# Communication Structure Matters: Hierarchical Centralized Planning for LLM-Based Multi-Agent Autonomous Driving

Anonymous ACL submission


## Abstract

LLM-based multi-agent systems have emerged as a promising approach for collaborative autonomous driving, enabling vehicles to resolve conflicts through natural language negotiation. However, existing systems adopt decentralized and sequential communication structures in which vehicles negotiate with many other vehicles, one at a time, causing negotiation latency to grow linearly and token usage to grow quadratically with the number of vehicles involved. In real-time driving, this scaling directly degrades driving performance. We identify four failure modes of this communication pattern, organized into two classes: negotiation quality failures arising from the absence of coordinating authority, and communication cost failures arising from sequential execution. To address these, we propose Hierarchical Centralized Planning (HCP), a communication structure that decomposes large conflict groups hierarchically and introduces hierarchically distributed centralized planning, reducing per-round latency from $\mathcal{O}(N \cdot T_{\text{LLM}})$ to $\mathcal{O}(T_{\text{LLM}})$. Experiments on a multi-vehicle driving benchmark across 46 scenarios and vehicle group sizes from 2 to 8 show that HCP improves driving score, reduces negotiation latency, and increases the consensus rate under matched model conditions, with structural efficiency advantages persisting across model configurations.


## 1 Introduction

Autonomous driving (AD) systems have the potential to significantly improve road safety and efficiency (Yurtsever et al., 2020). However, single-vehicle systems face fundamental limitations: their perception is bounded by their own sensors, they cannot anticipate the intentions of surrounding vehicles, and they resolve conflicts without any shared understanding of what other drivers plan to do (Wu et al., 2025). Vehicle-to-vehicle (V2V) multi-agent collaborative autonomous driving addresses these limitations directly by allowing vehicles to share observations and coordinate decisions, extending each vehicle's perceptual range and enabling joint resolution of conflicts that no single vehicle could handle alone (Malik et al., 2021). Traditional approaches to this multi-agent coordination problem include optimization-based methods that formulate multi-vehicle planning as constrained problems and learning-based methods that train cooperative policies through reinforcement learning. These achieve strong performance in the specific scenarios they are designed for, but generalize poorly to the diversity of real-world multi-vehicle interactions (Wu et al., 2025).

LLM-based multi-agent systems (LLM-MAS) provide a more flexible foundation for coordination by expressing context and constraints in a generalizable way and leveraging broad knowledge to reason about unseen scenarios (Wu et al., 2025). A central challenge in building such systems for V2V collaborative AD, however, is designing the communication structure: which agents talk to which and in what order. These structural choices directly address adaptability, responsiveness, latency, communication efficiency, and reasoning capabilities (Yan et al., 2025). The field has explored a range of designs. LanguageMPC (Sha et al., 2025) uses a centralized LLM agent to coordinate distributed vehicle agents. AgentsCoDriver (Hu et al., 2025a) and AgentsCoMerge (Hu et al., 2025b) take a fully decentralized approach, where each vehicle independently broadcasts messages to all others with no coordinating authority. CoLMDriver advances prior work by introducing a critic-feedback mechanism, dynamic grouping to identify relevant collaborators, and by linking language-based negotiation with executed waypoints whose outcomes inform subsequent decisions (Liu et al., 2025). Its actor negotiation, however, remains fully decentralized, as all vehicles in a conflict communicate with

all others in sequence.

This communication structure does not scale well. As the number of vehicles in a conflict grows, the time and computational cost required to reach a negotiated plan increase substantially, and driving performance degrades. Through theoretical, empirical, and qualitative analyses of how CoLMDriver's decentralized and sequential communication structure configuration performs across a variety of real-time driving scenarios, we identify the root causes of these structural limitations.

One class of problems concerns *negotiation quality.* Without a coordinating authority, vehicles lack a shared starting point and independently reason about actions, often defaulting to overly cautious behavior that delays conflict resolution. As more vehicles are involved, this issue compounds: uninvolved vehicles are drawn into negotiation, redundant requests are repeated, and the group fails to reach a workable agreement.

A second class of problems concerns *communication cost*. Due to sequential negotiation, the total time required to complete a negotiation grows with each additional vehicle in the conflict. The amount of information each vehicle must process also grows, as every vehicle reads the full history of what all prior vehicles have said. In real-time driving, where a vehicle cannot wait indefinitely for negotiation to finish, these costs directly hurt driving performance.

To solve these problems, a straightforward method is to introduce a single coordinating agent into the communication. However, having a single agent reason about all vehicles in a conflict becomes increasingly difficult as more vehicles are involved. Furthermore, this centralization creates a singular point of failure, leaving centralized planning errors uncorrected before execution.

Addressing these requirements jointly leads to our proposed communication structure, **Hierarchical Centralized Planning (HCP)**. Rather than having all vehicles negotiate freely, a centralized planner provides an initial plan to all vehicles simultaneously. Large conflict groups are decomposed into independent subgroups, keeping each planner's reasoning burden manageable and distributing planning authority across multiple subgroup planners rather than concentrating it in a single agent. A global leader then reconciles subgroup plans and resolves cross-subgroup conflicts before execution.

Our contributions are as follows:

- We identify and characterize four failure modes of decentralized and sequential communication in LLM-based collaborative autonomous driving, organized into two classes: negotiation quality and communication cost, supported through theoretical, empirical, and qualitative analysis.
- We propose HCP, a communication structure that eliminates linear latency scaling and quadratic token scaling through hierarchical group decomposition and centralized planning, addressing all identified failure modes.
- We evaluate HCP on the InterDrive benchmark across 46 scenarios, demonstrating that under matched model conditions, HCP improves driving score from 60.51 to 73.73, reduces negotiation latency from 3.32s to 1.32s, and increases the consensus rate from 0.642 to 0.709.

## 2 Related Work

### 2.1 LLM Multi-Agent Systems: Communication Structure

Communication structures in LLM-based multi-agent systems can be characterized along multiple dimensions (Yan et al., 2025). We focus on architecture (centralized, decentralized, hierarchical, hybrid) and strategy (sequential, simultaneous), as these two dimensions have the most direct implications for latency, scalability, and coordination quality in time-sensitive driving scenarios.

Flat/decentralized structures support rapid decision-making but scale poorly as agent count grows. Centralized and hierarchical structures improve coordination quality but risk overloading the central planner. Hybrid structures mitigate both concerns by combining hierarchical oversight with local agent reasoning (Yan et al., 2025).

### 2.2 LLM Multi-Agent Systems for Collaborative Autonomous Driving

Existing work varies primarily along two dimensions: communication architecture and strategy. On architecture, LanguageMPC employs a centralized structure (Sha et al., 2025), using a single LLM planner to resolve multi-vehicle conflicts. In contrast, AgentsCoDriver, AgentsCoMerge, CoDrivingLLM, and CoLMDriver adopt decentralized architectures, enabling direct peer-to-peer communication between vehicles (Wu et al., 2025; Liu

et al., 2025; Fang et al., 2025). KoMA and CoMAL take a different approach, building shared memory pools through which agents publish and retrieve information to coordinate behavior (Yao et al., 2025; Jiang et al., 2024); neither addresses real-time action-level negotiation between vehicles with conflicting trajectories.

On strategy, most prior work operates under a sequential pattern, where vehicles communicate in turn to converge on a coordinated policy. CoLMDriver and AgentsCoDriver exemplify this approach. AgentsCoMerge differs, using a round-based broadcast model where all agents simultaneously receive messages, reason, and broadcast their intentions within each timestep (Hu et al., 2025b). LanguageMPC represents a further contrast, as its distributed vehicle LLMs simultaneously report status to a central coordinating LLM, enabling parallel information exchange (Sha et al., 2025). The distinction between sequential and simultaneous patterns has direct implications for latency and scalability, which we analyze in Section 4.1.

Recent work has also explored complementary locations of centralization in collaborative driving. LLM+Debrief (Cui et al., 2025) improves cooperative driving through centralized post-episode debriefing, demonstrating that centralized structure benefits coordination; however, it operates offline and targets strategy learning rather than real-time negotiation. V2V-LLM (Chiu, 2026) centralizes at the perception-fusion layer, using a single LLM node to fuse multimodal sensor data from multiple CAVs for grounding and planning, rather than action-level negotiation between vehicles.

Despite this variety in architecture and strategy, no existing work investigates which communication structure achieves the best driving performance and communication efficiency, or how structure choice scales across different scenario types and group sizes.

## 3 Preliminaries

Among the systems referenced in Section 2.2, CoLMDriver represents a strong foundation for this work. Its negotiation module follows an actor-critic paradigm: in each round, $N$ vehicles negotiate sequentially, each generating a message conditioned on the accumulated dialogue history, followed by a critic module that evaluates the proposed joint behavior via safety and efficiency scores (Liu et al., 2025), and returns corrective feedback for the next round. In this work, we retain the same critic module unchanged when comparing communication structures. This loop repeats until consensus is reached or a maximum of $R_{\max}$ rounds is exhausted.

Central to the negotiation process is the concept of a *conflict group*: the set of vehicles whose planned trajectories are determined to be in mutual conflict and who therefore participate in a shared negotiation. CoLMDriver dynamically identifies conflict groups at each planning step using a check on predicted vehicle trajectories; vehicles whose trajectories intersect are grouped together. The size of this group, $N$, is the primary variable governing negotiation cost and is determined by traffic and scenario geometry.

We define the following quantities used throughout the analysis. Let $N$ be the group size, $R$ the number of negotiation rounds, $L_{\text{msg}}$ the average message length in tokens, $L_{\text{desc}}$ the per-vehicle scenario description length, and $T_{\text{LLM}}$ the average LLM inference time per call. These quantities are treated as roughly constant across vehicles within a negotiation and are measured empirically on the InterDrive benchmark in Section 4.1.3.

## 4 Methodology

### 4.1 Limitations of Decentralized and Sequential Communication

The actor phase in the CoLMDriver framework follows a decentralized and sequential communication structure. We analyze this structure theoretically, then verify the predictions qualitatively and empirically.

#### 4.1.1 Theoretical Analysis

**Token scaling.** In each round, vehicles speak sequentially. The $k$-th speaker receives the messages of all $k-1$ prior speakers appended to its prompt, in addition to the scenario description:

$$\text{Prompt}(k) = L_{\text{desc}} + (k-1) \cdot L_{\text{msg}}. \quad (1)$$

Summing across all $N$ speakers per round and $R$ rounds:

$$\text{Total tokens} = R \cdot \left( N \cdot L_{\text{desc}} + \frac{N(N-1)}{2} \cdot L_{\text{msg}} \right). \quad (2)$$

Total token cost scales as $O(R \cdot N^2 \cdot L_{\text{msg}})$, growing **quadratically** with $N$: each additional vehicle both contributes its own message and forces all subsequent speakers to read a longer history.

**Latency scaling.** The $k$-th speaker incurs a wait of $(k-1) \cdot T_{\text{LLM}}$ before its turn. Total latency per round is therefore $N \cdot T_{\text{LLM}}$, and across $R$ rounds:

$$\text{Latency} = \mathcal{O}(R \cdot N \cdot T_{\text{LLM}}). \quad (3)$$

Latency scales linearly with $N$. In real-time deployment, where vehicles are locked into their current plan for the duration of negotiation, this overhead directly degrades driving performance.

#### 4.1.2 Qualitative Analysis

Beyond the scaling costs captured by Equations 2 and 3, analysis of negotiation logs reveals two additional structural failure modes attributable to the current communication structure.

**Lack of Initial Direction.** With no coordinating authority, vehicles must reason from scratch at the start of each negotiation. This manifests in two ways. First, without an initial plan, each vehicle defaults to the most conservative safe action. As shown in Listing 1, all seven vehicles produce identical messages in round 0, wasting an entire negotiation round. Second, in some instances, vehicles issue instructions to others rather than committing to a policy for themselves. As shown in Listing 2, V3 directs other vehicles without stating its own action, making its final assigned action ambiguous.

**Listings 1–2: Negotiation Quality Failures**

| **1: Deadlock** | **2: Ambiguous Policy** |
|---|---|
| V0–V6: "I will yield and stop." | V3: "V0, yield to V1. V1, go ahead." V0: "V1, I will yield." V2: "V1, please proceed." |

**Convergence Failures in Large Groups.** As group size grows, the absence of coordination produces two related failure patterns. First, a conflict between a subset of vehicles forces the entire group through additional rounds of negotiation. In Listing 3, a persistent V1↔V6 disagreement causes the critic to re-flag identical conflicts across three rounds. Vehicles 2, 3, 4, and 5 are never flagged as conflict pairs, yet all four degrade their speed actions to `SLOWER` or `STOP` by round 2 and the critic's score falls from 5.97 to 3.33. Second, vehicles in this network possess no mechanism to determine whether a response is necessary or useful. In Listing 4, four vehicles issue identical yield requests to Vehicle 3, adding no new information, and the score degrades from 3.69 to 2.84.

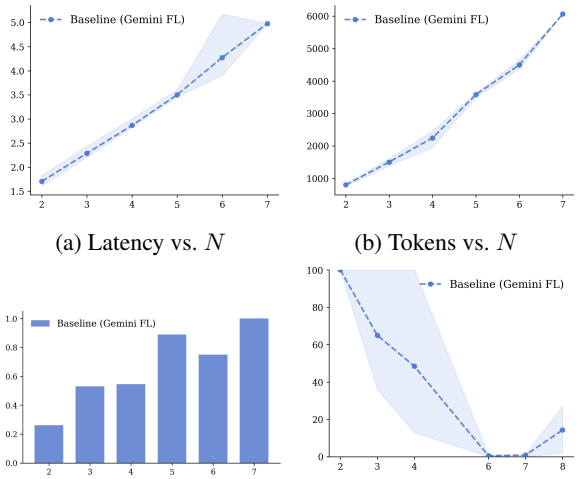


(a) Latency vs. $N$ (b) Tokens vs. $N$

(c) No-consensus rate vs. $N$ (d) Driving score vs. $N$

Figure 1: Empirical verification of baseline scaling behavior.

**Listings 3–4: Convergence Failures in Large Groups**

| **3: Collateral Damage** | **4: Redundant Requests** |
|---|---|
| **R1:** V1↔V6 disagree; V0,V2–V5: tell V1 to yield; score: 5.97→4.63 **R2:** V1↔V6 repeat; V2,V4=SLOWER; V3,V5=STOP; score: 4.63→3.33 | **R2:** V0,V1,V4,V5 [×4]: "V3, please yield" V3: "V1, please yield to me." score: 3.69→2.84 |

Together with the scaling costs derived theoretically, these observations motivate the design decisions of our proposed structure, which we evaluate against each failure mode shown in Table 1.

#### 4.1.3 Empirical Analysis

Figure 1a–1d confirm the theoretical predictions and qualitative observations on the InterDrive benchmark : latency scales linearly with $N$ (Pearson $r = 0.69$), token cost grows superlinearly ($R^2 = 0.69$, quadratic fit), and failure rate rises from 0.30 at $N = 2$ to 1.00 at $N = 7$. Pearson correlations between communication overhead and driving score ($r = -0.57$ for latency, $r = -0.65$ for tokens) suggest overhead as a contributing factor, though larger-$N$ scenarios are also inherently more complex. We set $k_{\max} = 4$ as failure rate exceeds 0.57 for all $N \geq 4$.

### 4.2 Proposed Structure: Hierarchical Centralized Planning (HCP)

The limitations identified in Section 4.1 motivate two core design decisions. First, the decline in negotiation quality and increased reasoning com-

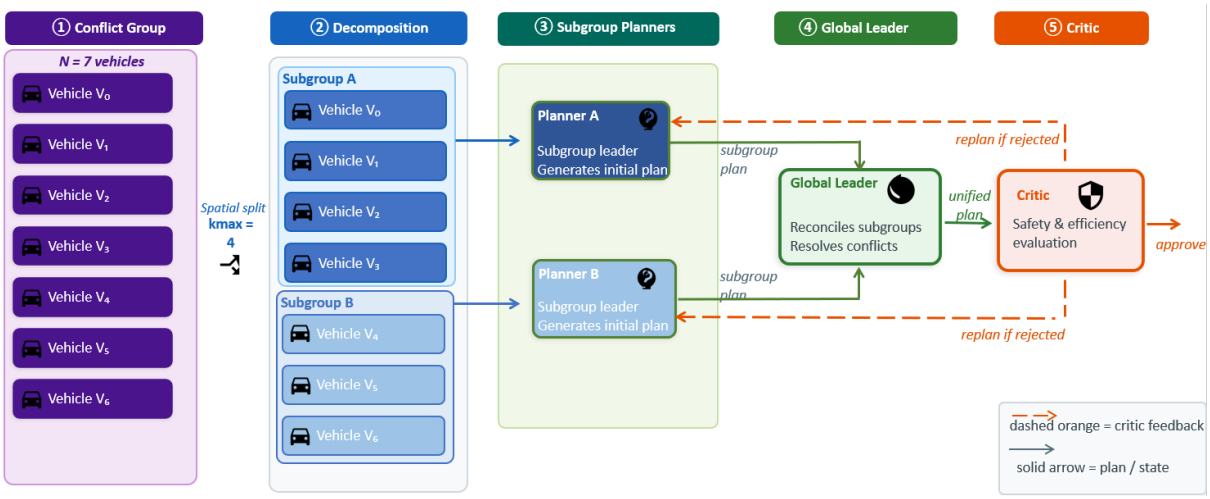


Figure 2: HCP communication structure for a conflict group of $N = 7$ vehicles. Vehicles whose predicted trajectories conflict are grouped and spatially partitioned into balanced subgroups of size at most $k_{\max} = 4$. Each subgroup planner generates an initial joint action plan concurrently. A global leader reconciles subgroup plans into a unified plan, resolving cross-subgroup conflicts. The critic evaluates safety and efficiency; if rejected, subgroup planning restarts, otherwise the plan is executed.

| Problem | Base | Central | Hier. | HCP |
|---|---|---|---|---|
| P1: No shared initial plan | × | ✓ | × | ✓ |
| P2: Convergence issues in large groups | × | ✓ | ∼ | ✓ |
| P3: Latency increases with group size | × | ✓ | ∼ | ✓ |
| P4: Token growth from long histories | × | ✓ | ∼ | ✓ |
| P5: High reasoning burden for LLMs | ✓ | × | ✓ | ✓ |

Table 1: Comparison of communication structures across identified failure modes. P1–P4 are observed limitations of the sequential baseline; P5 is a structural risk introduced by centralization. HCP is the only structure that fully addresses all failure modes. ✓ = mitigated, ∼ = partially mitigated, × = unaddressed.

plexity in large groups require methods that keep negotiation groups small. Second, the lack of initial direction, the risk of a single point of failure, and token and latency scaling constraints jointly motivate distributing planning authority across multiple centralized subgroup planners, each reasoning over a bounded set of vehicles. We refer to the resulting framework as Hierarchical Centralized Planning (HCP). Figure 2 illustrates the full structure.

#### 4.2.1 Design Decision 1: Hierarchical Group Decomposition

To address convergence failures in large groups, conflict groups larger than a threshold $k_{\max}$ are recursively partitioned into subgroups of size at most $k_{\max}$. Each subgroup negotiates independently, and subgroup plans are forwarded to a global leader who reconciles them into a unified plan. Because subgroups are spatially coherent, vehicles within a subgroup are physically proximate and therefore more likely to be in mutual conflict with each other than with vehicles in the other subgroup, limiting cross-subgroup dependencies that the global leader must resolve. This decomposition also bounds per-planner reasoning complexity, as subgroup planners reason over at most $k_{\max}$ vehicles, and the global leader reconciles already-negotiated subgroup plans rather than planning from scratch over the full vehicle set.

#### 4.2.2 Design Decision 2: Centralized Subgroup Planning

To address the lack of initial direction, we introduce a centralized planner within each subgroup that generates an explicit action plan for all vehicles before any negotiation occurs. This eliminates the deadlock observed in Section 4.1.2, as vehicles receive a concrete starting plan rather than reasoning independently from scratch.

Centralized planning within bounded subgroups also addresses the single point of failure risk inherent in fully centralized designs. Rather than concentrating planning authority in a single agent reasoning over all $N$ vehicles, HCP distributes it across multiple subgroup planners, each reasoning

over at most $k_{\max}$ vehicles. A global leader reconciles subgroup plans and resolves cross-subgroup conflicts before execution.

Replacing $N$ sequential vehicle turns with a single planner call per subgroup also eliminates the scaling costs derived in Section 4.1.1. Per-round latency reduces from $\mathcal{O}(N \cdot T_{\text{LLM}})$ to:

$$\text{Latency} = \mathcal{O}(R \cdot T_{\text{LLM}}). \tag{4}$$

Total token cost reduces from $\mathcal{O}(R \cdot N^2 \cdot L_{\text{msg}})$ to:

$$\text{Total tokens} = \mathcal{O}(R \cdot N \cdot L_{\text{msg}}). \tag{5}$$

Each of the $\lceil N/k_{\max} \rceil$ subgroup planners receives a scenario description and generates a single joint plan for at most $k_{\max}$ vehicles, with no accumulated dialogue history. The global leader reconciles $\lceil N/k_{\max} \rceil$ subgroup plans as input. Both contributions sum to $\mathcal{O}(R \cdot N \cdot L_{\text{msg}})$, growing linearly rather than quadratically with $N$.

#### 4.2.3 Mechanism

Figure 2 shows the full HCP structure. Initial conflict groups are identified using the same grouping procedure as CoLMDriver.

The conflict group is partitioned into spatially coherent subgroups of size at most $k_{\max}$. At each split, the two most spatially distant vehicles are selected as seeds, and remaining vehicles are assigned to the nearest seed while maintaining balanced subgroup sizes. Each subgroup planner generates an initial plan and distributes assigned actions directly to its vehicles.

These subgroup plans are forwarded to a global leader, which generates a unified global plan by resolving cross-subgroup conflicts while preserving local assignments. The global leader plans based on subgroup plans, prior round history, and critic feedback. If the critic finds the plan unsafe or inefficient, subgroup planning restarts from the beginning of the round. This process repeats until critic approval or $R_{\max}$ is reached. When $N \leq k_{\max}$, there is a single subgroup and the global leader reduces to the subgroup planner itself.

For all group sizes, safety and efficiency evaluations are performed using the original CoLMDriver critic, allowing structural changes to be isolated from changes in the evaluation module. We set $k_{\max} = 4$ based on the empirical analysis in Section 4.1, where failure rate and driving performance collapse consistently for $N > 4$ in the baseline structure.

Each negotiation round consists of two serial LLM hops regardless of group size: subgroup planners execute in parallel, followed by a single global leader call to reconcile subgroup plans.

### 4.3 Discussion

To situate HCP within the broader design space, we consider three reference structures: **Base**, following CoLMDriver's decentralized sequential design; **Centralized**, replacing sequential turns with a single LLM planner reasoning over all vehicles simultaneously; and **Hierarchical**, decomposing conflict groups into subgroups each running the baseline structure independently. Together, these structures span the centralized, hierarchical, and hierarchical-centralized points in the prior architectural design space, enabling controlled structural comparison within a unified pipeline. Table 1 analyzes how each addresses the failure modes from Section 4.1, together with the structural limitations introduced by a centralized planner; HCP is the only structure that fully addresses all five.

## 5 Experiments

Our experiments evaluate whether the communication structure design decisions motivated in Section 4.2 translate to performance improvements across four dimensions: negotiation quality, communication efficiency, driving performance, and robustness to model capability. We evaluate all communication structures on the InterDrive benchmark under both Gemini 2.5 Flash-Lite and Qwen2.5-3B configurations.

### 5.1 Experimental Settings

**Benchmark.** We evaluate all communication structures on the InterDrive benchmark introduced by Liu et al. (2025). InterDrive comprises three scenario types designed to induce multi-vehicle conflicts: intersection crossing (IC), highway lane changing (LC), and ramp merging (LM). We evaluate on the 46 all-LLM scenarios, with group sizes ranging from $N = 2$ to $N = 8$. We evaluate all structures in real-time mode such that inference latency directly impacts driving performance, making it a more faithful reflection of deployment conditions.

**Metrics.** For driving performance, we adopt the InterDrive metrics from Liu et al. (2025): Route Completion (RC), Infraction Score (IS), and Driving Score (DS = RC × IS) as the primary metric.

Table 2: Results under Gemini 2.5 Flash-Lite and Qwen2.5-3B conditions. Best values per model are **bold**.

| Model | Structure | Latency (s) | Tok/Rnd | Cons. Rate | Rounds | IS | RC | DS | IC | LM | LC |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Gemini 2.5 Flash-Lite | **HCP** | **1.32** | **430** | **0.709** | **1.0** | **0.74** | **98.89** | **73.73** | **66.67** | **86.43** | **57.89** |
| | Baseline | 3.32 | 1033 | 0.642 | 1.36 | 0.62 | 96.39 | 60.51 | 58.68 | 72.19 | 35.42 |
| Qwen2.5-3B | **HCP** | **0.76** | **430** | 0.536 | **1.0** | **0.70** | 88.19 | 64.09 | 61.88 | 74.76 | **42.36** |
| | Baseline | 1.64 | 1306 | **0.627** | 1.48 | 0.67 | **97.21** | **66.82** | **67.76** | **80.74** | 29.88 |

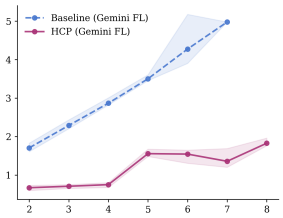


Figure 3: Median time per negotiation round vs. $N$ under Gemini Flash-Lite.

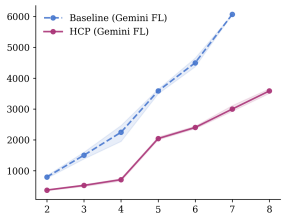


Figure 4: Prompt tokens per negotiation round vs. $N$ under Gemini Flash-Lite.

For communication efficiency, we measure average negotiation latency per conflict group and prompt tokens used per round. For negotiation quality, we report consensus rate and average rounds to consensus among successful negotiations.

**Models.** We use Gemini 2.5 Flash-Lite across all structures, evaluated over 3 independent seeds with metrics averaged across trials. We additionally evaluate under Qwen2.5-3B, the model used in the original CoLMDriver (Liu et al., 2025), to assess whether structural advantages persist across model configurations. If negotiation does not reach critic approval within $R_{\max}$ rounds, vehicles proceed with the plan from the final round. Observed conflict group sizes vary by structure, scenario type, and model configuration, as longer per-call inference times reduce the frequency of simultaneous large-group conflicts.

## 5.2 Results

HCP addresses each failure mode identified in Section 4.1 and outperforms the decentralized and sequential baseline on driving performance under matched model conditions. Table 2 summarizes results across both model configurations; unless otherwise noted, the following analysis refers to the Gemini Flash-Lite configuration.

### 5.2.1 Addressing Negotiation Quality

Section 4.1.2 identified two negotiation quality failures in the baseline structure: lack of initial direction (Listings 1–2) and convergence failures in large groups (Listings 3–4). Under matched model conditions, HCP achieves a consensus rate of 0.709 compared to 0.642 for the baseline, and resolves negotiations in 1.0 rounds on average compared to 1.36. The reduction to a single round indicates that the deadlock and redundant coordination patterns from Listings 1–4 are largely eliminated by design.

### 5.2.2 Addressing Communication Cost

Under matched model conditions, HCP achieves a mean negotiation latency of 1.32s compared to 3.32s for the baseline. Figure 3 shows group latency per negotiation as a function of $N$. HCP maintains lower absolute latency across all group sizes, with the separation growing with $N$, consistent with the theoretical prediction that HCP's latency advantage scales with conflict group size.

HCP also reduces total token usage per negotiation. As shown in Figure 4, total tokens grow quadratically with $N$ for the baseline, consistent with the $\mathcal{O}(R \cdot N^2 \cdot L_{\text{msg}})$ bound derived in Equation 2, while HCP scales linearly, consistent with the structural guarantee from Equations 4 and 5.

### 5.2.3 Driving Performance

Under matched model conditions, HCP achieves a driving score of 0.74 compared to 0.61 for the baseline, a 22% improvement. The gain is driven primarily by infraction score (0.74 vs. 0.62), with route completion near ceiling for both structures (98.9 vs. 96.4), indicating that HCP's centralized planning produces safer coordinated behavior rather than simply completing more routes. HCP outperforms the baseline across all three scenario types: intersection crossing (0.67 vs. 0.59), lane merging (0.86 vs. 0.72), and lane changing (0.58 vs. 0.35), with the largest absolute gain in lane changing.

### 5.2.4 Model Structure Interaction

To assess whether these advantages persist across model configurations, we evaluate both structures

Table 3: Ablation results (Gemini Flash-Lite). Each structure removes one design decision from HCP. Best values are **bold**.

| Struct. | DS(Total) | Lat. | Tok/Rd | Cons. Rt. | Rds |
|---|---|---|---|---|---|
| **HCP** | **73.73** | **1.32** | **430** | **0.709** | 1.0 |
| Hier. | 56.67 | 3.55 | 1115 | 0.657 | 1.0 |
| Central | 72.47 | 1.78 | 463 | 0.664 | 1.0 |

under Qwen2.5-3B, the lightweight model used in the original CoLMDriver (Liu et al., 2025).

Token efficiency advantages persist and are more pronounced under Qwen than under Flash-Lite. Tokens per round decrease from 1,306 to 430, a reduction of over 3×. Latency remains lower under HCP (0.76s vs. 1.64s). These efficiency gains are structural and model-independent.

Under Qwen, HCP achieves a slightly lower driving score than the baseline (0.64 vs. 0.67), in contrast to the clear advantage under Flash-Lite. Decomposing DS reveals that this gap is driven by route completion (88.2 vs. 97.2) rather than infraction score, where HCP marginally improves over the baseline (0.70 vs. 0.67). Unlike the baseline, where each vehicle reasons only about its own action, the HCP planner must simultaneously generate a coherent plan for all subgroup vehicles, occasionally producing plans that fail critic approval under limited model capacity. The consensus rate drop under Qwen (0.536 vs. 0.627) is consistent with this interpretation.

Taken together, these results demonstrate that HCP's efficiency advantages are model-independent, while achieving strong driving performance under centralized planning requires sufficient model capacity. The route completion gap under Qwen reflects a model capability threshold inherent to centralized planning rather than a structural weakness of HCP.

#### 5.2.5 Ablation Study

To isolate each design decision (DD), we evaluate two ablations of HCP. Hierarchical removes DD2 (centralized planner), running the baseline sequential structure within each subgroup. Centralized removes DD1 (hierarchical decomposition), using a single planner over all vehicles simultaneously.

As shown in Table 3, removing DD2 (Hierarchical) degrades driving score to match the baseline (0.57), confirming that decomposition alone provides no benefit without a coordinating planner. Token cost per round remains high (1,115), as sequential vehicle turns within each subgroup still accumulate dialogue history. This isolates the centralized planner as the primary driver of both driving performance and token efficiency.

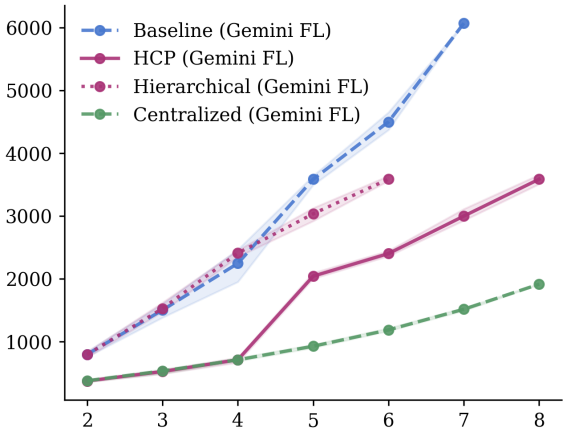


Figure 5: Prompt tokens per negotiation round vs. $N$ across ablation structures under Gemini Flash-Lite. The centralized planner is the primary driver of token efficiency.

Removing DD1 (Centralized) achieves competitive driving score (73.73 vs. 72.47) but lower consensus rate (0.664 vs. 0.709) and higher latency (1.78s vs. 1.32s). HCP's hierarchical decomposition bounds each planner's reasoning scope to at most $k_{\max}$ vehicles, avoiding the unbounded reasoning complexity that a single centralized planner incurs as $N$ grows.

Figure 5 illustrates that the centralized planner is the dominant factor in token efficiency: both HCP and Centralized achieve substantially lower tokens per round than Hierarchical (430 and 464 vs. 1,115). HCP's combination of both design decisions achieves the best performance across all metrics simultaneously.

## 6 Conclusion

Experiments on the InterDrive benchmark demonstrate that under matched model conditions, HCP outperforms the baseline across all metrics simultaneously, with the largest gains in lane-changing scenarios. Structural efficiency and driving performance advantages persist across model configurations, while achieving strong coordination quality under centralized planning requires sufficient model capability. These results establish communication structure as a critical and underexplored design dimension in LLM-based collaborative driving.

## 7 Limitations

This work evaluates fixed communication structures, but optimal structure selection likely depends on conflict complexity, group size, and scenario type. A natural extension is adaptive structure selection, where a controller dynamically chooses a communication topology based on observed conflict properties. Evaluation is conducted on the InterDrive benchmark, which covers three scenario types in a simulated urban environment. Performance may differ across more diverse traffic conditions, road geometries, and driving cultures not represented in the benchmark.

## References


Hsu-kuang et al. Chiu. 2026. V2v-llm: Vehicle-to-vehicle cooperative autonomous driving with multimodal large language models. *IEEE International Conference on Robotics and Automation (ICRA)*.

Jiaxun Cui, Chen Tang, Jarrett Holtz, Janice Nguyen, Alessandro G. Allievi, Hang Qiu, and Peter Stone. 2025. Towards natural language communication for cooperative autonomous driving via self-play. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Shiyu Fang, Jiaqi Liu, Mingyu Ding, Yiming Cui, Chen Lv, Peng Hang, and Jian Sun. 2025. Towards interactive and learnable cooperative driving automation: a large language model-driven decision-making framework. *IEEE Transactions on Vehicular Technology*.

Senkang Hu, Zhengru Fang, Yiqin Deng, Zihan Fang, Xianhao Chen, and Yuguang Fang. 2025a. Agentscodriver: Large language model empowered collaborative driving with lifelong learning. In *Proceedings of the 2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE.

Senkang Hu, Zhengru Fang, Zihan Fang, Yiqin Deng, Xianhao Chen, Yuguang Fang, and Sam Kwong. 2025b. Agentscomerge: Large language model empowered collaborative decision making for ramp merging. *IEEE Transactions on Mobile Computing*.

Kemou Jiang, Xuan Cai, Zhiyong Cui, Aoyong Li, Yilong Ren, Haiyang Yu, Hao Yang, Daocheng Fu, Licheng Wen, and Pinlong Cai. 2024. Koma: Knowledge-driven multi-agent framework for autonomous driving with large language models. *Preprint*, arXiv:2407.14239.

Changxing Liu, Genjia Liu, Zijun Wang, Jinchang Yang, and Siheng Chen. 2025. Colmdriver: Llm-based negotiation benefits cooperative autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 25951–25960.

Sumbal Malik, Manzoor Ahmed Khan, and Hesham El-Sayed. 2021. Collaborative autonomous driving—a survey of solution approaches and future challenges. *Sensors*, 21(11):3783.

Hao Sha, Yao Mu, Yuxuan Jiang, Li Chen, Chenfeng Xu, Ping Luo, Shengbo Eben Li, Masayoshi Tomizuka, Wei Zhan, and Mingyu Ding. 2025. Languagempc: Large language models as decision makers for autonomous driving. *Preprint*, arXiv:2310.03026.

Yaozu Wu, Dongyuan Li, Yankai Chen, Renhe Jiang, Henry Peng Zou, Wei-Chieh Huang, Yangning Li, Liancheng Fang, Zhen Wang, and Philip S. Yu. 2025. Multi-agent autonomous driving systems with large language models: A survey of recent advances. *Preprint*, arXiv:2502.16804.

Bingyu Yan, Zhibo Zhou, Litian Zhang, Lian Zhang, Ziyi Zhou, Dezhuang Miao, Zhoujun Li, Chaozhuo Li, and Xiaoming Zhang. 2025. Beyond self-talk: A communication-centric survey of llm-based multi-agent systems. *Preprint*, arXiv:2502.14321.

Huaiyuan Yao, Longchao Da, Vishnu Nandam, Justin Turnau, Zhiwei Liu, Linsey Pang, and Hua Wei. 2025. Comal: Collaborative multi-agent large language models for mixed-autonomy traffic. *Preprint*, arXiv:2410.14368.

Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. 2020. A survey of autonomous driving: common practices and emerging technologies. *IEEE Access*, 8:58443–58469.


## Appendix

### A Implementation Details

We build on the CoLMDriver codebase and implement all communication structures within its actor negotiation module, keeping all other pipeline components fixed. We follow CoLMDriver for all low-level components including the VLM intention planner and the intention-guided waypoint planner (Liu et al., 2025).

Experiments are run in CARLA 0.9.10.1 with a fixed simulation step of 0.05s (20 Hz) and an agent decision rate of 5 Hz. In real-time mode, vehicles that are mid-negotiation continue executing their previous plan until the negotiation cycle completes, effectively reducing the decision rate for the duration of negotiation. The communication range is 70m and the maximum number of negotiation rounds is $R_{\max} = 3$.

Malformed LLM outputs trigger a retry up to a fixed limit, after which a default action is assigned. For hierarchical structures, conflict groups larger than $k_{\max} = 4$ are recursively partitioned into spatially balanced subgroups using a 2-medoid algorithm, as described in Section 4.2.1.

## B Hardware and Concurrency

Gemini 2.5 Flash-Lite is accessed via the Vertex AI API, enabling true parallel subgroup planner execution. Qwen2.5-3B is run locally using A100 GPUs; as local GPU resources are shared, subgroup planner calls are executed sequentially and latency is computed as the maximum across subgroup planning times to simulate parallel execution. The global leader and critic execute after all subgroup planners complete.

## C Subgroup Decomposition Algorithm

Given a conflict group of $N > k_{\max}$ vehicles, subgroups are formed as follows: (1) select the two most spatially distant vehicles as seeds; (2) assign each remaining vehicle to its nearest seed; (3) rebalance by moving boundary vehicles between subgroups until $|\text{subgroup}_A| - |\text{subgroup}_B| \leq 1$; (4) recurse on any subgroup still exceeding $k_{\max}$. When $N \leq k_{\max}$, no decomposition occurs and the global leader reduces to the single subgroup planner itself.

## D Benchmark Scenario Distribution

Table 4 lists the 46 InterDrive all-LLM scenarios used in evaluation, grouped by scenario type and vehicle count. IC scenarios cover straight, signalized, and opposing-traffic intersections; LM scenarios cover ramp merging and roundabout configurations; LC scenarios cover highway lane changes and complex intersection crossings.

Table 4: InterDrive benchmark scenario distribution across type and vehicle count.

| Type | Routes | $N$ | Count |
|---|---|---|---|
| IC | 7–10 | 2 | 4 |
| | 11–16 | 2 | 6 |
| | 23, 26 | 3, 6 | 2 |
| | 31–32, 37–38, 41, 44 | 4–8 | 6 |
| | *Subtotal* | | *18* |
| LM | 1–6, 17–22 | 2 | 12 |
| | 24, 27, 29–30, 45 | 3–4 | 5 |
| | 33–34, 42 | 4 | 3 |
| | *Subtotal* | | *20* |
| LC | 25, 35–36, 43 | 4 | 4 |
| | 28, 39–40, 46 | 6–7 | 4 |
| | *Subtotal* | | *8* |
| **Total** | | | **46** |

## E $k_{\max}$ Sensitivity Analysis

Table 5 reports HCP performance under Gemini 2.5 Flash-Lite across $k_{\max} \in \{3, 4, 5\}$. Performance is broadly stable across all values, with $k_{\max} = 4$ achieving the best driving score (73.73), lowest latency (1.32s), and highest consensus rate (0.709). At $k_{\max} = 3$, smaller subgroups increase negotiation rounds (1.22 vs. 1.0) and latency (1.59s), as more cross-subgroup conflicts must be resolved by the global leader. At $k_{\max} = 5$, consensus rate drops slightly (0.658), consistent with larger subgroups approaching the reasoning complexity threshold identified in Section 4.1.3. These results confirm that $k_{\max} = 4$ is optimal within this range.

| $k_{\max}$ | DS | Lat. (s) | Tok/Rnd | Cons. Rate | Rounds | RC |
|---|---|---|---|---|---|---|
| 3 | 71.23 | 1.59 | 450 | 0.676 | 1.22 | 97.14 |
| 4 | **73.73** | **1.32** | **430** | **0.709** | **1.0** | 98.89 |
| 5 | 72.86 | 1.46 | 450 | 0.658 | **1.0** | 99.39 |

Table 5: HCP sensitivity to $k_{\max}$ under Gemini 2.5 Flash-Lite. $k_{\max} = 4$ performs best across driving score, latency, and consensus rate.