



SILO-BENCH: A Scalable Environment for Evaluating Distributed Coordination in Multi-Agent LLM Systems

Yuzhe Zhang¹ Feiran Liu¹ Yi Shan¹ Xinyi Huang¹
 Xin Yang² Yueqi Zhu¹ Xuxin Cheng³ Cao Liu³ Ke Zeng³
 Terry Jingchen Zhang^{4†} Wenyuan Jiang^{4†}

¹ Beijing University of Technology, Beijing, China

² Zhejiang University, Hangzhou, China ³ Meituan LongCat Interaction Team

⁴ ETH Zürich, Switzerland

Abstract

Large language models are increasingly deployed in multi-agent systems to overcome context limitations by distributing information across agents. Yet whether agents can reliably **compute** with distributed information, rather than merely exchange it, remains an open question. We introduce SILO-BENCH, a role-agnostic benchmark of 30 algorithmic tasks across three communication complexity levels, evaluating 54 configurations over 1,620 experiments. Our experiments expose a fundamental **Communication-Reasoning Gap**: agents spontaneously form task-appropriate coordination topologies and exchange information actively, yet systematically fail to synthesize distributed state into correct answers. The failure is localized to the reasoning-integration stage where agents often acquire sufficient information but cannot integrate it. This coordination overhead compounds with scale, eventually eliminating parallelization gains entirely. These findings demonstrate that naively scaling agent count cannot circumvent context limitations, and SILO-BENCH provides a foundation for tracking progress toward genuinely collaborative multi-agent systems. The code is available at <https://github.com/jwyjohn/acl26-silo-bench>.

1 Introduction

The rapid advancement of Large Language Models (LLMs) has demonstrated remarkable capabilities in individual inference and generation tasks (OpenAI, 2023; Touvron et al., 2023; DeepSeek-AI, 2024). However, as the scale and complexity of real-world problems continue to grow, a fundamental bottleneck has emerged: the limited context window of a single model restricts its ability to process global information (Li et al., 2024; Chen et al., 2024b; An et al., 2024). Even with recent progress in extending context lengths to millions of

[†]Corresponding Author. Email: wenyjiang@ethz.ch

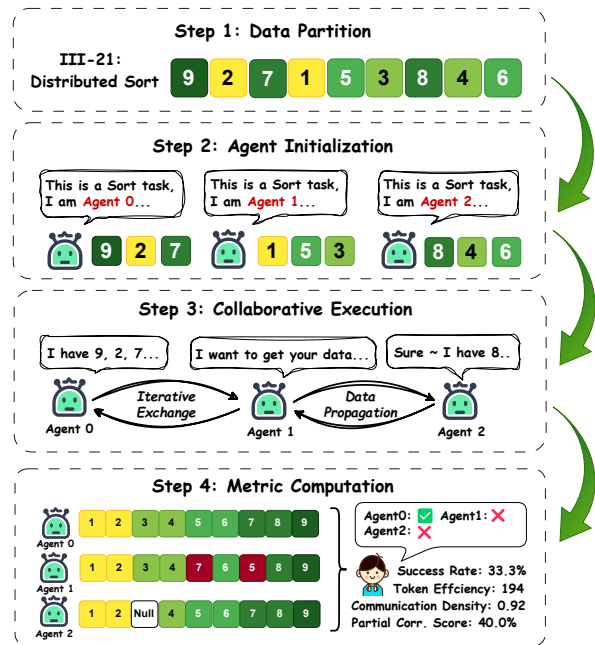


Figure 1: Pipeline of SILO-BENCH. Global information is partitioned across N agents, each holding only local data. Agents must communicate through the provided protocol to reconstruct global truth. Success requires effective collaboration strategies. This is an example of the III-21 Distributed Sort (Appendix E.)

tokens (Reid et al., 2024) as well as improving context robustness (Yang et al., 2026), the quadratic cost of attention (Ratner et al., 2023; Yen et al., 2024) makes centralized processing increasingly impractical for truly large-scale tasks.

Multi-Agent Systems (MAS) offer a compelling architectural paradigm to address this scalability challenge (Zhang et al., 2024a; Wang et al., 2024a, 2025b). By distributing global information across multiple agents that collaborate to compute results, MAS can theoretically overcome the token limitations of single models (Liu et al., 2024b). This distributed approach mirrors successful patterns in traditional computing, from MapReduce to modern distributed databases, where data partitioning

and coordinated computation across nodes achieve scales unattainable by a single machine (Dean and Ghemawat, 2008). In the realm of large models, we define the scenario where an individual agent has access only to partial information, thereby necessitating coordination to resolve token constraints, as information silos. However, a critical question remains underexplored: Can current LLM-based agents effectively collaborate within information silos to compute a globally correct answer (Qian et al., 2024a,b; Liu et al., 2024b)?

Existing multi-agent benchmarks either prescribe fixed communication structures (Li et al., 2023; Wu et al., 2023; Hong et al., 2024) or focus on social simulation rather than computational collaboration (Park et al., 2023; Lan et al., 2024). These approaches often introduce inductive bias into the agents’ final outputs (Baltaji et al., 2024). For instance, if an agent is assigned the role of a "doctor", it may exhibit poor performance in artistic domains, which contradicts the goal of developing general-purpose agents (An et al., 2024; Qian et al., 2024a,b; Liu et al., 2024b). Furthermore, most benchmarks to date target specific tasks (Deng et al., 2024; Chen et al., 2024a; Gioacchini et al., 2024; Wang et al., 2025c) and fail to address a significant gap in our understanding: whether logic-based models can autonomously discover and execute effective coordination strategies for distributed computing problems. Addressing this gap is a key objective for future AGI evaluation.

To bridge this gap, we propose SILO-BENCH—a pioneering benchmark for evaluating free-form communication and collaboration in multi-agent LLM systems (Liu et al., 2024b). In summary, our contributions are as follows:

- **We introduce SILO-BENCH, a role-agnostic configurable environment for evaluating distributed coordination under information silos.** Unlike static test suites that prescribe fixed roles and communication scripts, our framework can generate unlimited evaluation instances while providing high-level task hints and allowing observation of whether agents can translate structural understanding into effective coordination protocols (Press et al., 2022; Liu et al., 2024b).
- **We conduct the largest systematic study of multi-agent collaboration to date by instantiating 54 representative configurations.** Spanning diverse protocols and computing paradigms

(Zhao et al., 2024; Islam et al., 2024a), we propose a multi-dimensional metric suite to comprehensively quantify the trade-off between task success rate, token consumption, and communication density.

- **We expose critical scalability limitations and the *Communication-Reasoning Gap* in current LLMs.** Our results reveal that while agents can spontaneously discover task-appropriate communication topologies, they fail to translate effective coordination into correct distributed computation. This disconnection, coupled with inefficient information synthesis, causes performance to collapse as task complexity increases and the agent scale expands.

2 Related Work

Context Limitations and Distributed Reasoning.

The finite context window of LLMs constitutes a fundamental bottleneck for processing large-scale information. While recent advances have extended context lengths to millions of tokens (Reid et al., 2024; Liu et al., 2025), the quadratic computational complexity of attention mechanisms makes centralized processing increasingly resource-intensive and prone to “lost-in-the-middle” phenomena (Liu et al., 2024a). Although Retrieval-Augmented Generation (RAG) offers a palliative solution (Wang et al., 2024b; Islam et al., 2024b; Ji et al., 2025), it often fractures global context, struggling with tasks that require holistic reasoning across disjoint segments. Existing benchmarks like SCROLLS (Shaham et al., 2022), LongBench (Bai et al., 2024), and ∞ Bench (Zhang et al., 2024b) effectively evaluate single-agent retrieval but overlook the paradigm of *distributed collaboration*. We posit that overcoming the context barrier requires shifting from centralized attention to collaborative computation, where agents act as distributed processors, to digest partitioned information and synthesize global insights, which is a capability currently unmeasured by standard long-context evaluations.

Multi-Agent Architectures and Role-Agnosticism.

The paradigm of orchestrating multiple LLM agents has evolved from simple role-playing to complex problem-solving frameworks. Foundational works like CAMEL (Li et al., 2023) and MetaGPT (Hong et al., 2024) utilize *role-specialized agents* (e.g., assigning “Manager” or “Coder” personas) embedded within




Figure 2: Three complexity levels in SILO-BENCH characterized by their communication patterns. **Level I (Aggregation)**: A central agent collects data from all peers via a star topology. **Level II (Mesh Network)**: Agents exchange information with immediate neighbors through pairwise communication. **Level III (Global Shuffle)**: All agents must communicate with every other agent, requiring full mesh connectivity.

fixed hierarchical or waterfall workflows. While effective for domain-specific tasks like software engineering (Islam et al., 2024a) and interactive image generation (Ma et al., 2026), these approaches entangle the agents’ reasoning capabilities with semantic role priors, making it difficult to isolate the contribution of the communication architecture itself. Other efforts, such as debate-based systems (Du et al., 2024), Mixture-of-Agents (Wang et al., 2025a; Su et al., 2026), or one-shot topology generation frameworks (Sun et al., 2026), often prescribe static topological constraints or rely on offline topology optimization, limiting dynamic information flow. **We introduce SILO-BENCH, a role-agnostic configurable environment with task-structural guidance for evaluating distributed coordination under information silos.** Unlike static test suites that prescribe fixed roles and communication scripts, our framework dynamically generates unlimited evaluation instances while providing high-level task hints, allowing observation of whether agents can translate structural understanding into effective coordination protocols (Press et al., 2022; Liu et al., 2024b).

3 SILO-BENCH

This section presents the architecture of SILO-BENCH, a configurable environment for evaluating multi-agent collaboration under information silos. Each configuration is defined by three orthogonal dimensions: agent scale N , communication protocol P , and language model \mathcal{M} . We describe the task space, evaluation metrics, and execution pipeline.

3.1 Task Space

A central design goal of SILO-BENCH is to ground task difficulty in principled communication complexity theory, so that observed performance gaps can be attributed to coordination demands rather than ad hoc task choice. The theoretical foundation for analyzing distributed computation costs dates back to Yao’s seminal work on communication complexity (Yao, 1979), which established the framework for quantifying the minimum bits required for distributed parties to compute a function. Building on this foundation, we categorize tasks by their optimal communication complexity:

$$\tau_k = (f_k, \mathcal{X}_k, y_k^*) \quad (1)$$

where f_k specifies the computational function, \mathcal{X}_k is the global input data, and y_k^* is the ground-truth answer. Tasks are organized into three levels based on their optimal communication complexity (complete task specifications are provided in Appendix E).

Level I: Aggregation ($\mathcal{O}(N)$ communication). As illustrated in Figure 2 (left), these tasks exhibit embarrassingly parallel structure followed by reduction. Each agent processes its local shard independently, producing intermediate results aggregated through associative operations (e.g., max, sum, xor). The optimal topology is a star or tree structure where one agent collects all partial results. Representative tasks include global maximum (LC-414: “Third Maximum Number”), distributed voting (LC-169: “Majority Element”), and word frequency counting (LC-2085).

Level II: Mesh Network ($\mathcal{O}(N)$ communication). As shown in Figure 2 (center), these tasks exhibit

spatial locality: agent i 's computation depends primarily on neighboring agents $i - 1$ and $i + 1$. Information propagates through a structured mesh via pairwise exchanges, with optimal topology being a linear chain requiring $N - 1$ point-to-point exchanges. Representative tasks include prefix sum (LC-1480), moving average (LC-346: "Moving Average from Data Stream"), and trapping rain water (LC-42).

Level III: Global Shuffle ($\mathcal{O}(N \log N)$ to $\mathcal{O}(N^2)$ communication). As depicted in Figure 2 (right), these tasks feature irregular, potentially all-to-all communication patterns where any agent's output may depend on information from any other agent. The range $\mathcal{O}(N \log N)$ – $\mathcal{O}(N^2)$ spans from the classical lower bound for distributed reorganization to the full-consensus cost imposed by our evaluation criterion, where every agent must output the complete global answer. Representative tasks include distributed sorting (LC-912), graph connectivity (LC-323), and matrix multiplication (LC-311).

3.2 Task Construction Pipeline.

LeetCode problems serve solely as algorithmic inspiration and we *do not* transform raw LeetCode data. For each task *category* (e.g., "Global Maximum"), we independently implement a Python generator that programmatically produces random inputs and exact ground-truth answers. A *task instance* is one concrete input–output pair drawn from this generator under a fixed (N, P, \mathcal{M}) configuration, where N is the agent scale, P the communication protocol, and \mathcal{M} the language model, and a fixed random seed, ensuring reproducibility while allowing unlimited fresh instances.

To illustrate: for Level-I Global Maximum (inspired by LC-414, the "Third Maximum Number" problem), given agent count N and per-agent shard size k , the generator (i) samples $N \times k$ integers uniformly at random, (ii) partitions them into N equal shards $\mathcal{X}_1, \dots, \mathcal{X}_N$, (iii) computes $y^* = \max(\mathcal{X})$, and (iv) records the fixed seed for reproducibility. Each agent receives only its local shard \mathcal{X}_i and must coordinate to determine y^* . This pipeline generalizes directly to all 30 tasks: the generator encodes the task-specific function f_k , scales global input size proportionally with N to maintain constant per-agent workload, and produces exact ground-truth answers enabling fully objective evaluation.

3.3 Evaluation Metrics

We define four complementary metrics to capture both *what* agents achieve and *how* they coordinate. Let \hat{y}_i denote agent i 's submitted answer, and let m_i denote the total number of messages successfully transmitted outward by agent i during the entire collaboration.

Success Rate (\mathcal{S}). Measures the proportion of agents converging to the correct answer:

$$\mathcal{S} = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[\hat{y}_i = y^*] \quad (2)$$

A task instance is *successful* when $\mathcal{S} = 1$, indicating unanimous convergence.

Partial Correctness Score (\mathcal{P}). Binary success rate can understate partial progress. We introduce a continuous measure of answer quality tailored to each task category: for Level-I, \mathcal{P} is the fraction of agents within tolerance of the ground truth; for Level-II, the fraction of correctly computed elements per local segment; for Level-III, the longest correctly ordered subsequence relative to total length. Letting $q_i \in [0, 1]$ denote the per-agent quality score:

$$\mathcal{P} = \frac{1}{N} \sum_{i=1}^N q_i \quad (3)$$

Together with \mathcal{S} , this score allows us to isolate where coordination breaks down: the gap $\mathcal{P} - \mathcal{S}$ quantifies performance lost specifically at the reasoning-integration stage rather than at the communication stage.

Token Consumption (\mathcal{C}). Quantifies computational cost per communication round:

$$\mathcal{C} = \frac{\sum_{i=1}^N \sum_{r=1}^R t_i^{\text{out}}[r]}{R_{max}} \quad (4)$$

where $t_i^{\text{out}}[r]$ is the number of output tokens generated by agent i in round r , and R_{max} is the maximum number of rounds executed.

Communication Density (\mathcal{D}). Captures inter-agent interaction intensity. Here $N(N - 1)$ is the directed-edge count when each ordered pair exchanges exactly one message; since agents may send multiple messages to the same recipient across rounds, $\mathcal{D} \in [0, +\infty)$:

$$\mathcal{D} = \frac{\sum_{i=1}^N m_i}{N(N - 1)} \quad (5)$$

Values near 0 suggest sparse, targeted exchanges; $\mathcal{D} = 1$ indicates one message per directed pair on average; values exceeding 1 reflect iterative multi-round exchanges. For the SFS protocol (see Appendix A), m_i counts the number of times other agents successfully read files written by agent i , preserving the same “information actually transferred” semantics as direct message-passing.

Together, \mathcal{S} and \mathcal{P} measure *what* agents achieve, \mathcal{C} measures *at what cost*, and \mathcal{D} reveals *how* they coordinate.

3.4 Execution Pipeline

Given task $\tau = (f, \mathcal{X}, y^*)$ and configuration (N, P, \mathcal{M}) , the evaluation proceeds through four phases.

Phase 1: Data Partition. $\text{PARTITION}(\mathcal{X}, N) \rightarrow \{\mathcal{X}_1, \dots, \mathcal{X}_N\}$, where $|\mathcal{X}_i| \approx |\mathcal{X}|/N$ ensures equal partition and no agent holds privileged information.

Phase 2: Agent Initialization. Each agent i is initialized with model \mathcal{M} and receives $\text{INIT}(i) \leftarrow (\text{desc}(f, \mathcal{X}_i), P)$, specifying the core task logic, local data, and protocol constraint. The prompt provides task-structural guidance while preserving strategic autonomy (see Appendix B).

Phase 3: Collaborative Execution. Agents engage in iterative communication for up to R_{\max} rounds. All N agents are activated in parallel within each round: they receive incoming messages from the previous round, independently decide on actions, and execute them simultaneously. Messages or files written in round r become visible at the start of round $r + 1$. Execution terminates when all agents submit answers or the round limit is reached.

Phase 4: Metric Computation. The four metrics $(\mathcal{S}, \mathcal{P}, \mathcal{C}, \mathcal{D})$ are computed from submitted answers $\{\hat{y}_i\}_{i=1}^N$ and recorded communication logs.

4 Experiments

We systematically evaluate multi-agent coordination across three orthogonal axes: (i) agent scale, (ii) communication protocol, and (iii) language model, yielding a factorial design that covers qualitatively distinct coordination regimes.

4.1 Experimental Setup

Each evaluation instance in SILO-BENCH is specified by agent scale N , communication protocol P ,




Figure 3: The three communication protocols employed in SILO-BENCH.

and language model \mathcal{M} . All models are deployed locally with default temperature and 128K context windows.

Agent Scale (N). We vary team size across $N \in \{2, 5, 10, 20, 50, 100\}$, chosen to probe qualitatively distinct coordination regimes. The minimal team ($N = 2$) isolates fundamental pairwise coordination without overhead. Small groups ($N \in \{5, 10\}$) allow agents to feasibly track all peers simultaneously. Medium scale ($N = 20$) begins to make exhaustive peer tracking challenging, pushing agents toward selective communication. Large scale ($N \in \{50, 100\}$) makes hierarchical or highly selective coordination effectively necessary—and, as our results confirm, largely beyond the reach of current LLMs.

Communication Protocol (P). As shown in Figure 3, we instantiate three protocols: **P2P**—directed messaging where agents explicitly address individual recipients; **BP**—broadcast messaging where each transmission reaches all agents simultaneously; **SFS**—indirect coordination through a shared file system. Agents retain complete autonomy in deciding what to share, with whom, and when. Detailed specifications are provided in Appendix A.

Language Model (\mathcal{M}). All N agents within a configuration share the same model, isolating coordination capability from heterogeneity effects. We evaluate three frontier open-source models: **DeepSeek-V3.1** (DeepSeek-AI, 2024), **GPT-OSS-120B** (OpenAI, 2025), and **Qwen3-Next-80B-A3B** (Team, 2025).

Our Experimental Setup

Tasks: 30 (10 per difficulty level)
 Scales: 6 (2, 5, 10, 20, 50, 100)
 Protocols: 3 (P2P, BP, SFS)
 Models: 3 (DeepSeek, GPT, Qwen)

This yields $6 \times 3 \times 3 = 54$ unique configurations and $30 \times 54 = 1,620$ total experiments (see Appendix C for infrastructure details). To disentangle coordination overhead from intrinsic task difficulty, we additionally conduct $N=1$ **baseline** experiments where a single agent receives the complete global input and answers directly without communication. We define **Relative Coordination Cost (RCC)** $= 1 - \text{SR}(N=k)/\text{SR}(N=1)$, capturing the fraction of single-agent performance lost to coordination overhead. The $N=1$ oracle represents the upper bound; SILO-BENCH asks whether distributed agents can approach this bound through coordination alone.

4.2 Overall Performance

Table 1 summarizes performance across all models and configurations. DeepSeek-V3.1 achieves a 36.9% average success rate, followed by GPT-OSS-120B at 16.9% and Qwen3 at 8.2%—a $4.5\times$ spread. Even the strongest model fails nearly two-thirds of the time, establishing that current LLMs cannot reliably coordinate under information silos.

Coordination overhead, not task difficulty, drives the performance gap. To confirm that failures stem from coordination rather than intrinsic task hardness, we compare multi-agent success rates against the $N=1$ oracle. Table 2 reports results for GPT-OSS-120B (trends are consistent across models). Even at the smallest team size ($k=2$), multi-agent systems already lose 15–49% of single-agent performance, and RCC compounds steadily with scale, reaching 80–100% at $k=50$ for Level-II and Level-III tasks. Crucially, the single-agent success rate difference between Level-I and Level-III is modest with only about 15 percentage points, yet the multi-agent gap balloons to over 18 percentage points, confirming that performance collapse is driven by coordination failure, not by the tasks themselves being harder.

Agents gather information but fail to integrate it. While RCC reveals *that* coordination fails, the Partial Correctness Score (PCS) reveals *where*. PCS measures continuous answer quality (Section 3.3), and the divergence between PCS and SR isolates the reasoning-integration stage as the bottleneck. At $N \geq 50$ on Level-III tasks, SR drops to 0% while PCS remains at 8–16%, confirming that agents acquire partial global information but cannot synthesize it correctly. This dissociation appears even on simpler tasks: averaged across all scales on Level-I

tasks, DeepSeek-V3.1 achieves a PCS of 88.0% yet an SR of only 62.0% (Table 1), a gap of 26 percentage points indicating that agents collectively hold nearly all required information but still fail to produce a correct final answer.

Performance degrades multiplicatively with scale and complexity. Figure 5 and Table 3 show that task complexity and agent scale interact multiplicatively. DeepSeek-V3.1 drops from 62% on Level-I to 12% on Level-III, and Level-III tasks reach *zero* success at $N \geq 50$, while Level-I tasks remain above 40% even at 100 agents. As Figure 4 illustrates, all models degrade with agent count, and communication density *decreases* at larger scales and agents become sparser in interaction precisely when denser coordination is most needed.

4.3 Protocol Suitability

Having established that coordination fails broadly, we examine whether protocol choice modulates this failure. Figure 6 reveals distinct model-protocol affinities. DeepSeek-V3.1 prefers broadcast messaging (40% with BP vs. 32% with SFS), while GPT-OSS-120B performs best with targeted communication (20% under P2P vs. 14% under BP), suggesting that protocol suitability depends on how a model balances the cognitive cost of addressing decisions against the noise of undifferentiated broadcasts. SFS underperforms in most cases: despite comparable information transfer volume to BP, it consistently yields lower SR, indicating that the bottleneck lies in reasoning about shared state rather than in communication volume.

5 Analysis and Discussion

The preceding results establish that coordination broadly fails, that failures scale with complexity, and that agents accumulate partial information they cannot synthesize. We now investigate the mechanisms: first asking whether agents at least discover the *right structural approach*, then tracing exactly where execution breaks down.

5.1 Case Study: Emergent Coordination Patterns

Figure 7 visualizes the communication patterns that agents spontaneously adopt for three representative tasks. In the Global Max heatmap (Level-I, left), nearly all message traffic flows into column 0: Agent 0 emerges organically as a central aggregator, producing a near-perfect star topology. This

Dimension	DeepSeek-V3.1				GPT-OSS-120B				Qwen3-Next-80B-A3B			
	SR↑	PCS↑	Token↓	Den.	SR↑	PCS↑	Token↓	Den.	SR↑	PCS↑	Token↓	Den.
<i>By Communication Protocol</i>												
BP	40.4	50.7	297.3	0.93	14.4	33.1	148.7	0.78	4.5	12.2	846.9	0.13
P2P	38.9	50.4	363.1	1.52	20.3	42.6	579.8	2.25	10.5	26.5	909.8	0.62
SFS	31.5	41.0	308.5	1.13	16.0	34.7	212.8	0.74	9.5	18.7	864.4	0.21
<i>By Difficulty Level</i>												
Level I	62.0	88.0	184.0	0.71	27.4	56.8	187.1	1.05	20.7	44.4	747.1	0.62
Level II	35.1	59.7	355.9	0.98	14.5	35.3	330.1	1.18	2.9	11.9	990.1	0.19
Level III	11.7	27.9	439.2	1.93	8.8	22.7	424.2	1.54	1.0	1.5	881.6	0.16
<i>By Agent Scale</i>												
N = 2	61.2	78.4	12.1	2.80	34.4	52.0	6.3	2.82	17.2	23.7	41.2	0.54
N = 5	48.5	68.2	44.2	1.94	28.0	47.9	30.6	1.78	9.1	15.7	112.4	0.42
N = 10	39.9	59.1	91.3	1.19	14.0	36.8	77.6	1.35	8.6	19.0	261.4	0.38
N = 20	33.6	60.2	211.0	0.72	13.2	37.0	194.1	0.88	7.4	20.1	549.6	0.35
N = 50	19.0	46.8	510.3	0.25	5.2	27.3	549.8	0.46	5.1	9.1	1466.3	0.14
N = 100	18.1	46.5	1093.8	0.14	6.4	28.7	1024.2	0.24	1.3	5.1	2901.1	0.08
Average	36.9	47.1	323.0	0.82	16.9	38.3	313.8	1.01	8.2	19.8	873.6	0.25

Table 1: Overall performance summary across all models and configurations. SR = Success Rate (%), PCS = Partial Correctness Score (%), Token = Token Consumption (tokens/round), Den. = Communication Density. Best results per section in **bold**.

Scale k	Level I (Aggregation)			Level II (Mesh)			Level III (Global Shuffle)		
	SR($N=1$)	SR($N=k$)	RCC	SR($N=1$)	SR($N=k$)	RCC	SR($N=1$)	SR($N=k$)	RCC
$k = 2$	96.7	82.0	15.2%	90.0	62.0	31.1%	80.0	41.0	48.8%
$k = 5$	93.3	65.0	30.3%	70.0	47.0	32.9%	73.3	22.0	70.0%
$k = 10$	76.7	51.3	33.1%	73.3	22.0	70.0%	60.0	9.0	85.0%
$k = 20$	63.3	48.0	24.2%	36.7	14.0	61.8%	43.3	7.0	83.8%
$k = 50$	33.3	18.0	45.9%	30.0	6.0	80.0%	26.7	0.0	100%
$k = 100$	20.0	10.0	50.0%	13.3	5.0	62.4%	10.0	0.0	100%

Table 2: Single-agent baseline SR (%), multi-agent SR (%), and Relative Coordination Cost (RCC = $1 - \text{SR}(N=k)/\text{SR}(N=1)$) for GPT-OSS-120B across difficulty levels and scales. RCC columns (shaded) quantify the fraction of single-agent performance lost to coordination overhead. Trends are consistent across all three models; full results in Appendix F.

Level	N=2	N=5	N=10	N=20	N=50	N=100
I	85.0	72.0	68.7	65.7	38.1	40.6
II	61.7	55.3	28.3	29.5	17.4	14.3
III	36.2	17.2	10.0	5.7	0.0	0.0
Avg	61.2	48.5	39.9	33.6	19.0	18.1

Table 3: Success Rate (%) by agent count and difficulty level for DeepSeek-V3.1 (averaged across all protocols).

self-organized structure closely matches the theoretically optimal pattern and yields high task success. In the Prefix Sum heatmap (Level-II, center), a prominent diagonal band reflects agents communicating primarily with their immediate neighbors, correctly capturing the sequential dependency of the prefix computation. However, off-diagonal scat-

ter reveals that agents also broadcast beyond their neighbors, generating redundant overhead rather than the clean chain the task requires. In the Distributed Sort heatmap (Level-III, right), the matrix is uniformly dense: agents exchange messages with nearly every other agent, which is precisely what global data reorganization demands, yet the high density comes with highly uneven per-agent loads, and some senders dominate entire rows, suggesting uncoordinated flooding rather than structured exchange.

Taken together, these patterns confirm that agents can translate high-level task descriptions into broadly appropriate coordination topologies without explicit instruction. Yet the heatmaps also reveal a consistent gap between structural intent




Figure 4: Scaling behavior across agent counts. (a) Success rates decline for all models as team size increases, with sharp drops beyond $N = 20$. (b) Token consumption scales roughly linearly with agent count. (c) Communication density decreases at scale, suggesting coordination sparsification.




Figure 5: Success rate by difficulty level.




Figure 6: Success rate by communication protocol.

and execution quality: even when the right topology emerges, agents over-communicate, distribute load unevenly, or fail to adhere to the optimal pattern. This raises the core question addressed next: given that agents communicate in approximately the right way, why do they still so often fail?

5.2 The Communication-Reasoning Gap

To classify failure systematically, we apply a two-stage hybrid procedure: rule-based detection identifies *Premature Submission* (agent submits before reaching the task-specific minimum peer count), *Consensus Failure* ($|\{\hat{y}_i\}_{i=1}^N| > 1$), and *Computation Error* (full data receipt confirmed in log but answer incorrect). Two independent annotators then reviewed stratified runs, achieving Cohen’s $\kappa = 0.87$; disagreements were resolved by discussion.

Analyzing execution logs under this scheme, we identified three distinct failure modes (Table 4). *Premature Submission* (37.2%) is the most preva-

lent: agents submit before gathering sufficient information. Agent-77 in Task I-06, for instance, submitted after contacting only 28 of 100 agents, yielding answer 208 vs. the expected 114. *Consensus Failure* (29.9%) occurs when agents communicate actively but cannot converge; one 100-agent XOR checksum run produced 12 distinct answers, with 86 agents converging on 146 while the remainder submitted values ranging from 42 to 238. *Computation Error* (28.6%) arises when agents collect all required data yet compute incorrectly, such as submitting 619 instead of 620 due to an off-by-one error during final aggregation. These modes frequently co-occur: 67 runs exhibit both Premature Submission and Consensus Failure, as early exits prevent subsequent consensus-building and widen the convergence gap for remaining agents.

Together, these three modes define the *Communication-Reasoning Gap*: agents exhibit proficiency in the social mechanics of coordination like formatting messages, responding to peers, organizing information flow, while failing at the computational core of determining information sufficiency and synthesizing distributed state. This is not a failure of effort: behavioral comparison shows successful runs complete in fewer rounds, while verification behaviors appear in over 95% of runs regardless of outcome. The bottleneck is reasoning quality at the integration stage, not communication intent.

5.3 Implications and Future Directions

The analyses jointly reveal a structural asymmetry with practical consequences. Coordination overhead does not merely reduce parallelization gains. For Level-III tasks at $N \geq 50$, it eliminates them entirely, leaving a coordinated team outperformed by




Figure 7: Communication heatmaps for representative tasks with $N = 20$ agents using DeepSeek-V3.1. **Left:** I-01 (Global Max) shows emergent leader pattern. **Center:** II-11 (Prefix Sum) exhibits diagonal pattern indicating partial spatial locality discovery. **Right:** III-21 (Distributed Sort) shows dense all-to-all communication.

Failure Mode	Count	Percent
Success	153	50.8%
Premature Submission	112	37.2%
Consensus Failure	90	29.9%
Computation Error	86	28.6%

Table 4: Failure mode distribution (categories not mutually exclusive).

a single agent with full data access. Perhaps most counterintuitively, spontaneous leader emergence which is conventionally assumed to help, actively hurts performance on Level-III tasks, because the aggregator becomes overwhelmed by the volume of global data it must process.

Three directions follow. First, agents need mechanisms to detect information sufficiency before committing to a final answer. Second, the explicit synchronization checkpoints present in successful runs should be formalized as consensus protocols. Third, adaptive protocol selection based on task structure could unlock model-protocol co-optimization, given the model-dependent affinities observed. SILO-BENCH provides the evaluation foundation for tracking progress along all three.

6 Conclusion

We introduce SILO-BENCH to evaluate distributed coordination in multi-agent LLM systems across 1,620 experiments. The results are unambiguous: current LLMs cannot reliably escape their information silos through coordination alone.

The *Communication-Reasoning Gap* identifies the precise fault line: agents are competent communicators but poor distributed reasoners. They spontaneously form task-appropriate topologies and exchange information actively, yet consistently fail

to integrate what they have gathered—a dissociation made concrete by the PCS–SR divergence and the RCC analysis showing that coordination overhead eliminates parallelization gains entirely at high complexity. Most strikingly, spontaneous leader emergence actively hurts performance on complex tasks, revealing that self-organized centralization creates bottlenecks rather than resolving them.

Closing this gap will require mechanisms for information sufficiency detection, explicit consensus protocols, and adaptive coordination strategies. SILO-BENCH provides the evaluation infrastructure to track progress along these directions.

Limitations

While SILO-BENCH provides a comprehensive framework for evaluating multi-agent collaboration, it has several limitations. Our evaluation only covers three fundamental communication protocols and does not include other coordination mechanisms such as hierarchical protocols, gossip-based dissemination and hybrid approaches due to limited budget. We adopt agent configurations with uniform underlying models, whereas real-world multi-agent systems usually involve heterogeneous compositions with distinct coordination patterns. Closed-source models are not evaluated in this work due to their high cost at our scale and unverifiable, incomparable reported token usage. In addition, our assessment focuses on three frontier LLMs, which may not capture the full spectrum of failure modes across all LLMs since each model has unique characteristics in reasoning logic, communication strategies and error propagation that lead to distinct performance limitations.

References

- Chenxin An, Shansan Gong, Ming Zhong, Xingjian Zhao, Mukai Li, Jun Zhang, Lingpeng Kong, and Xipeng Qiu. 2024. [L-eval: Instituting standardized evaluation for long context language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 14388–14411. Association for Computational Linguistics.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024. [Longbench: A bilingual, multi-task benchmark for long context understanding](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 3119–3137. Association for Computational Linguistics.
- Razan Baltaji, Babak Hemmatian, and Lav R. Varshney. 2024. [Conformity, confabulation, and impersonation: Persona inconstancy in multi-agent LLM collaboration](#). *CoRR*, abs/2405.03862.
- Junzhe Chen, Xuming Hu, Shuodi Liu, Shiyu Huang, Wei-Wei Tu, Zhaofeng He, and Lijie Wen. 2024a. [Llmarena: Assessing capabilities of large language models in dynamic multi-agent environments](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 13055–13077. Association for Computational Linguistics.
- Longze Chen, Ziqiang Liu, Wanwei He, Yinhe Zheng, Hao Sun, Yunshui Li, Run Luo, and Min Yang. 2024b. [Long context is not long at all: A prospector of long-dependency data for large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 8222–8234. Association for Computational Linguistics.
- Jeffrey Dean and Sanjay Ghemawat. 2008. [Mapreduce: simplified data processing on large clusters](#). *Commun. ACM*, 51(1):107–113.
- DeepSeek-AI. 2024. [Deepseek-v3 technical report](#). *CoRR*, abs/2412.19437.
- Shihan Deng, Weikai Xu, Hongda Sun, Wei Liu, Tao Tan, Jianfeng Liu, Ang Li, Jian Luan, Bin Wang, Rui Yan, and Shuo Shang. 2024. [Mobile-bench: An evaluation benchmark for llm-based mobile agents](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 8813–8831. Association for Computational Linguistics.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2024. [Improving factuality and reasoning in language models through multiagent debate](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, Proceedings of Machine Learning Research, pages 11733–11763. PMLR / OpenReview.net.
- Luca Gioacchini, Giuseppe Siracusano, Davide Sanvito, Kiril Gashteovski, David Friede, Roberto Bifulco, and Carolin Lawrence. 2024. [Agentquest: A modular benchmark framework to measure progress and improve LLM agents](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: System Demonstrations, NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pages 185–193. Association for Computational Linguistics.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. 2024. [Metagpt: Meta programming for A multi-agent collaborative framework](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Md. Ashraful Islam, Mohammed Eunus Ali, and Md. Rizwan Parvez. 2024a. [Mapcoder: Multi-agent code generation for competitive problem solving](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 4912–4944. Association for Computational Linguistics.
- Shayekh Bin Islam, Md. Asib Rahman, K. S. M. Tozammel Hossain, Enamul Hoque, Shafiq Joty, and Md. Rizwan Parvez. 2024b. [Open-rag: Enhanced retrieval augmented reasoning with open-source large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024, Miami, Florida, USA, November 12-16, 2024*, Findings of ACL, pages 14231–14244. Association for Computational Linguistics.
- Zimo Ji, Daoyuan Wu, Wenyuan Jiang, Pingchuan Ma, Zongjie Li, and Shuai Wang. 2025. [Measuring and augmenting large language models for solving capture-the-flag challenges](#). In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security, CCS 2025, Taipei, Taiwan, October 13-17, 2025*, pages 603–617. ACM.
- Yihuai Lan, Zhiqiang Hu, Lei Wang, Yang Wang, Deheng Ye, Peilin Zhao, Ee-Peng Lim, Hui Xiong, and Hao Wang. 2024. [Llm-based agent society investigation: Collaboration and confrontation in avalon gameplay](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16,*

- 2024, pages 128–145. Association for Computational Linguistics.
- Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. [CAMEL: communicative agents for "mind" exploration of large language model society](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Jiaqi Li, Mengmeng Wang, Zilong Zheng, and Muhan Zhang. 2024. [Loogle: Can long-context language models understand long contexts?](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 16304–16333. Association for Computational Linguistics.
- Hao Liu, Wilson Yan, Matei Zaharia, and Pieter Abbeel. 2025. [World model on million-length video and language with blockwise ringattention](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024a. [Lost in the middle: How language models use long contexts](#). *Trans. Assoc. Comput. Linguistics*, 12:157–173.
- Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, and 3 others. 2024b. [Agentbench: Evaluating llms as agents](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Shichao Ma, Yunhe Guo, Jiahao Su, Qihe Huang, Zhengyang Zhou, and Yang Wang. 2026. [Talk2image: A multi-agent system for multi-turn image generation and editing](#). In *Fortieth AAAI Conference on Artificial Intelligence, Thirty-Eighth Conference on Innovative Applications of Artificial Intelligence, Sixteenth Symposium on Educational Advances in Artificial Intelligence, AAAI 2026, Singapore, January 20-27, 2026*, pages 32437–32445. AAAI Press.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- OpenAI. 2025. [gpt-oss-120b & gpt-oss-20b model card](#). *CoRR*, abs/2508.10925.
- Joon Sung Park, Joseph C. O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. [Generative agents: Interactive simula-cra of human behavior](#). In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, UIST 2023, San Francisco, CA, USA, 29 October 2023- 1 November 2023*, pages 2:1–2:22. ACM.
- Ofir Press, Noah A. Smith, and Mike Lewis. 2022. [Train short, test long: Attention with linear biases enables input length extrapolation](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024a. [Chatdev: Communicative agents for software development](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 15174–15186. Association for Computational Linguistics.
- Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024b. [Chatdev: Communicative agents for software development](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 15174–15186. Association for Computational Linguistics.
- Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Karpas, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. [Parallel context windows for large language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 6383–6402. Association for Computational Linguistics.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy P. Lillicrap, Jean-Baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, Ioannis Antonoglou, Rohan Anil, Sebastian Borgeaud, Andrew M. Dai, Katie Millican, Ethan Dyer, Mia Glaese, Thibault Sottiaux, Benjamin Lee, and 34 others. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *CoRR*, abs/2403.05530.
- Uri Shaham, Elad Segal, Maor Ivgi, Avia Efrat, Ori Yoran, Adi Haviv, Ankit Gupta, Wenhan Xiong, Mor Geva, Jonathan Berant, and Omer Levy. 2022. [SCROLLS: standardized comparison over long language sequences](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 12007–12021. Association for Computational Linguistics.
- Jinwei Su, Qizhen Lan, Yinghui Xia, Lifan Sun, Weiyu Tian, Tianyu Shi, and Lewei He. 2026. [Difficulty-aware agentic orchestration for query-specific multi-agent workflows](#). In *Proceedings of the ACM Web*

- Conference 2026, WWW 2026, Dubai, United Arab Emirates, originally scheduled for April 13-17, 2026, rescheduled for June 29 - July 3, 2026*, pages 2060–2070. ACM.
- Rui Sun, Jie Ding, Chenghua Gong, Tianjun Gu, Yihang Jiang, Juyuan Zhang, Liming Pan, and Linyuan Lü. 2026. [Topodim: One-shot topology generation of diverse interaction modes for multi-agent systems](#). *CoRR*, abs/2601.10120.
- Qwen Team. 2025. [Qwen3 technical report](#). *CoRR*, abs/2505.09388.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *CoRR*, abs/2302.13971.
- Junlin Wang, Jue Wang, Ben Athiwaratkun, Ce Zhang, and James Zou. 2025a. [Mixture-of-agents enhances large language model capabilities](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Qineng Wang, Zihao Wang, Ying Su, Hanghang Tong, and Yangqiu Song. 2024a. [Rethinking the bounds of LLM reasoning: Are multi-agent discussions the key?](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 6106–6131. Association for Computational Linguistics.
- Wenhao Wang, Peizhi Niu, Zhao Xu, Zhaoyu Chen, Jian Du, Yaxin Du, Xianghe Pang, Keduan Huang, Yanfeng Wang, Qiang Yan, and Siheng Chen. 2025b. [Mcp-flow: Facilitating LLM agents to master real-world, diverse and scaling MCP tools](#). *CoRR*, abs/2510.24284.
- WenHao Wang, Zijie Yu, Rui Ye, Jianqing Zhang, Guangyi Liu, Liang Liu, Siheng Chen, and Yanfeng Wang. 2025c. [Fedmabench: Benchmarking mobile GUI agents on decentralized heterogeneous user data](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, EMNLP 2025, Suzhou, China, November 4-9, 2025*, pages 26387–26408. Association for Computational Linguistics.
- Zheng Wang, Shu Xian Teo, Jieer Ouyang, Yongjun Xu, and Wei Shi. 2024b. [M-RAG: reinforcing large language model performance through retrieval-augmented generation with multiple partitions](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 1966–1978. Association for Computational Linguistics.
- Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. 2023. [Autogen: Enabling next-gen LLM applications via multi-agent conversation framework](#). *CoRR*, abs/2308.08155.
- Xin Yang, Letian Li, Abudukelimu Wuerkaixi, Xuxin Cheng, Cao Liu, Ke Zeng, Xunliang Cai, and Wenyuan Jiang. 2026. [Towards self-robust llms: Intrinsic prompt noise resistance via coipo](#). *CoRR*, abs/2603.03314.
- Andrew Chi-Chih Yao. 1979. [Some complexity questions related to distributive computing \(preliminary report\)](#). In *Proceedings of the 11th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1979, Atlanta, Georgia, USA*, pages 209–213. ACM.
- Howard Yen, Tianyu Gao, and Danqi Chen. 2024. [Long-context language modeling with parallel context encoding](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 2588–2610. Association for Computational Linguistics.
- Jintian Zhang, Xin Xu, Ningyu Zhang, Ruibo Liu, Bryan Hooi, and Shumin Deng. 2024a. [Exploring collaboration mechanisms for LLM agents: A social psychology view](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 14544–14607. Association for Computational Linguistics.
- Xinrong Zhang, Yingfa Chen, Shengding Hu, Zihang Xu, Junhao Chen, Moo Khai Hao, Xu Han, Zhen Leng Thai, Shuo Wang, Zhiyuan Liu, and Maosong Sun. 2024b. [inftybench: Extending long context evaluation beyond 100k tokens](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 15262–15277. Association for Computational Linguistics.
- Xiutian Zhao, Ke Wang, and Wei Peng. 2024. [An electoral approach to diversify llm-based multi-agent collective decision-making](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 2712–2727. Association for Computational Linguistics.

A Communication Protocols

This appendix provides complete specifications for the three communication protocols implemented in SILO-BENCH. Each protocol defines a distinct coordination substrate constraining the *mechanism* of information exchange while preserving full agent autonomy over *content* and *strategy*.

A.1 Protocol Overview

Table 5 summarizes the three protocols. They span the spectrum of coordination paradigms along three axes: (1) *explicit vs. implicit addressing*—P2P requires agents to name recipients, BP eliminates addressing entirely, SFS routes coordination through shared state; (2) *direct vs. indirect communication*—P2P and BP involve direct message exchange, SFS agents never “speak” to each other explicitly; (3) *default density*—P2P encourages sparse targeted exchanges, BP defaults to dense all-to-all dissemination, SFS density depends entirely on read/write behavior.

A.2 Peer-to-Peer Protocol (P2P)

The P2P protocol implements directed messaging through SQLite-backed mailboxes. Each agent maintains a private inbox; messages are delivered asynchronously to the recipient’s buffer until their next activation. Agents must decide not only *what* to communicate but *whom* to contact, enabling evaluation of task-appropriate routing strategy discovery. Available actions are: `send_message(target_id, content)`, which delivers a message to the specified agent; `receive_messages()`, which retrieves all pending messages; `wait()`, which signals completion of the agent’s decision for the current round (see below); and `submit_result(answer)`, which submits the final answer. Messages are stored in an in-memory SQLite database recording sender ID, recipient ID, content, timestamp, and read status. Delivery ordering within each sender-recipient pair is guaranteed; no global ordering is enforced.

A.3 Broadcast Protocol (BP)

The BP protocol implements broadcast messaging where each transmission reaches all other agents simultaneously. An implicit aggregator collects messages each round and distributes the compiled history to all participants. Available actions are: `broadcast_message(content)`, `receive_messages()`, `list_agents()`, `wait()`,

and `submit_result(answer)`. Broadcast messages are stored centrally, tagged with sender ID and timestamp, and delivered as a chronologically ordered compiled view at each round.

A.4 Shared File System Protocol (SFS)

The SFS protocol implements indirect coordination through a shared key-value store visible to all agents. Rather than exchanging messages directly, agents read and write to a common namespace, enabling asynchronous coordination analogous to blackboard architectures. Available actions are: `list_files(prefix)`, `read_file(path)`, `write_file(path, content)`, `delete_file(path)`, `wait()`, and `submit_result(answer)`. The shared file system is backed by an in-memory SQLite database storing path, content, creation time, and last modification time. Writes are immediately visible to subsequent reads; concurrent writes to the same path follow last-writer-wins semantics.

A.5 The `wait()` Action: Formal Specification

The `wait()` action is semantically uniform across all three protocols and serves as an explicit round-boundary signal. When an agent invokes `wait()`, all remaining operations in its current round are skipped, and the agent’s decision phase for round r is marked complete. The agent is then suspended until the start of round $r+1$.

At the start of round $r+1$, the following activation and message-delivery rules apply:

- **P2P and BP:** The agent’s inbox is populated with all messages sent to it (P2P) or broadcast by any agent (BP) during round r . These are delivered atomically at round start; no message sent in round r is visible before round $r+1$.
- **SFS:** The agent observes the full shared file system state as of the end of round r , including all writes committed by other agents during round r .
- **No blocking:** `wait()` does not block on a specific event or agent. It is a “pass” that yields control to the synchronous round scheduler. If an agent never calls `wait()` explicitly, the runtime automatically advances it to the next round after its action budget is exhausted.
- **Post-submit behavior:** Once an agent has called `submit_result()`, subsequent rounds are no-ops for that agent—it neither receives new messages nor is activated again.

Protocol	Description	Available Actions
P2P	Directed messaging via agent-addressed mailboxes	send_message, receive_messages, wait, submit_result
BP	Broadcast messaging to all agents simultaneously	broadcast_message, receive_messages, list_agents, wait, submit_result
SFS	Coordination through shared key-value storage	list_files, read_file, write_file, delete_file, wait, submit_result

Table 5: Comparison of communication protocols in SILO-BENCH.

This synchronous round-based model ensures that all N agents observe a consistent snapshot of the communication state at each round boundary, making execution reproducible and analysis tractable.

B Prompt Design: Structural Guidance without Role Prescription

This appendix details the prompt templates used to initialize agents in SILO-BENCH. Our core design philosophy is to provide *task-structural information* while preserving *strategic autonomy*: prompts convey high-level dependency patterns and potential coordination approaches, but do not prescribe mandatory execution sequences or assign semantic roles.

B.1 Base Prompt Template

Each agent receives an initialization prompt following this structure:

Agent Initialization Prompt

```
You are Agent {agent_id} in a multi-agent
system consisting of {N} agents (IDs range
from 0 to {N-1}).
Task Description:
{task_description}
Your Local Data:
{data_shard}
Communication Protocol:
{protocol_description}
Available Actions:
- {protocol_specific_actions}
- submit_result(answer): Submit your final
answer when confident
Your goal is to coordinate with other agents
to compute the globally correct answer. No
single agent has sufficient information to
solve this task independently. When you
have determined the answer, submit it using
submit_result().
```

The `{protocol_specific_actions}` placeholder is instantiated according to the protocol specifications in Appendix A.

B.2 Task Description Examples

Task descriptions convey task structure and potential coordination patterns while leaving concrete implementation decisions to the agents themselves. The three examples below illustrate how descriptions scale from simple aggregation to global reorganization tasks.

Example: Global Maximum (Task I-01)

Find the maximum value across all data distributed among the agents. Each agent holds a portion of a larger array. The correct answer is the single largest integer across the entire distributed dataset. You must coordinate with other agents to determine this global maximum.

Example: Prefix Sum (Task II-11)

Compute the prefix sum array for a sequence distributed across agents. Agent 0 holds elements $[0, k)$, Agent 1 holds elements $[k, 2k)$, and so on. The prefix sum at position i is the sum of all elements from position 0 to i . You must coordinate to compute the correct prefix sum for your portion, accounting for cumulative sums from preceding agents.

Example: Distributed Sort (Task III-21)

Sort the entire distributed array in ascending order. Each agent holds a portion of the unsorted data. The final result should be the complete sorted sequence. You must coordinate to exchange data and determine the correct global ordering.

B.3 Design Principles

Our prompts are carefully calibrated to provide structural information without prescribing behavior. On the one hand, descriptions convey whether tasks involve aggregation, sequential dependencies, or global data reorganization (e.g., “accounting for cumulative sums from preceding agents”); some prompts suggest possible coordination patterns as

options rather than requirements (e.g., “consider establishing a coordinator” or “you may exchange with neighbors”); and for complex tasks, prompts may mention general algorithmic paradigms (e.g., “sample-based partitioning”) without specifying concrete steps. On the other hand, we do not assign semantic roles—no agent is designated “Manager,” “Worker,” or “Coordinator”—and prompts never specify “Step 1: Agent 0 does X, Step 2: Agent 1 does Y.” Agents decide their own message timing and recipients, and must discover consensus and verification mechanisms independently. The distinction is illustrated below:

- ✗ **Prescribed (NOT our approach):** “You are the leader. Step 1: Collect data from all agents. Step 2: Compute the result. Step 3: Broadcast the answer.”
- ✓ **Our approach:** “Can you identify a leader to collect and compare results? How would agents coordinate to reach consensus?”

All agents receive structurally identical prompts (modulo their ID and data shard), ensuring no agent holds implicit leadership status. The phrase “No single agent has sufficient information” is included explicitly to prevent premature submission of partial results. This design tests whether agents can translate high-level task understanding into concrete coordination protocols—a capability that, as our results demonstrate, remains largely absent in current LLMs.

C Experimental Details

C.1 Agent Scale Rationale

The six agent counts are chosen to probe qualitatively distinct coordination regimes. The minimal team ($N = 2$) isolates fundamental pairwise coordination without overhead. Small groups ($N \in \{5, 10\}$) allow agents to feasibly track all peers simultaneously. Medium scale ($N = 20$) begins to make exhaustive peer tracking challenging, pushing agents toward selective communication. Large scale ($N \in \{50, 100\}$) makes hierarchical or highly selective coordination effectively necessary—and, as our results confirm, largely beyond the reach of current LLMs.

C.2 Execution Parameters

Each configuration is allocated a maximum of $R_{\max} = 100$ communication rounds. Within each

round, all agents are activated in parallel: they receive incoming messages from the previous round, independently decide on actions, and execute them simultaneously. Messages or files written in round r become visible to all agents at the start of round $r + 1$. An agent exits the coordination loop upon invoking `submit_result(answer)`; agents that fail to submit within R_{\max} rounds are assigned a null answer counted as incorrect. Due to computational constraints, each configuration is executed once with fixed random seeds for data generation.

C.3 Infrastructure

Experiments were conducted on a GH200 cluster, with up to 50 concurrent configurations executed simultaneously. The compute amounted to approximately 500+ GPU-hours equivalent for each experiment. Complete conversation histories, token counts, and timing information were recorded for all runs.

C.4 Model Licenses and Intended Use

All language models used in this study are open-source and deployed locally on our infrastructure. DeepSeek-V3.1 is released under the MIT License¹, GPT-OSS-120B under the Apache 2.0 License², and Qwen3-Next-80B-A3B under the Apache 2.0 License³, all permitting research and commercial use.

D Token Budget Feasibility Analysis

To verify that SILO-BENCH operates within practical token limits, we profile token consumption across all 54 configurations, decomposing total usage into three components: the base initialization prompt (T_{base}), the local data shard (T_{data}), and accumulated communication messages (T_{comm}). Table 6 summarizes the fixed components, and Table 7 reports model-dependent communication costs.

Context window utilization, shown in Table 8, remains low on average: DeepSeek-V3.1 uses 7.5% of the 128K budget on average, while GPT-OSS-120B and Qwen3 stay below 3%. The 95th percentile cases—driven by redundant broadcasting, failed convergence with extended verbose rounds,

¹<https://huggingface.co/deepseek-ai/DeepSeek-V3.1>

²<https://huggingface.co/openai/gpt-oss-120b>

³<https://huggingface.co/Qwen/Qwen3-Next-80B-A3B-Instruct>

Component	Min	Mean	Max	Std
Base Prompt (T_{base})	612	748.8	989	87.3
Data Shard (T_{data})	45	312.4	1,856	298.6

Table 6: Token consumption of base prompt and data shard per agent (across all configurations).

Model	Min	Mean	Max	Std
DeepSeek-V3.1	124	8,498.7	98,432	12,847
GPT-OSS-120B	89	2,049.3	45,218	5,632
Qwen3-Next-80B-A3B	156	2,299.3	52,847	6,891

Table 7: Communication (T_{comm}) token consumption per agent by model (across all configurations).

or agents copy-pasting full message histories—are precisely the coordination inefficiencies SILO-BENCH is designed to expose. Overall, frontier models (128K–200K context) can run all configurations comfortably; mid-tier models (32K) handle over 90% of configurations, with Level-III at $N \geq 50$ potentially requiring truncation; and smaller models (8K) are suitable for $N \leq 10$ and Level I–II tasks.

E Complete Task Specifications

Table 9 provides the complete mapping between SILO-BENCH tasks and their algorithmic foundations, including the distributed adaptation approach for each.

F Detailed Results

F.1 Task-Level Breakdown for DeepSeek-V3.1

Table 10 provides a comprehensive breakdown of DeepSeek-V3.1’s success rate across all 30 tasks and three communication protocols. Tasks achieving $\geq 50\%$ success rate under any protocol are highlighted in green. Level-I aggregation tasks cluster at the top, with Distributed Vote (I-03) and Any Match (I-04) achieving near-perfect performance across all protocols. Performance degrades sharply for Level-III tasks, with K-Means Iteration (III-25), Collaborative Filtering (III-27), PageRank Step (III-28), and Matrix Multiply (III-30) achieving zero success across all protocols—these tasks require precise numerical computation over all data shards simultaneously, which proves beyond current distributed LLM capabilities.

Model	Mean Util.	95th Pctl.	Max Util.
DeepSeek-V3.1	7.5%	28.4%	76.9%
GPT-OSS-120B	2.4%	8.2%	35.3%
Qwen3-Next-80B-A3B	2.6%	9.1%	41.3%

Table 8: Context window utilization (%) for 128K context models.

F.2 Results by Model, Protocol, and Difficulty

Table 11 provides success rates for all model-protocol-difficulty combinations, and Table 12 reports success rates by agent count across models. Together, they confirm that the patterns observed for DeepSeek-V3.1 are consistent across all three models: P2P generally outperforms or matches BP for GPT-OSS-120B and Qwen3, SFS consistently underperforms, and performance degrades monotonically with both complexity level and agent count.

F.3 Communication Density Analysis

Table 13 reports communication density across configurations. A consistent pattern emerges across all models: P2P yields substantially higher densities than BP, reflecting agents’ tendency to send multiple targeted messages per pair across rounds; BP densities cluster below 1.0, consistent with one-to-all single broadcasts; and SFS yields notably lower densities than P2P across all models and difficulty levels, indicating that file-based coordination generates sparser cross-agent information flow under our operational definition (read-based transfer counting)—which further explains SFS’s systematic under performance on SR despite non-trivial write activity.

G Failure Mode Analysis

G.1 Representative Failure Cases

Table 14 presents representative examples of each failure mode extracted from DeepSeek-V3.1 execution logs. The cases illustrate how failures manifest in practice: premature submission occurs even after reasonable communication volume (Case 4: only 28 of 100 peers contacted before submitting); consensus failure can persist despite near-unanimous agreement, with a single outlier agent preventing full success (Cases 1 and 3); and computation error strikes even when agents have gathered the complete required data (Case 2: off-by-one arithmetic during aggregation).

ID	Task Name	Reference	Distributed Adaptation
<i>Level I: Aggregation (Optimal: Star/Tree Topology, $\mathcal{O}(N)$ messages)</i>			
I-01	Global Maximum	LC-414	Array partitioned; local max \rightarrow global aggregation
I-02	Word Frequency	LC-2085	Word lists distributed; count target word globally
I-03	Distributed Vote	LC-169	Vote records partitioned; aggregate to find majority
I-04	Any Match	LC-28	String collection split; detect pattern in any shard
I-05	Range Count	LC-327	Count elements in range across shards
I-06	Checksum (XOR)	LC-136	Data blocks distributed; compute global XOR
I-07	Average Value	LC-1491	Array partitioned; combine local sums and counts
I-08	Set Union Size	LC-217	Elements distributed; compute $ \bigcup_i \mathcal{D}_i $
I-09	Top-K Selection	LC-215	Array partitioned; merge local top-K candidates
I-10	Standard Deviation	—	Two-phase: global mean \rightarrow global variance
<i>Level II: Mesh Network (Optimal: Chain Topology, $\mathcal{O}(N)$ messages)</i>			
II-11	Prefix Sum	LC-1480	Sequential dependency; cumulative offset propagation
II-12	Moving Average	LC-346	Sliding window spans boundaries; neighbor exchange
II-13	Longest Palindrome	LC-5	String partitioned; palindromes may cross boundaries
II-14	1D Life Game	LC-289	Cellular automaton; boundary cells need neighbor states
II-15	Pattern Search	LC-392	Subsequence matching across partitioned sequence
II-16	Trapping Rain	LC-42	Global max-left/max-right propagation required
II-17	Diff Array	LC-1094	Difference array with boundary handling
II-18	List Ranking	LC-542	Linked list ranking requires predecessor chain
II-19	Merge Neighbors	—	Boundary element merging between adjacent agents
II-20	Pipeline Hash	—	Sequential hash with chained dependencies
<i>Level III: Shuffling (Optimal: Varies, $\mathcal{O}(N \log N)$ to $\mathcal{O}(N^2)$ messages)</i>			
III-21	Distributed Sort	LC-912	Sample-sort or merge-sort across partitions
III-22	Median of Medians	LC-295	Iterative median selection across distributed data
III-23	Graph Components	LC-323	Edges distributed; iterative union-find
III-24	BFS Distance	LC-542	Graph BFS with distributed edge list
III-25	K-Means Iteration	LC-296	One K-means iteration with distributed points
III-26	Global Distinct	LC-349	Hash-based global deduplication
III-27	Collab. Filtering	LC-1	User-item matching with distributed vectors
III-28	PageRank Step	LC-207	One PageRank iteration with distributed edges
III-29	Load Balance	LC-410	Task redistribution to minimize load variance
III-30	Matrix Multiply	LC-311	Row/column partitioned matrix multiplication

Table 9: Complete specification of SILO-BENCH tasks with LeetCode references and distributed adaptation descriptions.

G.2 Failure Mode Definitions and Co-occurrence

We formally define the three failure modes as follows. *Premature Submission* occurs when an agent invokes `submit_result()` before receiving information from a sufficient subset of peers—where “sufficient” means the minimum number of agents whose data is required to compute the correct answer. *Consensus Failure* occurs when agents submit multiple distinct answers ($|\{\hat{y}_i\}_{i=1}^N| > 1$), indicating that coordination failed to synchronize agents’ understanding of global state. *Computation Error* occurs when an agent receives sufficient information but submits an incorrect answer, iso-

lating failures in the reasoning phase from those in the communication phase.

These modes frequently co-occur within single runs, as shown in Table 15. The high co-occurrence between premature submission and consensus failure (67 cases) suggests a cascading effect: agents submitting early cannot participate in subsequent consensus-building, leaving remaining agents with incomplete information and widening the convergence gap.

Task	BP	P2P	SFS	Avg
<i>Level I: Aggregation</i>				
I-01 Global Max	100	100	80	93.3
I-02 Word Frequency	100	52	70	73.9
I-03 Distributed Vote	100	100	100	100.0
I-04 Any Match	100	100	99	99.6
I-05 Range Count	83	67	63	71.0
I-06 Checksum (XOR)	17	0	2	6.1
I-07 Average Value	99	83	46	76.2
I-08 Set Union Size	50	50	42	47.5
I-09 Top-K Select	36	67	21	41.2
I-10 Standard Deviation	17	17	17	16.7
<i>Level II: Structured Mesh</i>				
II-11 Prefix Sum	84	80	43	68.8
II-12 Moving Average	0	0	0	0.0
II-13 Longest Palindrome	49	66	48	54.5
II-14 1D Life Game	27	47	24	32.4
II-15 Pattern Search	0	17	17	11.1
II-16 Trapping Rain	33	33	40	35.6
II-17 Diff Array	48	62	20	43.3
II-18 List Ranking	0	0	0	0.0
II-19 Merge Neighbors	59	60	66	61.8
II-20 Pipeline Hash	52	36	55	47.6
<i>Level III: Global Shuffle</i>				
III-21 Distributed Sort	33	20	20	24.4
III-22 Median of Medians	20	20	0	13.3
III-23 Graph Components	40	40	14	31.3
III-24 BFS Distance	0	0	0	0.0
III-25 K-Means Iteration	0	0	0	0.0
III-26 Global Distinct	33	33	0	22.2
III-27 Collab. Filtering	0	0	0	0.0
III-28 PageRank Step	0	0	0	0.0
III-29 Load Balance	32	3	63	32.9
III-30 Matrix Multiply	0	0	0	0.0

Table 10: Success Rate (%) by task and communication protocol for DeepSeek-V3.1. Tasks with $\geq 50\%$ success are highlighted in gray background.

G.3 Behavioral Patterns and Leader Emergence

Table 16 compares behavioral metrics between successful and failed runs. Successful runs complete in notably fewer rounds (8.3 vs. 12.7), suggesting that effective coordination converges quickly while failed runs engage in extended but ultimately unproductive communication loops. Verification behaviors appear in over 95% of runs regardless of outcome, confirming that the bottleneck is not communication intent but reasoning quality.

We also examined whether spontaneous leader emergence correlates with task success, classifying an agent as an emergent leader if it receives more than $1.5\times$ the average number of messages. The results in Table 17 are counterintuitive: leader emergence does not consistently improve outcomes, and for Level-III tasks, runs *with* an emergent leader

Model	Protocol	L-I	L-II	L-III
DeepSeek-V3.1	BP	69.7	34.5	14.7
	P2P	62.9	39.9	11.5
	SFS	53.5	30.6	9.0
GPT-OSS-120B	BP	19.6	13.9	9.7
	P2P	34.9	18.5	7.5
	SFS	27.7	11.0	9.1
Qwen3-Next-80B-A3B	BP	9.0	3.2	1.1
	P2P	27.6	3.4	0.3
	SFS	25.5	2.1	1.7

Table 11: Success Rate (%) by model, protocol, and difficulty level.

Model	N=2	N=5	N=10	N=20	N=50	N=100
DeepSeek-V3.1	61.2	48.5	39.9	33.6	19.0	18.1
GPT-OSS-120B	34.4	28.0	14.0	13.2	5.2	6.4
Qwen3-Next-80B-A3B	17.2	9.1	8.6	7.4	5.1	1.3

Table 12: Success Rate (%) by model and agent count.

achieve 0% success versus 33.3% without one. This suggests that spontaneous centralization at high complexity creates coordination bottlenecks—the designated aggregator becomes overwhelmed by the volume of global data—rather than resolving them.

G.4 Successful Coordination Examples

To contrast with the failure modes above, we document two illustrative successful patterns. In Case S-1 (Task I-07, $N=5$), Agent-0 emerged as coordinator organically: all agents broadcast local results to Agent-0, which computed and rebroadcast the global answer. All agents verified and submitted identically within 4 rounds (100% success). In Case S-2 (Task I-01, $N=10$), agents adopted a distributed verification strategy, with each agent confirming understanding with two neighbors before submission. This redundant verification eliminated consensus failures despite higher message overhead. Both successful patterns share a key property: explicit synchronization checkpoints where agents confirm mutual understanding before proceeding—a discipline entirely absent in failed runs.

H Prompt Scaffold Ablation

To assess the sensitivity of our results to prompt design, we ran a controlled ablation under DeepSeek-V3.1 + P2P with three scaffolding conditions beyond the standard neutral prompt: (a) **planning round**—a dedicated strategy-discussion round be-

Model	Protocol	L-I	L-II	L-III
DeepSeek-V3.1	BP	0.56	0.82	1.46
	P2P	0.76	1.06	2.83
	SFS	0.82	1.08	1.51
GPT-OSS-120B	BP	0.54	0.84	0.94
	P2P	1.86	2.15	2.73
	SFS	0.73	0.54	0.95
Qwen3-Next-80B-A3B	BP	0.23	0.10	0.06
	P2P	1.15	0.38	0.33
	SFS	0.48	0.07	0.09

Table 13: Communication Density by model, protocol, and difficulty level.

fore data exchange begins; (b) **protocol reminder**—a brief restatement of available communication actions injected at each round; and (c) **scratchpad hint**—a suggestion to maintain a shared intermediate workspace.

The planning round yields the most consistent gains (~ 5 – 8 on Level-II/III); the protocol reminder helps primarily on Level-I; the scratchpad hint benefits intermediate scales ($N=10$ – 20) but cannot prevent collapse at $N \geq 50$. Critically, qualitative failure patterns remain stable across all conditions: agents continue to communicate actively while failing to translate interaction into correct distributed computation, and the Communication-Reasoning Gap persists regardless of scaffolding. This confirms that the bottleneck reflects genuine LLM limitations in distributed information synthesis rather than a prompting artifact.

Case	Failure Mode	Description	Key Evidence
1	Consensus Failure	In task I-05, agents communicated extensively but failed to converge, submitting three distinct answers: {1176, 1182, 1167}.	97 of 100 agents submitted 1182; Agent-3 submitted 1176; Agent-55 submitted 1167.
2	Computation Error	Agent-10 in task I-05 received complete data from all peers but computed an incorrect range count.	Submitted 619 instead of correct answer 620. Arithmetic error during final aggregation.
3	Consensus Failure	In task I-05 (different instance), 50 agents split between two answers despite communication.	49 agents submitted 619; Agent-27 submitted 631.
4	Premature Submission	Agent-77 in task I-06 submitted before collecting sufficient data.	Submitted after receiving data from only 28 of 100 agents. Answer: 208; Expected: 114.
5	Consensus Failure	100 agents in task I-06 produced 12 distinct answers for XOR checksum task.	Answers ranged from 42 to 238. Majority (86 agents) converged on 146.

Table 14: Representative failure cases from DeepSeek-V3.1 experiments.

	Premature	Consensus	Compute
Premature	112	67	45
Consensus	–	90	52
Compute	–	–	86

Table 15: Co-occurrence of failure modes. Diagonal: total occurrences; off-diagonal: joint occurrences.

Metric	Success	Failed
Verification Rate	98.7%	95.9%
Strategy Discussion Rate	93.5%	87.2%
Avg. Messages per Agent	31.2	27.4
Avg. Rounds to Completion	8.3	12.7

Table 16: Behavioral comparison between successful and failed runs.

Level	Leader Rate	w/ Leader	w/o Leader
I	27.5%	56.8%	62.1%
II	21.8%	23.5%	59.0%
III	23.8%	0.0%	33.3%

Table 17: Leader emergence rates and associated success rates by complexity level.

Scaffold	L-I SR	L-II SR	L-III SR
No scaffold (baseline)	62.9	39.9	11.5
+Planning round	64.3	47.2	17.8
+Protocol reminder	65.1	41.3	12.1
+Scratchpad hint	63.7	44.6	14.9

Table 18: Prompt scaffold ablation results under DeepSeek-V3.1 + P2P. SR = Success Rate (%).