Orchestrating Reasoning and Reaction: An Asynchronous Hierarchical Framework for LLM-driven Traffic Signal Control

*Fansheng Sun , Jiyu Wang , Zhidan Liu**

The Hong Kong University of Science and Technology (Guangzhou)

** Corresponding author: zhidanliu@hkust-gz.edu.cn*

KDD2026 August 9–13, 2026 Jeju, Korea

1 Background & Motivation

TSC dynamically allocates green-light durations across urban intersections. Its goal is to reduce waiting time, improve traffic throughput, and prevent congestion from propagating across the road network..

MARL-based Methods

Fast adaptive control through local interaction and neighborhood communication.
However, limited observation leads to myopic coordination and weak global awareness.

LLM-driven Agents

Strong semantic reasoning and zero-shot strategic planning for network-wide TSC.
However, direct LLM control suffers from latency, high context cost, and topological hallucination.

Our Gogal

Bridge strategic LLM reasoning and reactive MARL control across different time scales.
Achieve topology-grounded network-wide TSC through hierarchical asynchronous synergy.

Key Challenges

- **C1** Computational burden vs. real-time response
- **C2** Spatial-topological mismatch
- **C3** Asynchronous coordination instability

2 Our Proposal: Astra

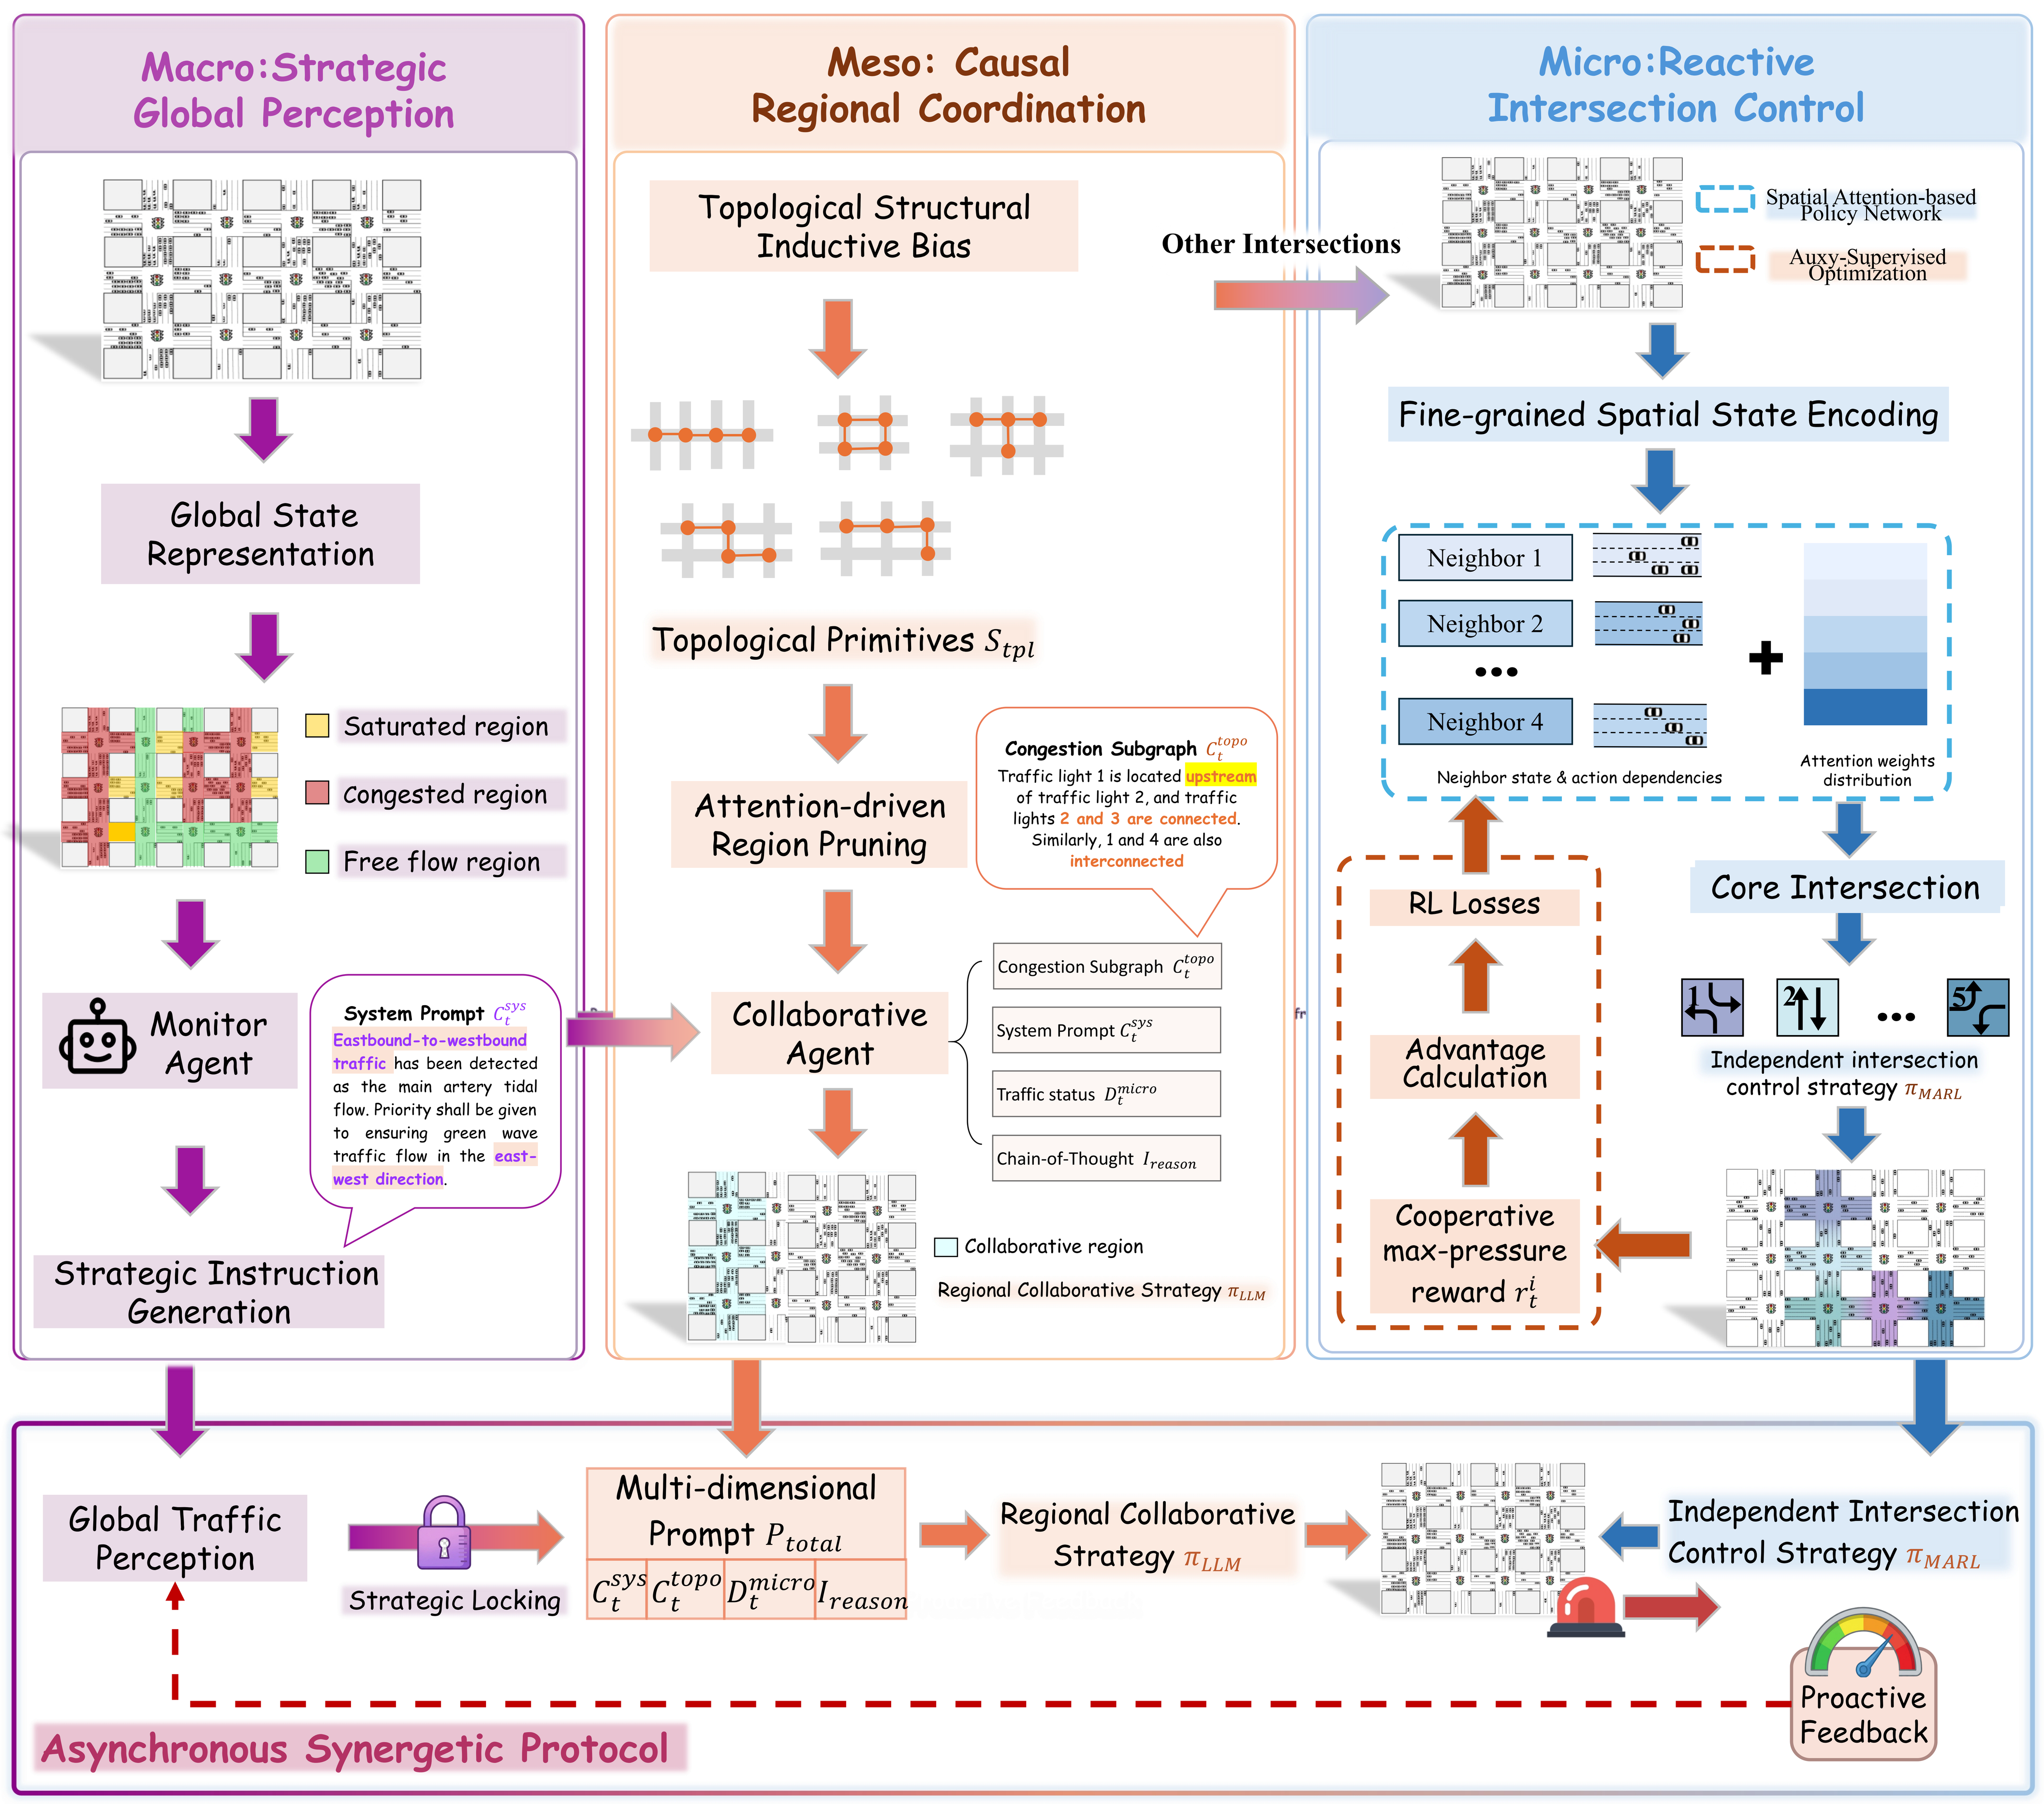

1. **Macro:** Global bottleneck detection via low-frequency LLM reasoning.
2. **Meso:** Topology-grounded causal coordination over congestion subgraphs.
3. **Micro:** High-frequency MARL control with spatial attention.
4. **Protocol:** Asynchronous locking and feedback for stable re-planning.

3 Experimental Results

A. Overall Performance

Model	Jinan		Hangzhou	
	ATT	AWT	ATT	AWT
FixedTime	358.49	297.23	470.92	297.5
Advanced CoLight	283.94	55.16	336.12	50.44
LLMLight	275.74	58.42	306.53	39.27
Traffic-R1	264.48	47.93	308.46	43.94
Dec-POMDP	285.36	59.24	310.14	43.32
Astra (Gemini-3-Flash)	**253.28**	**45.03**	**301.95**	**21.66**

Astra consistently achieves **SOTA** performance across 5 real-world datasets.

B. Efficiency: Tokens vs. Time

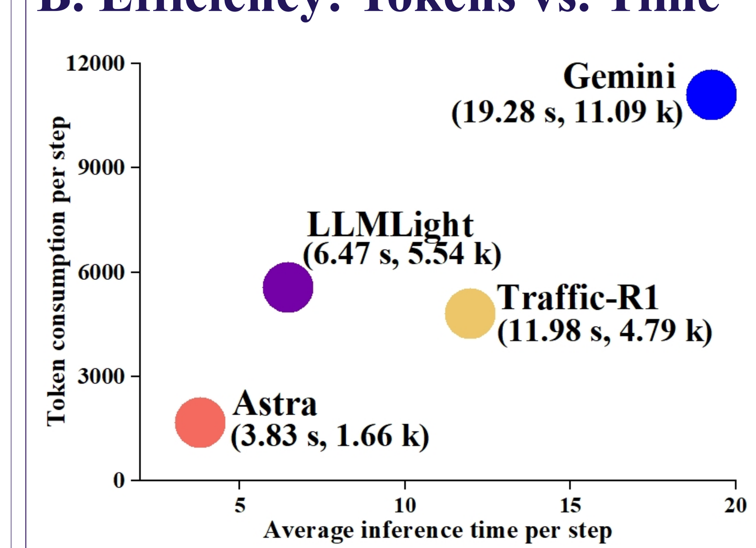

C. Component Analysis

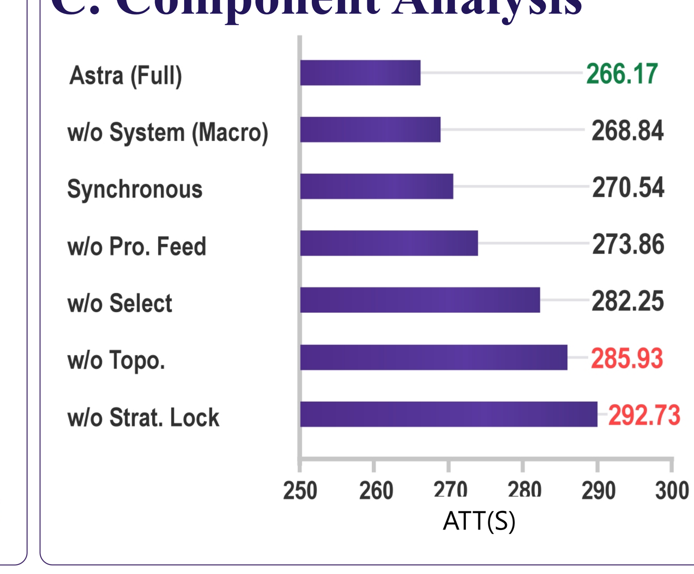

D. Large-scale Evaluation (Lower is Beter)

196 INTERSECTIONS

Method	New York-1 (196 intersections)		New York-2 (196 intersections)	
	ATT (s)	AWT (s)	ATT (s)	AWT (s)
Advanced-CoLight	1180.56	141.46	1467.89	261.37
Traffic-R1	932.74	77.24	1194.69	197.04
CoLLMLight	971.88	93.38	1235.37	198.49
Astra (Qwen3-30B)	**923.18**	**70.36**	**1140.82**	**154.84**

Astra sustains **SOTA** performance on city-scale networks.

E. Sensitivity

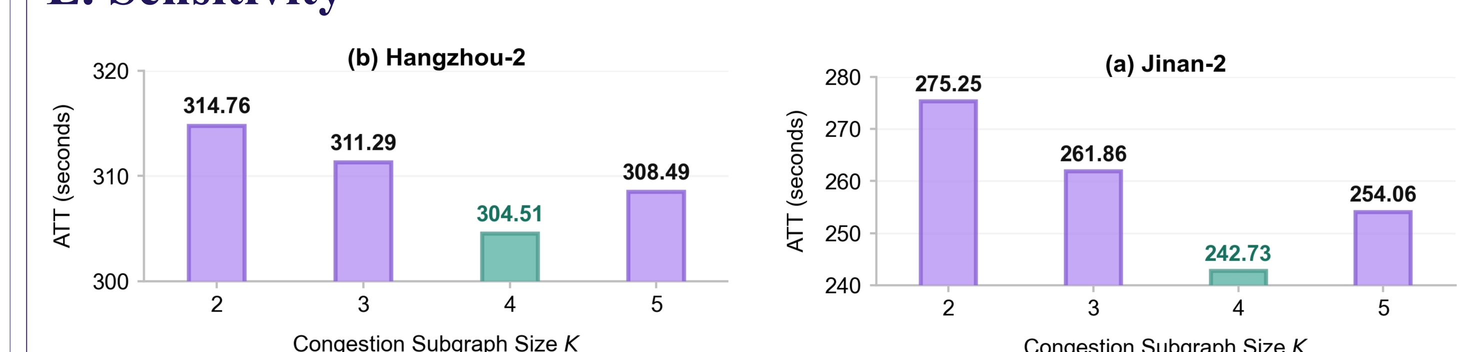

4 Key Contributions

Macro–Meso–Micro Hierarchy
A principled three-level hierarchy that integrates strategic reasoning and reactive control.

Topology-grounded LLM Coordination
Semantic causal primitives align LLM reasoning with the physical road structure.l.

Asynchronous Synergetic Protocol
Strategic locking and proactive feedback synchronize planning and reaction effectively.

Extensive Real-world Validation
Evaluated on 5 real-world datasets and 2 large-scale networks (196 intersections), achieving consistent SOTA results.

Astra is the SOTA across 5 real-world datasets and 2 city-scale networks.