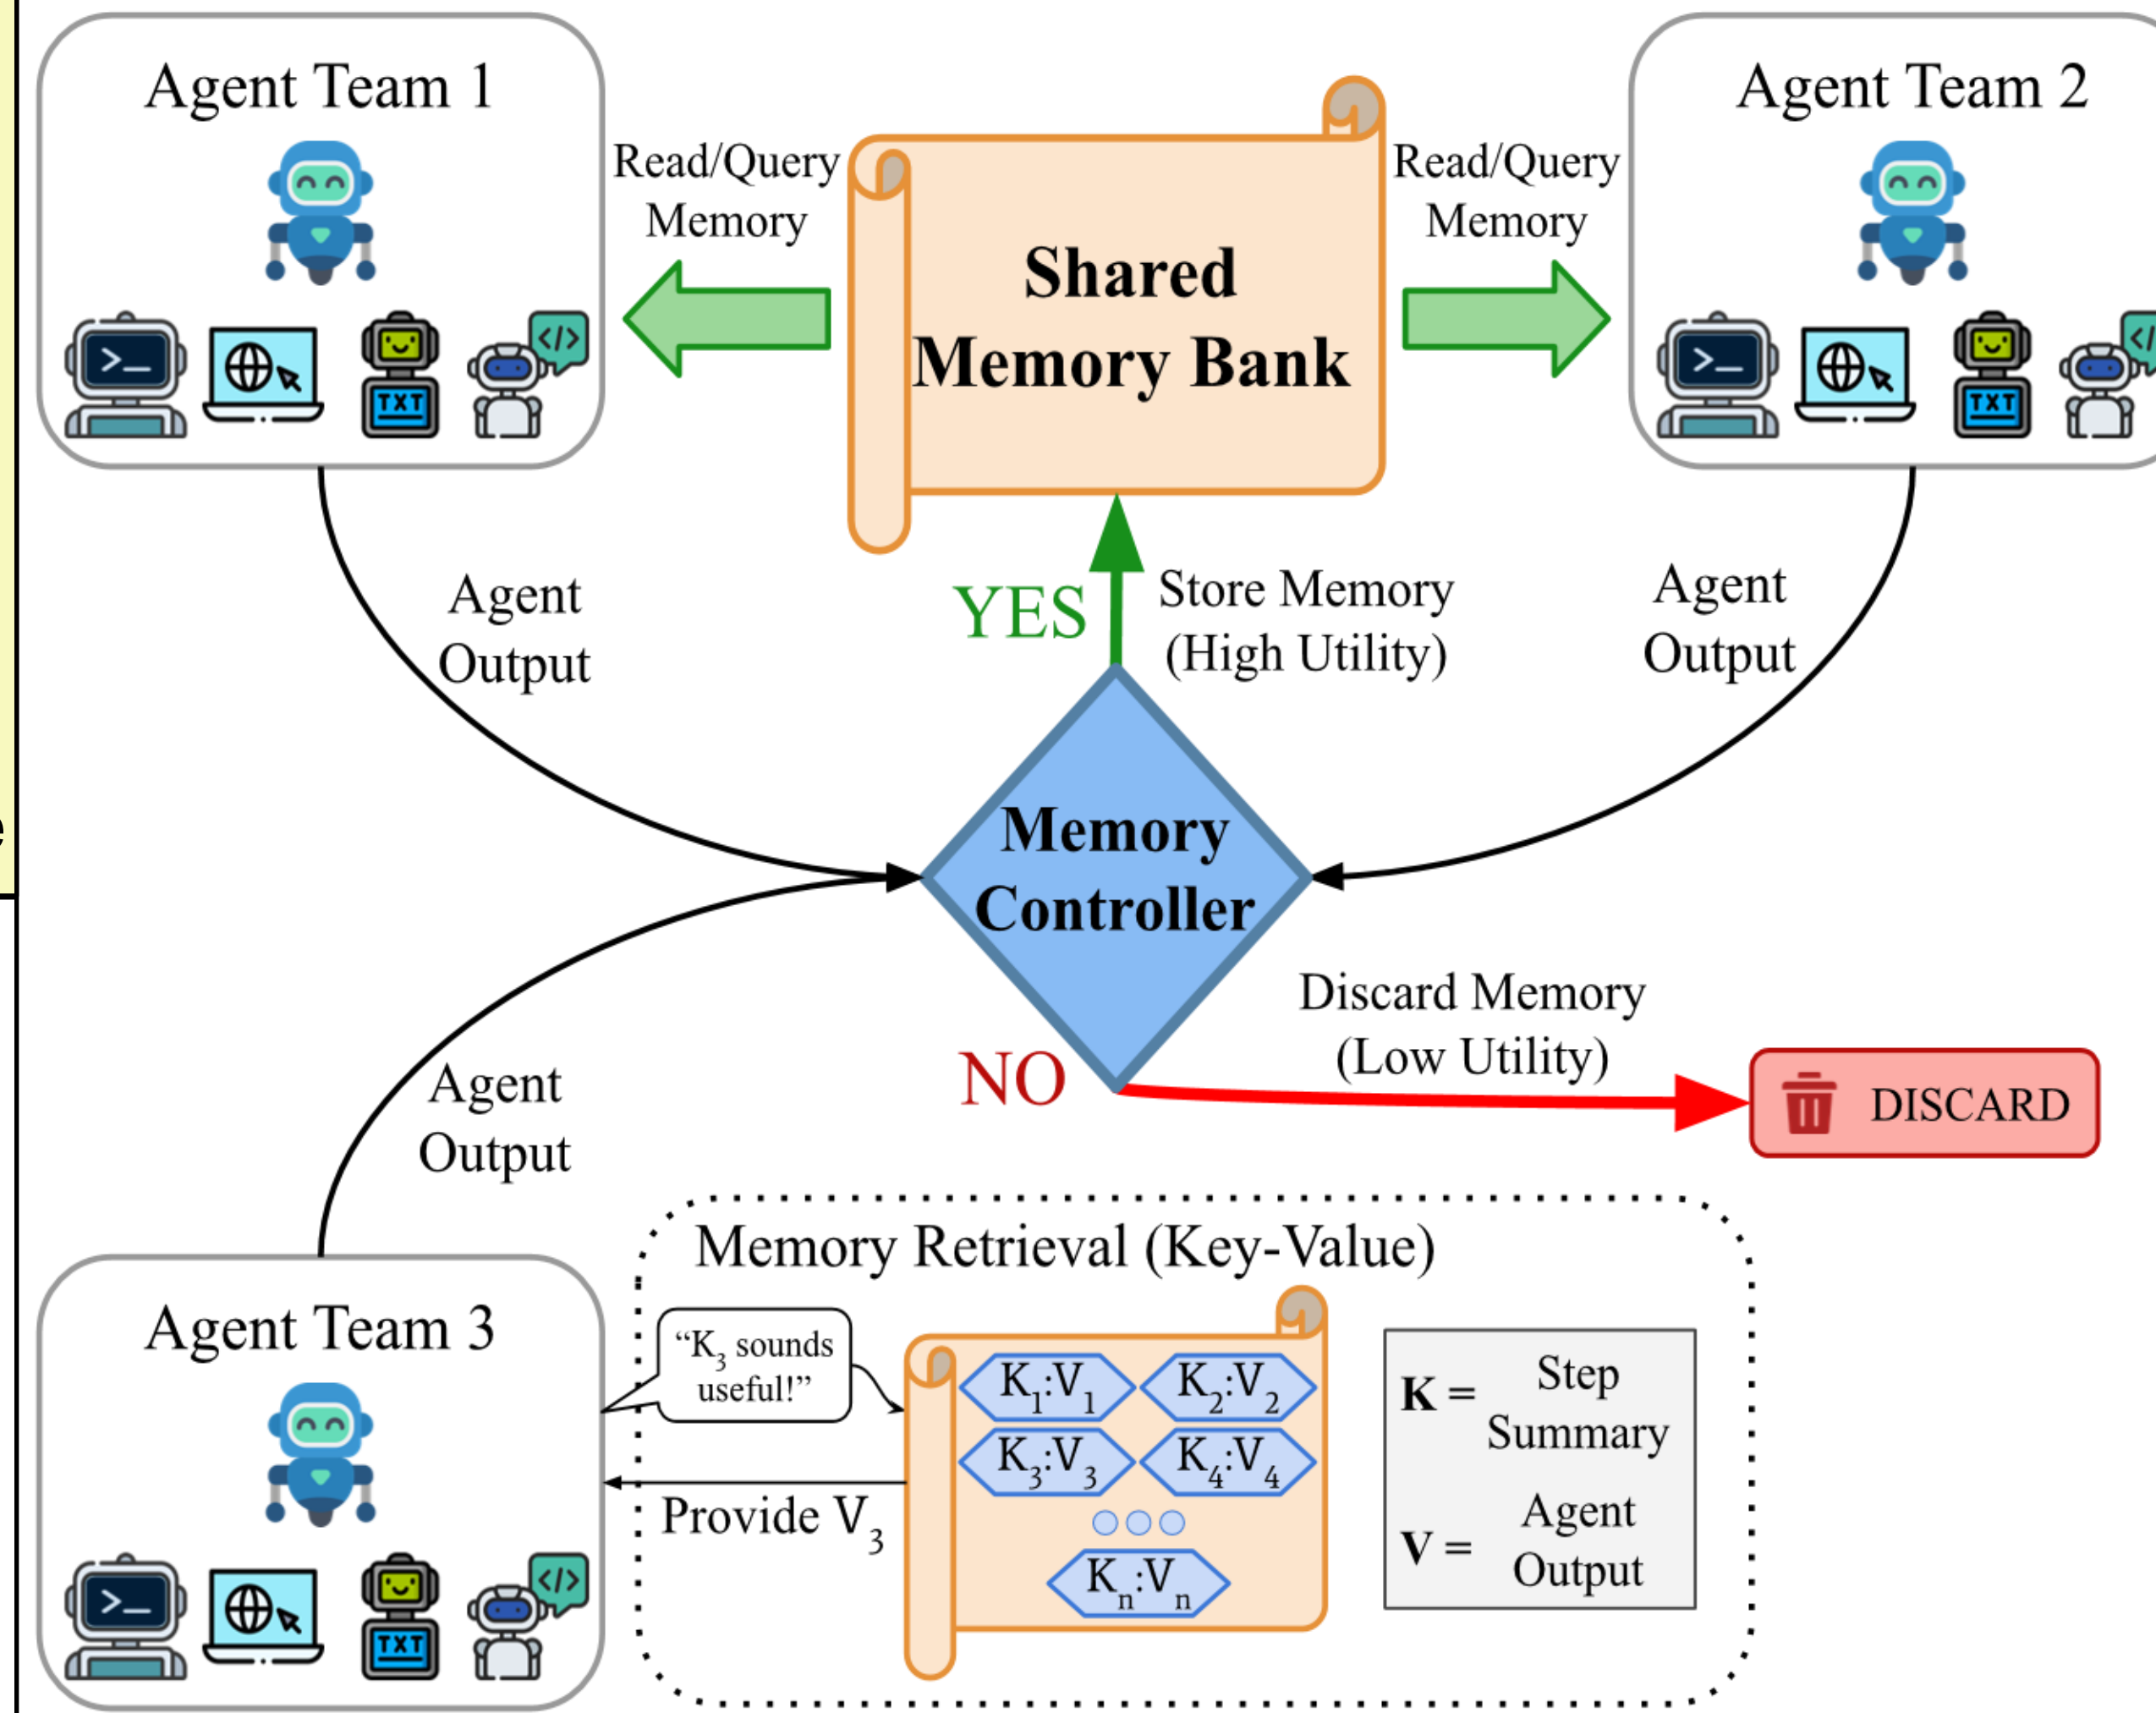
#### (a) Without Shared Memory



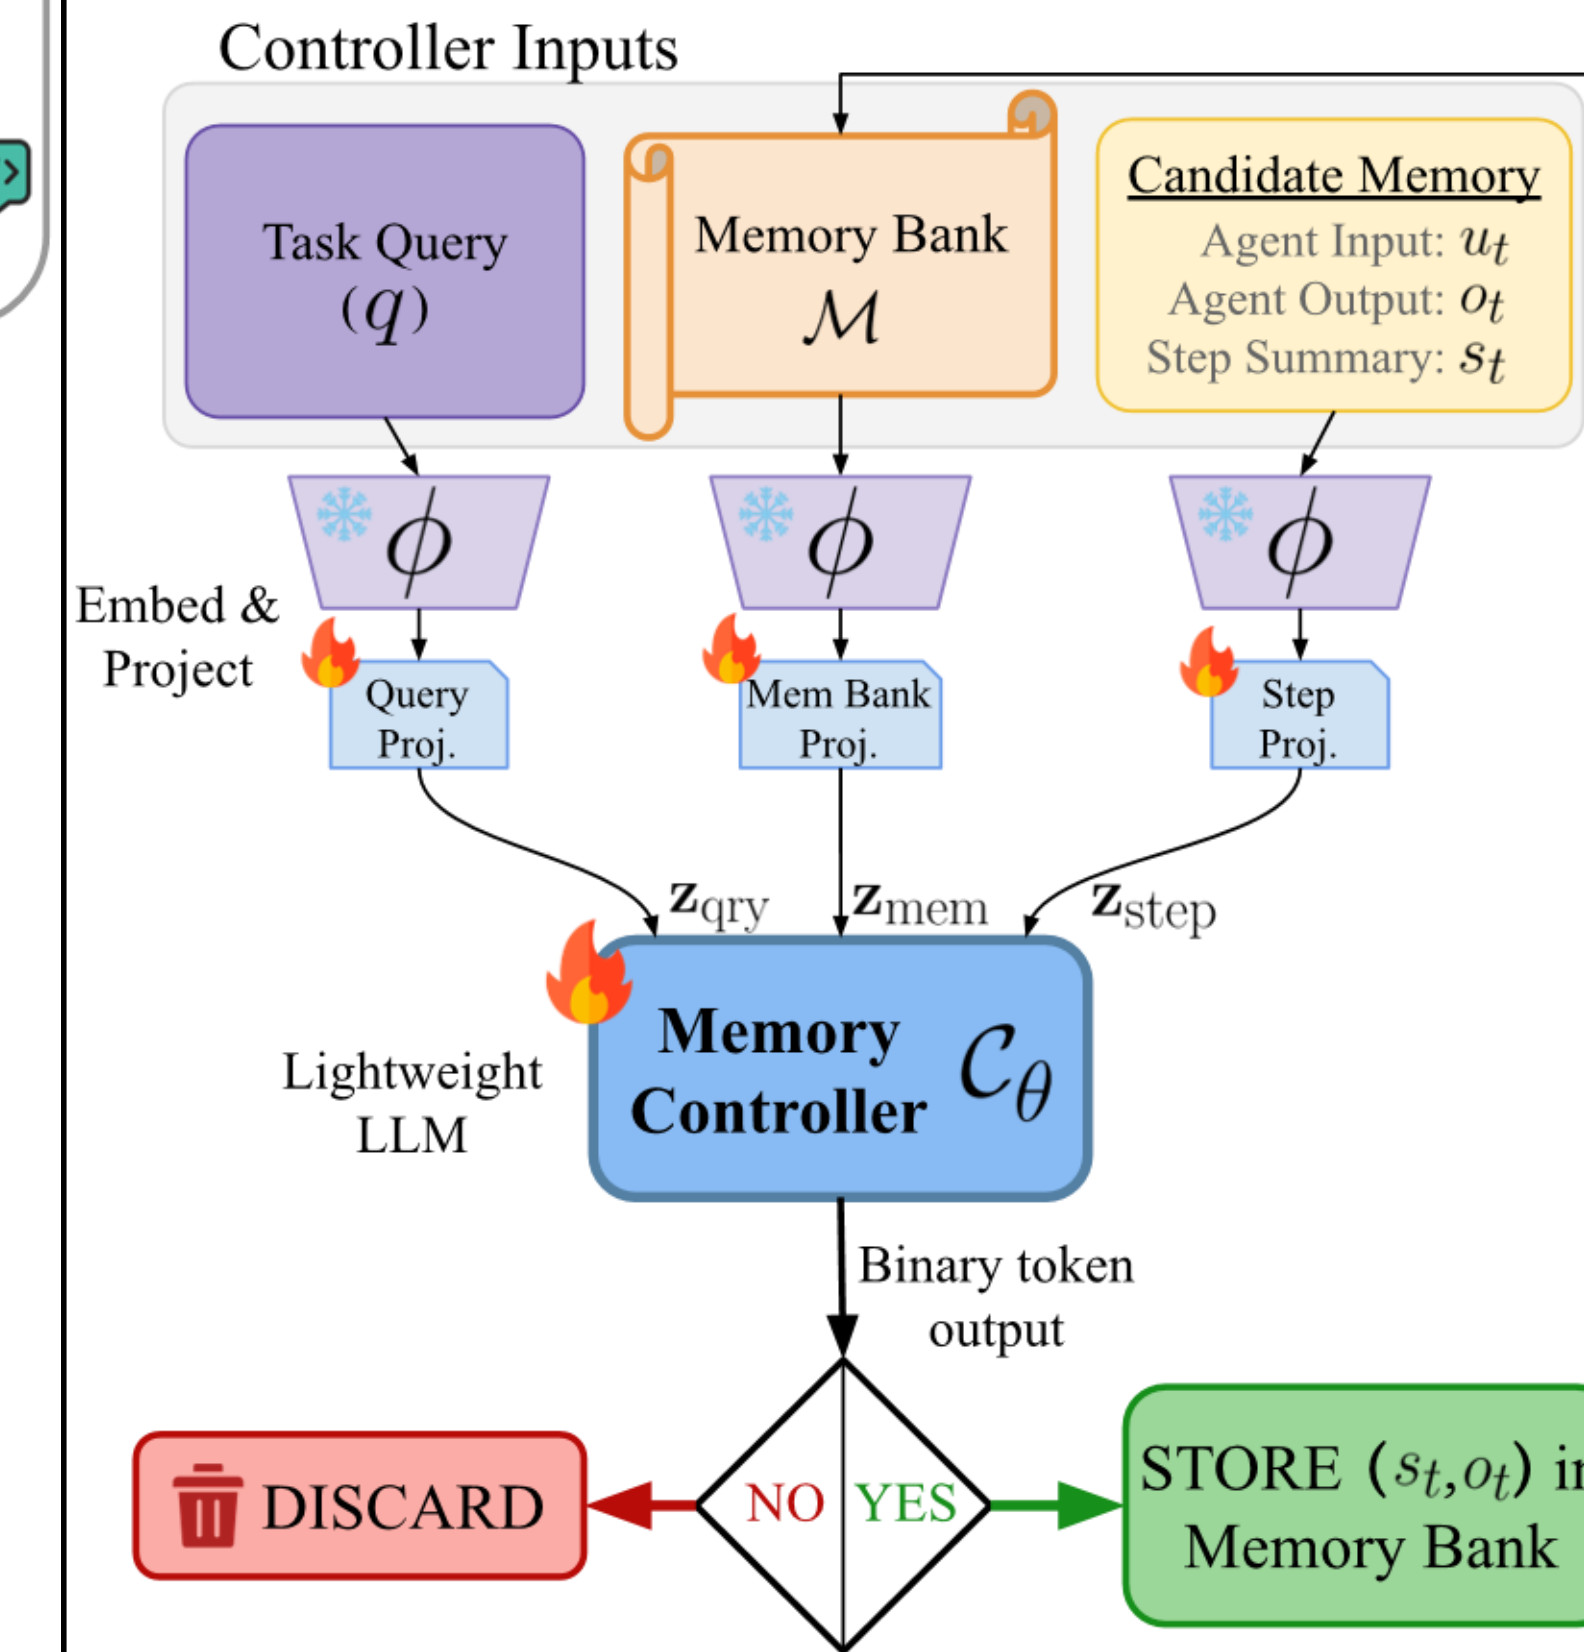
#### (b) With Shared Memory



### Workflow Diagram



### Memory Controller Design/Inputs



### Reward Modeling + Usage-Aware Shaping

Base reward:  $R(\tau) \in [0, 1]$  (benchmark score for trajectory  $\tau$ )

$$R(\tau) = R_{\text{agg}}(\tau) + \lambda_{\text{first}} R_{\text{first}}(\tau)$$

Reward of final aggregated answer      Scaled reward of first completing team's answer

$$\text{Per-trajectory advantage score: } A_{\text{base}}^{(i)} = \frac{R(\tau^{(i)}) - \mu_R}{\sigma_R + \epsilon}$$

**Shaped step-wise advantage function:**

$$\hat{A}_t^{(i)} = A_{\text{base}}^{(i)} + \beta \cdot \mathbb{I}(t \in \mathcal{U}^{(i)} \wedge R(\tau^{(i)}) > 0)$$

- Increases reward for saving retrieved memories

Advantage applied in **policy loss:**

$$L_t^{\text{policy}}(\theta) = -\log \pi_{\theta}(z_t | c_t) \cdot \hat{A}_t$$

**Sparsity penalty** to avoid degenerate “always YES” policy:

$$L_t^{\text{sparse}}(\theta) = \pi_{\theta}(z_t = \text{YES} | c_t)$$

- Admitting all steps increases usage reward (reward hacking)

**Complete optimization objective:**

$$\mathcal{L}(\theta) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[ \sum_{t=1}^T \left( L_t^{\text{policy}}(\theta) + \lambda_{\text{sparse}} L_t^{\text{sparse}}(\theta) \right) \right]$$

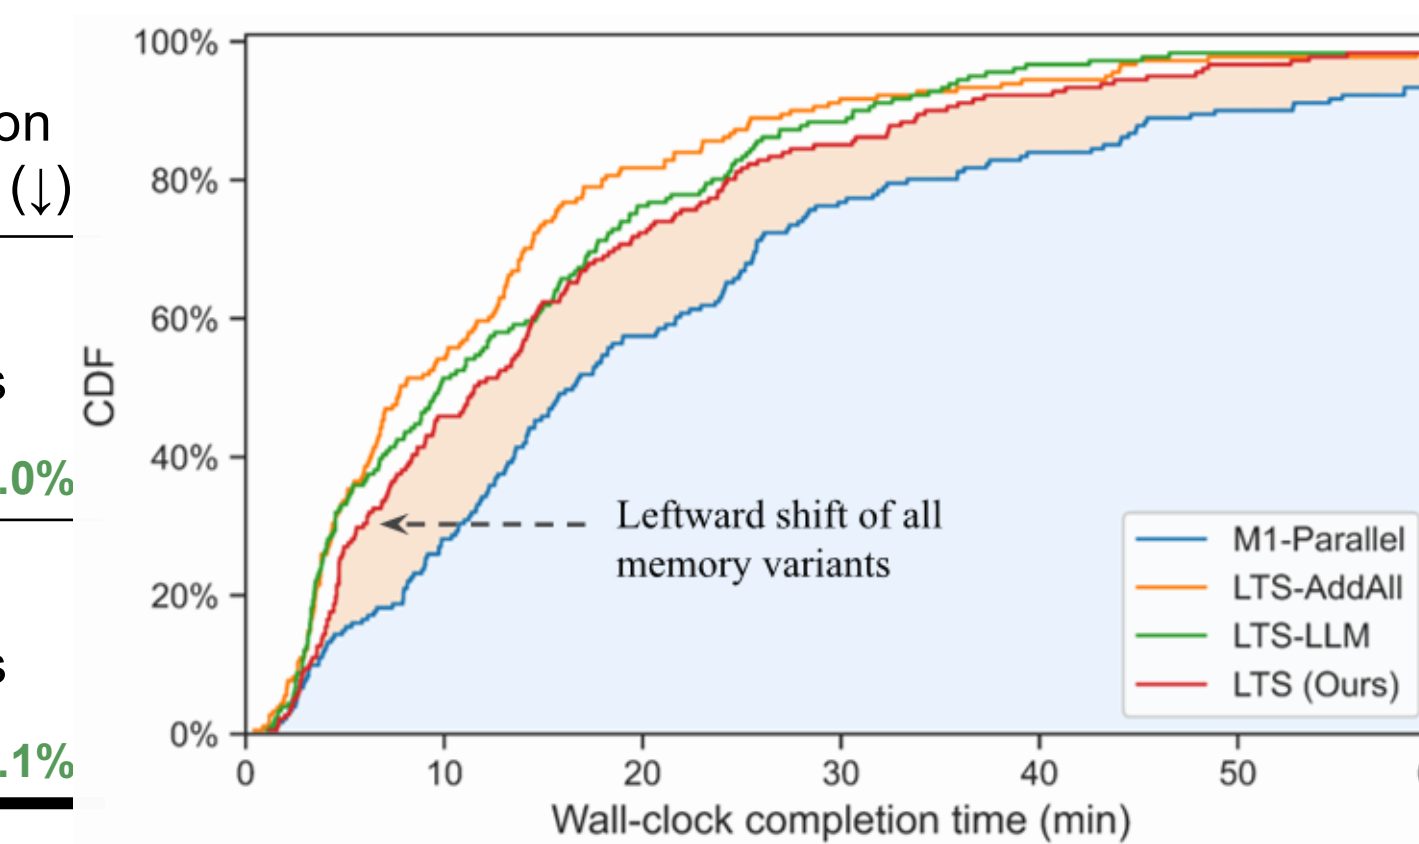
**Component Ablation:**

Variant	GAIA		
	Task Acc. (%)	Runtime (s)	Memories Saved (%)
<b>LTS (Ours)</b>	<b>49.1</b>	<b>792s</b>	<b>84.9</b>
w/o Usage-Aware Shaping	47.7	925s	89.5
w/o Sparsity Loss ( $\lambda_{\text{sparse}} = 0$ )	47.4	983s	98.9

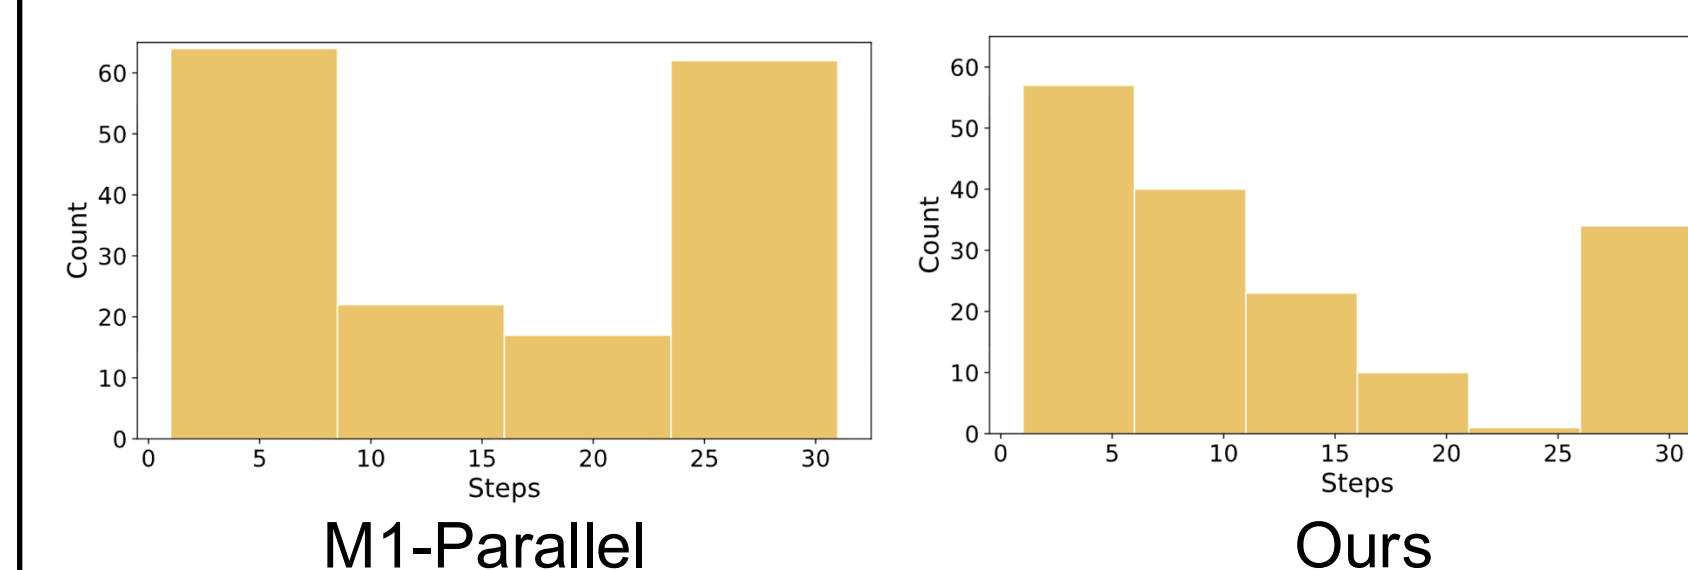
### Results

Method	Shared Memory	Model	AssistantBench				Execution Runtime (s)	GAIA				Execution Runtime (s)
			Easy Acc.	Med. Acc.	Hard Acc.	All Acc.		Lvl 1 Acc.	Lvl 2 Acc.	Lvl 3 Acc.	All Acc.	
MagenticOne [1]	<b>X</b>	Qwen3 (32B)	53.1	17.4	8.5	13.4	1084s	41.5	16.3	3.8	22.5	758s
M1-Parallel [2]	<b>X</b>		52.6	21.4	8.5	14.7	2239s	43.4	18.6	7.7	24.8	1569s
<b>LTS (Ours)</b>	<b>✓</b>		<b>65.8</b>	<b>25.2</b>	<b>14.4</b>	<b>20.3</b>	<b>1479s</b> ↓40.8%	<b>47.2</b>	<b>25.6</b>	<b>11.8</b>	<b>30.4</b>	<b>892s</b> ↓55.0%
MagenticOne [1]	<b>X</b>	GPT-5.1	43.6	31.5	15.3	21.7	724s	48.1	37.3	16.7	37.5	815s
M1-Parallel [2]	<b>X</b>		57.3	<b>35.2</b>	16.0	24.0	1389s	60.4	46.5	<b>26.9</b>	47.9	1005s
<b>LTS (Ours)</b>	<b>✓</b>		<b>61.0</b>	<b>35.1</b>	<b>20.0</b>	<b>26.7</b>	<b>882s</b> ↓44.6%	<b>62.3</b>	<b>47.7</b>	<b>26.9</b>	<b>49.1</b>	<b>781s</b> ↓25.1%

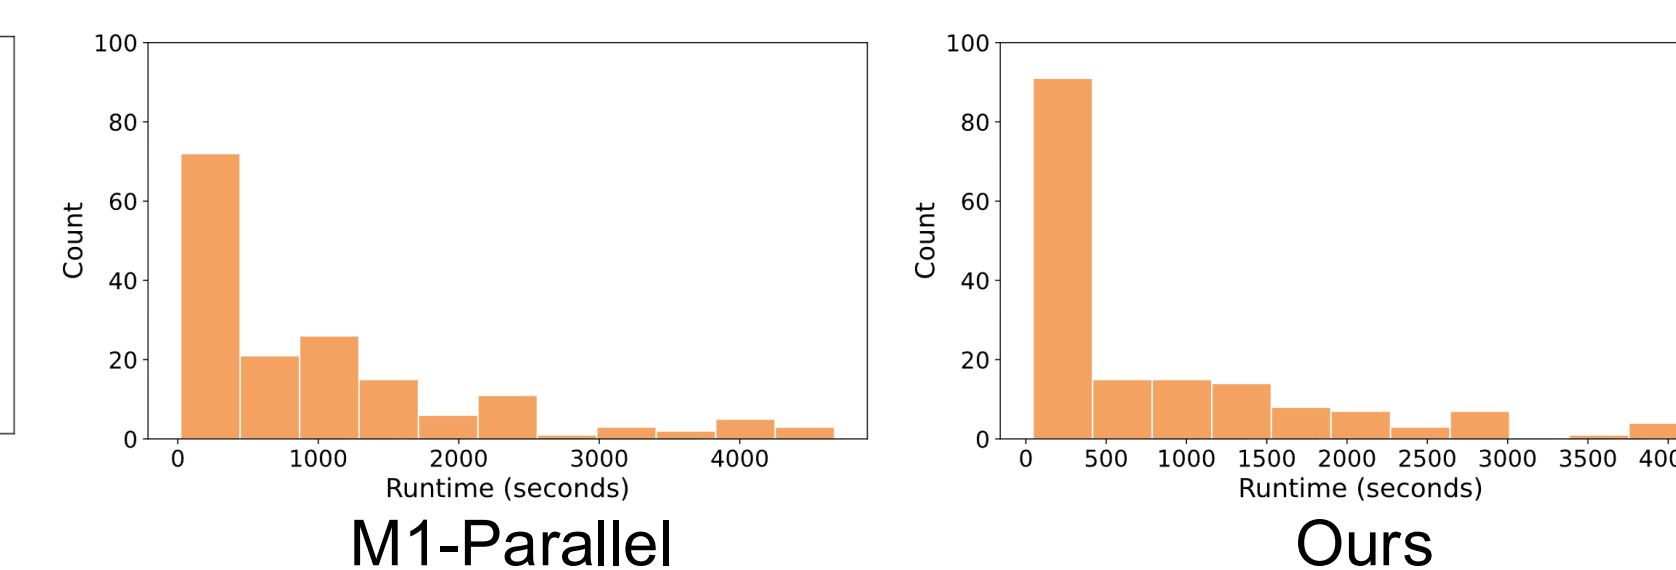
### Runtime CDF



### Steps Taken / Task Distribution



### Runtime / Task Distribution



**Shared Memory Improves Average Runtime!**

### References

- [1] Fourney, Adam, et al. "Magentic-one: A generalist multi-agent system for solving complex tasks." arXiv preprint arXiv:2411.04468 (2024).  
[2] Zhang, Enhao, et al. "Optimizing sequential multi-step tasks with parallel llm agents." arXiv preprint arXiv:2507.08944 (2025).

**Webpage/Code:**

