GridSCAPE: An LLM-Based Multi-Agent Framework for Real-Time Intelligent Supply Chain Planning in Electric Power Enterprises

Haoyang He[1], Yan Gu[2,3*]
[1]State Grid Beijing Procurement Company, State Grid Beijing Electric Power Company, Beijing, China
[2]School of Control and Computer Engineering, North China Electric Power University, Beijing, China
[3]School of Computer Science and Statistics, Trinity College Dublin, Dublin, Ireland
hehaoyang@bj.sgcc.com.cn, yangu@ncepu.edu.cn, YAGU@tcd.ie

***Abstract*—Supply chain planning is a core challenge in supply chain management, orchestrating heterogeneous item categories across multiple vendors, and jointly optimizing end-to-end material flows, from sourcing and inbound logistics, through warehousing, to on-time delivery at demand locations. Electric power enterprises, as large and safety-critical consumers of engineering materials, require a comprehensive supply chain planning system, spanning demand planning, supplier and contract management, inventory control, to ensure reliable and cost-effective material availability. Moreover, recent advances in artificial intelligence (AI), particularly large language models (LLMs), have opened new directions for supply chain management, enabling knowledge-grounded reasoning, tool-augmented integration with enterprise platforms, and multi-agent coordination for closed-loop, adaptive planning. Accordingly, we present GridSCAPE, an LLM-based multi-agent framework, built to orchestrate closed-loop, real-time, intelligent supply chain planning for electric power enterprises. GridSCAPE deploys coordinated LLM agents with tool support, covering the entire supply chain, from demand and sourcing to inventory, warehousing, and logistics. We present implementation details of GridSCAPE and validate its efficacy via controlled experiments, showing reduced end-to-end cycle time, higher plan accuracy, and an increased in-stock rate.**

***Index Terms*—Large language model, multi-agent system, supply chain planning**

I. Introduction

Supply chain systems coordinate material, information, and financial flows across suppliers, warehouses, transport, and demand points on rolling horizons. With the rise of e-commerce platforms and big-data infrastructure, supply chain must provide real-time visibility, fine-grained demand sensing, rapid re-planning, and scalable, data-driven decision support. These expectations further entail multi-echelon synchronization, explicit policy compliance, and auditable decision trails under demand and lead-time uncertainty. For electric power enterprises, these capabilities are mission-critical to reliability, outage response, and budget compliance. In practice, the supply chain system is realized through enterprise platforms that unify data, codify policies, and support auditable, closed-loop planning. However, operational data are frequently fragmented across different modules and external feeds, complicating end-to-end reasoning and timely re-planning. Therefore, rigorous supply chain management is essential.

Supply chain planning (SCP) is a core function of supply chain management that converts strategy into time-phased purchasing and stocking actions while governing material flows across echelons. It links forecasting to execution policies, sets buffers that protect service targets, and coordinates movements within budget. In practice, SCP operates on a rolling, multi-period horizon that synchronizes procurement, inventory positioning, and inter-facility logistics under capacity, contractual, and budget constraints. In electric power enterprises, SCP is central to reliability and outage response. Long and uncertain lead times, capital-intensive spares, and mandated approvals create high exposure to disruption. Scenario-driven, disruption-aware planning is therefore essential, helping curb emergency buys, prevent idle inventory, and mitigate compliance risk. To ensure traceability and governance, SCP should integrate with enterprise platforms and produce auditable, KPI-driven plans that balance in-stock rate, plan accuracy, and end-to-end cycle time. A general framework about SCP in electric power enterprise is shown in Fig. 1.

Recent advances in AI—especially large language models (LLMs)—have enabled tool-using agents with retrieval, code execution, and stepwise reasoning. In multi-agent form, LLMs decompose objectives into role-specific tasks, coordinate under uncertainty, and interface safely with enterprise data and solvers. This capability aligns with supply-chain needs for policy-aware reasoning, rapid re-planning, and auditable decisions. By coupling LLM policies with retrieval-augmented grounding (RAG) and whitelisted tool adapters, agents generate evidence-backed, policy-conformant proposals with explicit citations, unit/period harmonization, and contract compliance checks. We therefore propose GridSCAPE, an LLM-based multi-agent framework for SCP in electric power enterprises. GridSCAPE fields role agents for planning, procurement, contract management, delivery, and stock management, under an overall manager that fuses enterprise-platform data, retrieves SOPs, and drives optimization and simulation.

In general, the main contributions of this work are as

*Corresponding author: Yan Gu (E-mail: yangu@ncepu.edu.cn).

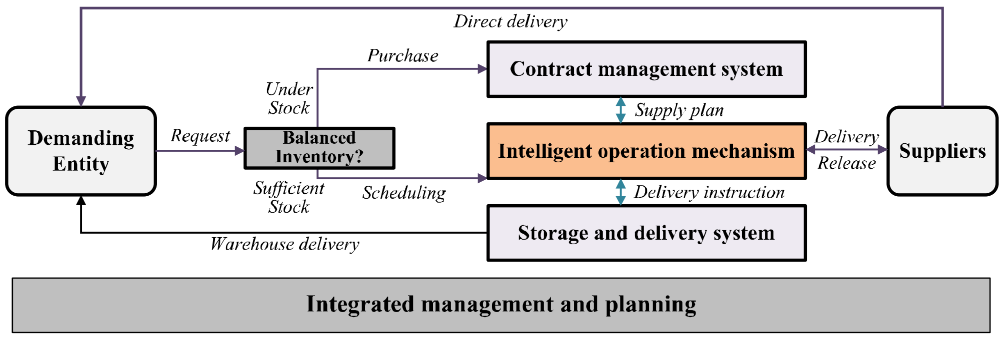

Fig. 1: The general framework about supply chain planning in electric power enterprise.

follows:

- We introduce GridSCAPE, an LLM-based multi-agent framework for supply chain planning in electric power enterprises, tightly integrated with the enterprise platforms to enable a closed-loop plan–execute–diagnose–replan workflow across demand, procurement, inventory, and logistics.
- We present a rigorous formulation and system design, including a multi-period supply chain planning model with uncertainty, a Dec-POMDP abstraction for role-based agents, and a tool-augmented pipeline that supports human-in-the-loop decision making.
- We detail the implementation of agent roles and coordination protocols, and we conduct a utility case study with controlled experiments, showing reduced cycle time, higher plan accuracy, and increased in-stock rate.

The rest of this paper is organized as follows: Section II reviews the literature. Section III introduces the system model and formulates the problem. Section IV describes the implementation of GridSCAPE. Section V evaluates the performance of GridSCAPE and Section VI concludes the work.

II. Related Work

A. Supply chain planning in asset-intensive industries

Classical supply chain planning has been extensively studied in the operations research literature, covering demand forecasting, material requirements planning, safety stock optimization, and lot-sizing, which remain fundamental to replenishment design [1]. Multi-echelon inventory systems are commonly modeled using frameworks such as the Clark–Scarf approach [2], and in asset-intensive industries, planning emphasizes reliability-driven availability and coordination of spares and maintenance materials across echelons, as seen in METRIC-type inventory control models [3]. Demand in these environments is often intermittent and event-driven, motivating forecasting methods such as Croston's approach and its refinements [4]. Although robust and stochastic optimization techniques can hedge uncertainty [5], their practical adoption remains limited in enterprise contexts. Recently, research has explored how large language models can be integrated into supply chain reasoning and tooling to enhance adaptability and automation in planning workflows [6].

B. LLMs and multi-agent systems

Large language models have enabled a new class of intelligent agents capable of reasoning and multi-step planning. Approaches such as ReAct demonstrate how LLMs can interleave reasoning with external actions to iteratively gather information and refine decisions [7], while Toolformer shows that LLMs can acquire tool usage capabilities in a self-supervised manner [8]. Retrieval-augmented generation provides access to external knowledge, enabling agents to ground outputs in up-to-date policies and domain facts [9]. Multi-agent frameworks such as CAMEL and AutoGen support coordinated reasoning and tool orchestration among multiple LLM agents [10]. These advances facilitate automated plan generation and iterative correction in workflows such as supply chain planning without replacing classical optimization models, instead acting as orchestration layers that select analytical tools, configure inputs, interpret outputs, and enforce constraints.

C. Agents in supply chains

Agent-based designs have a long history in supply chain management for distributed decision-making across organizational boundaries, including negotiation protocols like the Contract Net Protocol [11] and blackboard architectures [12], which enabled decentralized task allocation and coordination. Subsequent work applied multi-agent coordination to supply chain dynamics, demonstrating enhanced scalability and robustness in decentralized planning [13]. Recent studies explore LLM-augmented agents for collaborative planning and optimization across stakeholders, aiming to structure cooperation and maintain auditability in decision making [14]. Such agent architectures show promise for governance-aware, human-in-the-loop planning in complex enterprise supply chains. However, most existing approaches seldom integrate formal planning models or enforce operational constraints and contractual compliance. This work addresses these gaps by grounding agent reasoning in enterprise data and coupling decisions with structured multi-period planning objectives.

III. System Model and Problem Statement

This section introduces the overall framework of GridSCAPE, formalizes a multi-period supply chain planning model for electric power enterprises, and defines the corresponding planning problem. The proposed model serves as the structured feasibility and optimization backbone that reconciles LLM-based multi-agent proposals into executable, auditable supply chain plans.

A. The Framework of GridSCAPE

To operationalize rolling-horizon supply chain planning at scale, as illustrated in Fig. 2, the GridSCAPE framework comprises three tightly integrated components: enterprise data ingestion aligned to the moving planning horizon, retrieval-augmented grounding of policies and contracts, and an LLM-based multi-agent system coordinated by an overall manager. Role agents cover core supply chain functions including planning, procurement, contract management, delivery, and inventory management. Each agent consumes scoped operational data, queries the RAG corpus to retrieve relevant policy and contractual context, and proposes time-phased actions accompanied by traceable citations.

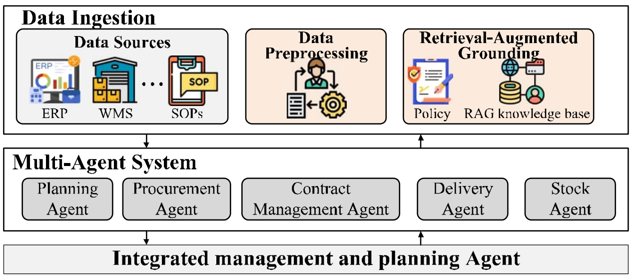

Fig. 2: The framework of GridSCAPE.

These role agents interact with multiple suppliers, warehouses, and delivery lanes to generate end-to-end, multi-period planning decisions, including purchase releases, safety-stock targets, transshipments, and schedule assignments. At each planning cycle, agent proposals are reconciled into a feasible plan that respects operational constraints and enterprise policies; period-t decisions are executed, feedback from execution is recorded, and the planning horizon advances to $t+1$. This rolling process yields a closed-loop, adaptive, and auditable supply chain planning workflow tailored to the operational requirements of Electric Power Enterprises.

B. Multi-Period Planning Model

We model supply chain planning on a rolling horizon as a time-phased program that unifies item, location, and period semantics, enabling consistent integration of role-agent proposals. A planning instance at roll t is defined as:

$$\mathcal{P}_t = \big(\mathcal{I}, \mathcal{L}, \mathcal{M}, \mathcal{T};\ D,\ L,\ \bar{C},\ \bar{U},\ \bar{Q},\ \bar{B},\ Q^{\min/\max},\ (\tilde{x}, \tilde{y})\big). \tag{1}$$

where $\mathcal{I}$ denotes the set of items, $\mathcal{L}$ the locations (including suppliers $\mathcal{S} \subset \mathcal{L}$, warehouses, and demand sites), $\mathcal{M}$ the set of transport modes, and $\mathcal{T} = \{t, \dots, t+T\}$ the planning horizon. The inputs comprise stochastic demand summaries Di,l,τ, lead-time parameters $L^m l \to l'$ and $L^{\text{sup}} i, s \to l$, warehouse capacities $\bar{C}_l$, lane capacities $\bar{U}^m l \to l', \tau$, supplier capacities $\bar{Q}s, \tau$, period budget bands $\bar{B}_\tau$, contract bounds $Q^{\min} i, s, Q^{\max} i, s$, and the published baseline plan $(\tilde{x}, \tilde{y})$ for the same calendar horizon.

Each decision is time-phased at the item–location–period granularity, defined as:

$$\Pi_t = \{x, y, I, u, e, SS\} \tag{2}$$

where $x_{i,s,l,\tau}$ denotes releases from supplier s to location l, $y^m i, l \to l', \tau$ denotes shipments between locations by mode m, Ii, l, τ is end-of-period inventory, $u_{i,l,\tau}$ is unmet demand (stockout/backorder), $e_{i,l,\tau}$ is emergency procurement, and $SS_{i,l,\tau}$ is an optional safety-stock target. Items may carry size coefficients v_i for storage accounting.

Material conservation incorporates transport and supplier lead-time lags. For each item i, location l, and period τ, inventory evolves as:

$$\begin{aligned} I_{i,l,\tau} = I_{i,l,\tau-1} &+ \sum_{s \in \mathcal{S}} x_{i,s,l,\tau - L^{\text{sup}}_{i,s\to l}} + \sum_{m \in \mathcal{M}} \sum_{l' \in \mathcal{L}} y^m_{i,l'\to l,\tau - L^m_{l'\to l}} \\ &- \sum_{m \in \mathcal{M}} \sum_{l' \in \mathcal{L}} y^m_{i,l\to l',\tau} - \big(D_{i,l,\tau} - u_{i,l,\tau}\big) + e_{i,l,\tau}. \end{aligned} \tag{3}$$

where inbound flows are offset by outbound shipments and realized demand, and emergency procurement provides an additional recourse mechanism. Feasibility is further ensured through capacity, contract, and budget constraints:

$$\begin{aligned} \sum_{i \in \mathcal{I}} v_i\, I_{i,l,\tau} &\le \bar{C}_l, && \forall\, l \in \mathcal{L},\ \tau \in \mathcal{T}, \\ \sum_{i \in \mathcal{I}} y^m_{i,l\to l',\tau} &\le \bar{U}^m_{l\to l',\tau}, && \forall\, m \in \mathcal{M},\ l, l' \in \mathcal{L},\ \tau \in \mathcal{T}, \\ \sum_{i \in \mathcal{I}} \sum_{l \in \mathcal{L}} x_{i,s,l,\tau} &\le \bar{Q}_{s,\tau}, && \forall\, s \in \mathcal{S},\ \tau \in \mathcal{T}. \end{aligned} \tag{4}$$

$$\begin{aligned} \sum_{i \in \mathcal{I}} \sum_{s \in \mathcal{S}} \sum_{l \in \mathcal{L}} \pi_{i,s}\, x_{i,s,l,\tau} &+ \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} \sum_{l \in \mathcal{L}} \sum_{l' \in \mathcal{L}} \chi^m_{l\to l'}\, y^m_{i,l\to l',\tau} \\ &\le \bar{B}_\tau, \quad \forall \tau \in \mathcal{T}. \end{aligned} \tag{5}$$

where v_i is the storage size coefficient, π and χ denote contracted purchase and transportation costs, respectively. Contractual obligations are enforced by:

$$Q^{\min} i, s \le \sum \tau \in \mathcal{T} \sum_{l \in \mathcal{L}} x_{i,s,l,\tau} \le Q^{\max}_{i,s}. \tag{6}$$

Because electric power enterprises place strong emphasis on plan stability, deviations from the published baseline plan are explicitly modeled as slack variables:

$$\delta^x_{i,s,l,\tau} = |x_{i,s,l,\tau} - \tilde{x}_{i,s,l,\tau}|, \quad \delta^y_{i,l\to l',\tau} = |y^m_{i,l\to l',\tau} - \tilde{y}^m_{i,l\to l',\tau}|. \tag{7}$$

The optimization objective jointly reflects enterprise-level KPIs. End-to-end cycle time is proxied by unmet demand, in-stock rate improves as unmet demand decreases, and plan accuracy is captured by deviations from the baseline plan. The resulting objective is:

$$\min_{\Pi_t,\delta} \quad \beta_1 \sum_{i,l,\tau} u_{i,l,\tau} + \beta_2 \sum_{i,s,l,\tau} \delta^x_{i,s,l,\tau} + \beta_3 \sum_{i,l,l',\tau} \delta^y_{i,l\to l',\tau}. \tag{8}$$

subject to the above constraints, where $\beta_1, \beta_2, \beta_3 \geq 0$ balance service performance, plan stability, and operational adherence. Only period-t decisions are executed, while future periods serve as planning targets. Execution feedback updates system states and the baseline plan, and the planning horizon advances to $t+1$.

C. Problem Statement

We formalize each roll as computing a time-phased plan given the current information state. We define the information bundle at period t as:

$$\Xi_t = \big(D,\ L,\ \bar{C},\ \bar{U},\ \bar{Q},\ \bar{B},\ Q^{\min/\max},\ (\tilde{x},\tilde{y})\big), \tag{9}$$

where D are demand summaries, L are lead-time parameters, $\bar{C}$,$\bar{U}$,$\bar{Q}$ are warehouse, lane adn supplier capacities, $\bar{B}$ is the period budget band, $Q^{\min/\max}$ are contract bounds, $(\tilde{x},\tilde{y})$ is the published baseline plan for the horizon $\mathcal{T} = \{t,\dots,t+T\}$. The decision is a plan, which defined as:

$$\Pi_t = x, y, I, u, e, SS, \tag{10}$$

that jointly optimizes service level, plan stability, and cost while satisfying operational, contractual, and budget constraints. The goal of GridSCAPE is to solve this rolling-horizon planning problem in a closed-loop manner by coordinating LLM-based agents with a structured optimization model, enabling real-time, intelligent supply chain planning for Electric Power Enterprises.

IV. Method Implementation

This section describes how GridSCAPE is implemented as a retrieval-augmented, tool-enabled multi-agent system for real-time supply chain planning in electric power enterprises. The implementation combines role-specialized LLM agents with enterprise data access, policy grounding, and a centralized feasibility mechanism based on the multi-period planning model, enabling closed-loop and auditable decision making.

A. LLM Usage and Retrieval-Augmented Grounding

GridSCAPE leverages large language models as role-specialized planners that operate over horizon-aligned enterprise states and a curated corpus of policies and contracts. Operational data, including inventory levels, open orders, capacities, budgets, and baseline plans, are serialized into structured, time-indexed representations. In parallel, standard operating procedures, contracts, and approval rules are indexed for retrieval with explicit version and effective-date controls. Given a role-specific objective, the LLM reasons over retrieved evidence and issues structured tool calls to generate time-phased planning proposals with traceable policy citations.

Role behavior is constrained through prompts that encode decision scope, KPI priorities (cycle time, plan accuracy, and in-stock rate), and enterprise-specific guardrails relevant to electric power enterprises. Retrieval filters ensure that only in-force and role-relevant clauses are considered, while output validators harmonize units and calendars and screen candidate decisions against capacity, budget, and contractual constraints. As a result, each agent produces an evidence-backed, policy-conformant proposal expressed in a common item–location–period format for coordination.

B. LLM-Based Agent Design

Each GridSCAPE agent follows a unified design that combines retrieval-augmented reasoning with structured tool use. Given its role, the agent retrieves relevant policy and contractual context, assembles a horizon-aligned enterprise state, and invokes external tools—including the multi-period planning model—to generate and assess candidate decisions. Outputs are validated for consistency across units, periods, capacities, budgets, and commitments, and provenance information (retrieved evidence, tool interactions, and deviations from the baseline plan) is recorded to support auditability.

Five role agents specialize this design according to organizational functions in Electric Power Enterprises, covering planning, procurement, contract management, delivery, and stock management. All agents express decisions using consistent item–location–period semantics, enabling coherent system-level coordination. An overall controller schedules agent execution, aggregates proposals, and enforces centralized feasibility via the multi-period model; when constraints bind, it provides diagnostic feedback to guide minimal-churn revisions, thereby combining decentralized, policy-aware reasoning with centralized feasibility and end-to-end traceability.

C. The Process of LLM Agents

Each LLM agent operates according to a repeatable five-stage process, forming a closed-loop rolling-horizon planning cycle. The process of LLM agents is shown in Fig. 3. In the task orchestration stage, the agent translates its role objective into an executable plan by defining the required tool sequence, output schema, KPI priorities, and validation checks. This plan specifies the structure of the final proposal, including decisions, assumptions, bounds, citations, and a validation summary.

Each LLM agent operates in a closed-loop manner by first retrieving horizon-aligned operational states from enterprise systems and grounding its reasoning in relevant SOPs, contracts, and approval rules through filtered retrieval. It then synthesizes operational and policy context to generate candidate, time-phased decisions, invoking the multi-period planning model for feasibility and target assessment and applying validation routines to detect capacity, budget, and contractual violations while explicitly controlling deviations from the published baseline. Validated proposals are reconciled and executed for the current period through a centralized feasibility operator, with all decisions and supporting evidence logged for auditability. Execution feedback, binding-constraint diagnostics, and KPI telemetry are subsequently incorporated

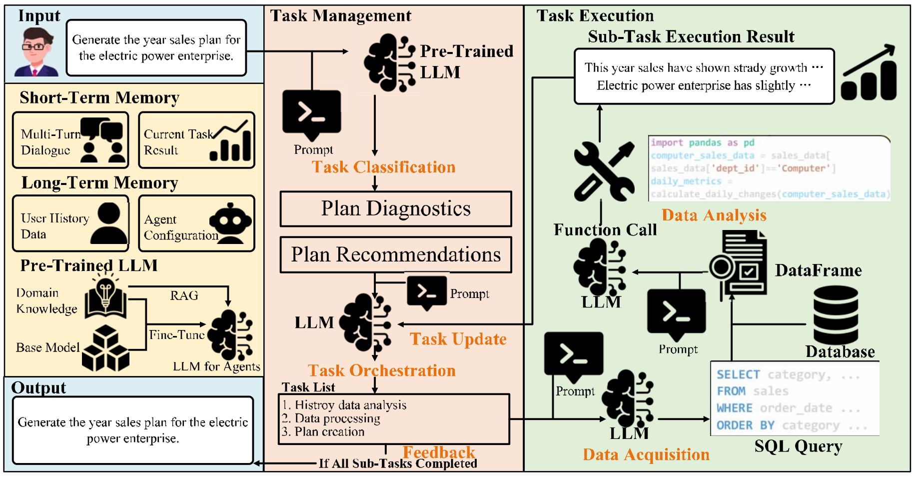

Fig. 3: The Process of LLM Agents.

to update assumptions and refresh the baseline plan, closing the rolling plan–execute–monitor–replan loop.

V. Experiments

This section evaluates GridSCAPE on realistic supply chain planning tasks. We describe the experimental setup, including datasets and baselines; introduce performance metrics aligned with the planning model objectives; and present results across varying problem sizes, different methodological approaches, and different LLM backbone sizes.

A. Setup

1) Datasets: We evaluate on three benchmark datasets representing supply chain planning scenarios with increasing scale and complexity:

- SupplyGraph Planning Benchmark: A publicly released graph-structured supply chain dataset including temporal production and demand data across realistic multi-tier networks [15].
- SynDelay Delivery Prediction Dataset: A synthetic dataset designed for delivery delay prediction and related planning challenges[16].
- Internal Power Enterprise Simulator: A simulated dataset capturing electric power enterprise supply chain operations.

Each dataset is instantiated at small, medium, and large problem sizes, controlling the number of items, locations, and planning horizon to test scalability.

B. Metrics

We evaluate performance using three primary indicators corresponding to the planning model objectives:

- Average Cycle Time (ACT): Average backlog exposure per period, capturing delays in meeting demand.
- Plan Accuracy (PA): Percentage deviation from baseline decisions in key dimensions (purchase and shipment decisions).
- In-Stock Rate (ISR): Proportion of periods where stock-out rates are below a predefined threshold.

These metrics align with operational KPIs for electric power enterprises and allow comparison across methods and models.

C. Baselines

We compare GridSCAPE against several recent methods:

- SupplyGraph-GNN Planner [15]: A graph neural network based planner using graph topology for forecasting and flow optimization, representing structured deep learning approaches to supply chain planning.
- PROPEL (ML+Optimization) [17]: A hybrid supervised and reinforcement learning framework that integrates learning with optimization to accelerate large-scale planning.
- LLM Planner No RAG: A baseline LLM agent that uses large language models directly for planning without retrieval-augmented grounding or centralized feasibility reconciliation.
- Traditional Rolling MILP: A mixed-integer linear programming solver applied in a rolling horizon fashion with no LLM or agent support, representing a classical optimization baseline.

D. Results

1) Different Dataset Sizes: We first analyze how methods scale with problem size. The results of average cycle time, plan accuracy, and in-stock rate are shown in TABLE I, TABLE II. and TABLE III, respectively.

TABLE I: Results of Average Cycle Time

Method	Small	Medium	Large
Traditional MILP	12.5	18.9	27.4
GNN Planner	11.2	17.8	26.1
LLM No RAG	10.4	16.9	24.8
GridSCAPE	**8.7**	**13.5**	**19.3**

TABLE II: Results of Plan Accuracy

Method	Small	Medium	Large
Traditional MILP	0.92	0.88	0.82
GNN Planner	0.93	0.89	0.84
LLM No RAG	0.90	0.85	0.80
GridSCAPE	**0.96**	**0.92**	**0.88**

TABLE III: Results of In-Stock Rate

Method	Small	Medium	Large
Traditional MILP	0.81	0.74	0.62
GNN Planner	0.83	0.76	0.65
LLM No RAG	0.85	0.79	0.71
GridSCAPE	**0.90**	**0.87**	**0.81**

As shown in TABLE I, GridSCAPE consistently yields lower cycle times, particularly for large problems, because coordinated agents with centralized feasibility avoid costly stockouts. As shown in TABLE II, GridSCAPE's RAG grounding and guideline enforcement yields higher accuracy relative to baselines. As shown in TABLE III, GridSCAPE achieves higher in-stock rates due to better constraint satisfaction and planning stability.

2) Different Methodological Approaches: We compare methods across metrics on the medium size dataset, results are shown in TABLE IV.

TABLE IV: Different Methodological Approaches

Metric	MILP	GNN Planner	LLM No RAG	GridSCAPE
ACT	18.9	17.8	16.9	**13.5**
PA	0.88	0.89	0.65	**0.92**
ISR	0.74	0.76	0.79	**0.87**

As shown in TABLE IV, GridSCAPE outperforms alternative approaches by a significant margin across all metrics, showing the benefit of RAG-grounded reasoning and feasibility reconciliation.

3) Different LLM Sizes: We evaluate GridSCAPE with three backbone model sizes commonly used in LLM research:

- Small (7B): 7-billion parameter backbone.
- Medium (30B): 30-billion parameter backbone.
- Large (70B): 70-billion parameter backbone.

TABLE V: Performance by LLM Size

Model	ACT	PA	ISR
Small 7B	15.2	0.89	0.83
Medium 30B	14.1	0.91	0.85
Large 70B	**13.5**	**0.92**	**0.87**

As shown in TABLE V, although larger backbones slightly improve performance, even the small 7B model achieves strong results, suggesting that GridSCAPE's structured framework reduces dependency on extreme model scale.

Overall, the experiments show that GridSCAPE achieves substantial improvements over both classical optimization and recent learning-based planning methods, especially in large-scale and realistic planning scenarios. Its design, combining structured optimization with RAG-grounded reasoning, is particularly effective for Electric Power Enterprise supply chain challenges, yielding lower backlogs, higher plan accuracy, and better inventory fulfillment.

VI. Conclusion

In this paper, we proposed GridSCAPE, an LLM-based multi-agent framework for real-time, closed-loop supply chain planning tailored to Electric Power Enterprises, which couples retrieval-augmented grounding of enterprise policies and contracts with a structured multi-period planning model to produce coordinated, auditable, KPI-aligned decisions; through comprehensive experiments on realistic large-scale planning scenarios, GridSCAPE demonstrated consistent improvements over classical optimization and recent learning-based baselines in average cycle time, plan accuracy, and in-stock rate, even with smaller LLM backbones, and ablation studies confirmed the critical role of RAG grounding, centralized feasibility coordination, and multi-objective planning in achieving these gains, highlighting the promise of hybrid systems that integrate structured optimization and intelligent, policy-aware reasoning for complex enterprise decision tasks.

References

[1] H. M. Wagner and T. M. Whitin, "Dynamic version of the economic lot size model," *Management Science*, vol. 5, no. 1, pp. 89–96, 1958.

[2] A. J. Clark and H. Scarf, "Optimal policies for a multi-echelon inventory problem," *Management Science*, vol. 6, no. 4, pp. 475–490, 1960.

[3] C. C. Sherbrooke, "Metric: A multi-echelon technique for recoverable item control," *Operations Research*, vol. 16, no. 1, pp. 122–141, 1968.

[4] J. D. Croston, "Forecasting and stock control for intermittent demands," *Operational Research Quarterly*, vol. 23, no. 3, pp. 289–303, 1972.

[5] D. Bertsimas, D. B. Brown, and C. Caramanis, "The theory of robust optimization," *SIAM Review*, vol. 53, no. 3, pp. 464–501, 2011.

[6] D. Simchi-Levi, I. Menache, J. Pathuri, and T. Linton, "Large language models for supply chain reasoning and tooling," *arXiv preprint*, 2025.

[7] S. Yao, J. Zhao *et al.*, "React: Synergizing reasoning and acting in language models," *arXiv preprint arXiv:2210.03629*, 2023.

[8] T. Schick, J. Dwivedi-Yu *et al.*, "Toolformer: Language models can teach themselves to use tools," *arXiv preprint arXiv:2302.04761*, 2023.

[9] P. Lewis, E. Perez, A. Piktus *et al.*, "Retrieval-augmented generation for knowledge-intensive nlp tasks," *NeurIPS*, 2020.

[10] G. Li *et al.*, "Camel: Communicative agents for mind exploration of large scale language model society," *arXiv preprint arXiv:2303.17760*, 2023.

[11] R. G. Smith, "The contract net protocol," *IEEE Transactions on Computers*, vol. 29, no. 12, pp. 1104–1113, 1980.

[12] B. Hayes-Roth, "A blackboard architecture for control," *Artificial Intelligence*, vol. 26, no. 3, pp. 251–321, 1985.

[13] J. M. Swaminathan, S. F. Smith, and N. M. Sadeh, "Modeling supply chain dynamics: A multiagent approach," *Decision Sciences*, vol. 29, no. 3, pp. 607–632, 1998.

[14] V. Jannelli, "Agentic llms in the supply chain: towards autonomous collaborative planning," *International Journal of Production Research*, 2025.

[15] A. T. Wasi, M. S. Islam, and A. R. Akib, "Supplygraph: A benchmark dataset for supply chain planning using graph neural networks," *arXiv preprint arXiv:2401.15299*, 2025.

[16] L. Xu, Y. Long, and A. Brintrup, "Syndelay: A synthetic dataset for delivery delay prediction," *arXiv preprint arXiv:2509.05325*, 2025.

[17] V. E. Akhlaghi, R. Zandehshahvar, and P. V. Hentenryck, "Propel: Supervised and reinforcement learning for large-scale supply chain planning," *arXiv preprint arXiv:2504.07383*, 2025.