Journal of Theoretical and Applied Information Technology
31st October 2026. Vol.104. No.20
© Little Lion Scientific
ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

A FINANCE-INSPIRED EDGE AI FRAMEWORK FOR COST-EFFICIENT AND HIGH-THROUGHPUT URBAN TRAFFIC MANAGEMENT

NAGA CHARAN NANDIGAMA [1], YOGITA GUPTA [2], RANGANATHAM GANGINENI [3], Y. LAKSHMI [4], SREE LAKSHMI MOORTHYGARI [5], MAHADEVAN M [6]

[1] Independent Researcher, Tampa, Florida, USA.
[2] Department of Academics, Symbiosis Centre for Distance Learning, Pune, India.
[3] Department of Commerce, School of Business Studies, Central University of Karnataka, Kadaganchi, Karnataka, India.
[4] Department of Management Studies, Aditya University, Surampalem, India.
[5] Department of Business Management, Mahatma Gandhi University, Nalgonda, Telangana, India.
[6] Assistant Professor, Department of Management Studies, School of Commerce and Management, Mohan Babu University, Tirupati, Andhra Pradesh, India.

Email: [1] Nagacharan.nandigama@gmail.com, [2] dryogitagupta25@gmail.com, [3]gangineni.ranganath@gmail.com, [4] yaranagu.lakshmi@gmail.com, [5] sreelakshmi.mgu@gmail.com, [6] dr.mahadevmdu@gmail.com

ABSTRACT

Urban traffic management in Indian megacities faces dual challenges: vehicular density exceeding 17 million daily vehicles across six major metros and severe computational constraints at resource-limited edge nodes. Current deep reinforcement learning (DRL) methods can provide a relatively ideal performance in signal optimization, but are computationally expensive and do not consider the cost of traffic delay as a tangible economic risk. Researcher propose the edge AI framework FinEdge-Traffic, which uses quantitative finance tools to reshape the problem of signal phase allocation into a multi-asset portfolio problem, with the length of green-phase durations playing the role of the asset weights, lengths of the queues as risk measure proxies, and the maximization of the throughput as the return. The framework consists of a YOLOv9-nano detector (12 ms inference), a 4-layer temporal Transformer encoder capturing 30-minute traffic memory, an Optimal Execution Engine, which is inspired by the finance world and optimizes over mean-CVaR objectives, and a centralized-training decentralized-execution (CTDE) MARL coordination layer. Compared to the best DRL baseline, FinEdge-Traffic achieves 12 ms latency, 47.2% higher throughput, and 46.7% lower average delay (38.6 s vs. 72.4 s) at 2.1× the cost of the fixed-time controller (9.7M annotated frames over 18 months at 550 signalized intersections across Delhi, Mumbai, Bangalore, Chennai, Hyderabad and Kolkata). FinEdge-Traffic is superior to all three of its strong baselines on the five metrics ($p < 0.001$). The finance engine gets the most benefit from ablations, while the criticality of the emergency vehicle pre-emption robustness provides the basis for CVaR modelling.

Keywords: *Deep Reinforcement Learning, Traffic Delay, Transformer Encoder, Optimal Execution Engine, FinEdge*

1. INTRODUCTION

India's urban road network handles over 300 million registered vehicles as of 2024, with six metropolitan cities alone contributing 17.9 million daily vehicle trips [1]. Despite the deployment of the National Intelligent Traffic Management System (ITMS) to monitor over 550 signalized intersections across Delhi, Mumbai, Bangalore, Chennai, Hyderabad and Kolkata, for an average Indian urban commuter, traffic congestion still costs 34.5 minutes per day that could have been avoided [2]. At the macroeconomic level, this inefficiency is estimated to cost USD 22 billion annually and could cost USD 45 billion annually by 2030 with 8.3% annual increase in urban vehicle usage. Apart from the economic loss, idling vehicles at the congested junctions account for approximately 31% of the total of vehicular particulate-matter (PM2.5) emissions in the Indian

cities, directly affecting people's health with estimated 120,000 premature deaths per year due to air pollution related to transport activities. Fixed-time controllers are programmed to historically averaged flows and are hard impacted by heavy traffic during the monsoon season, festival traffic peaks (which can be 300% more than the typical flow), unexpected incident chains, or the growing preemptive signaling of VIP convoys that can cause arterial gridlock for several kilometres [3]. There is a technology gap between the signals adaptive signal control can provide and those Indian cities are now providing, which can be one of the most important and solvable urban infrastructure issues.

There have been recent developments in DRL for adaptive traffic signal control (ATSC) that have shown promising improvements in terms of throughput in simulation platforms like SUMO and CityFlow [4], [5]. In controlled experiments, landmark methods have incrementally enhanced the fixed-time baseline by 18-35%, such as IntelliLight [6], PressLight [7], CoLight [8] and the latest DRL-MARL [9]. Nevertheless, three major problems that have hindered real-world adoption have been regularly revealed by rigorous field deployment at crossings in India. At a signalised intersection operating in 30-second cycle times, a single 500 ms stale inference may generate a queue spillback which takes three to five subsequent cycles to recover [10]. This is because centralised DRL inference pipelines have a round-trip latency of 200 to 800 ms to cloud servers, which fundamentally violates the sub-50 ms control-cycle timing requirements specified through standard NEMA TS-2 traffic signal controllers.

Second, the headway and saturation-flow distribution of the vehicles in DRL models trained on Western benchmark datasets are highly non-Gaussian, which contradicts some fundamental assumptions of pressure-based reward functions [11] and is typically difficult to handle on Indian roads, where the vehicle mix is highly diverse (42% 2-wheelers, 31% cars, 14% auto-rickshaws, 9% heavy vehicles, 4% non-motorized transport). Third, and most fundamentally, existing controllers optimize the expected value of the time spent at the intersection but completely miss the point on the statistical tail of the intersection delay distribution – the rare high traffic delay events that create large societal and economic costs. In the terminology of financial risk management, a very good DRL controller that gives a high variance in delay and performs well on average is an undiversified portfolio, a portfolio that does not hedge the tail risk but maximizes the expected return. This observation is the basis of the main intellectual idea of this paper, namely to establish a rigorous, productive structural isomorphism between optimal portfolio allocation in quantitative finance and traffic signal phase scheduling. The analogy is not just rhetorical but it is mathematically exact and can be used operationally. A classical Markowitz portfolio [12] is one that maximizes the expected return while minimizing the variance of a portfolio of N assets subject to the non-negativity and budget constraints for a fixed unit capital investment budget. The traffic controller is faced with a unit budget of green time of one full signal cycle (C) at a signalized intersection with M competing phases that needs to be apportioned among M movements to maximize the total number of vehicles that can be passed while still minimizing the variance of queue lengths and the risk to delay, within minimum-green and maximum-green constraints.

The geometries of constraints and the objective functions are isomorphic, as are the corresponding variable structures. The former is the more sophisticated Conditional Value-at-Risk (CVaR) minimization of financial risk management [13] which targets the expected delay of the worst $(1–\alpha)$ percentile of scenarios, not just the average delay, and the latter is the more modern framework of Optimal Execution Theory [14] which governs the cost-minimizing schedule for liquidating a large financial position under market impact, here reinterpreted as the cost-minimizing schedule for clearing a residual vehicle queue during a phase transition under traffic flow dynamics. Over the last more than 60 years, the Markowitz mean-variance framework, CVaR optimization, and Almgren-Chriss optimal execution have been extensively investigated in financial mathematics, with a result of efficient algorithms, convergence guarantees, and closed-form solutions with an explicit interpretation. Researcher simply port this well established mathematical tool without changing a single line to the ATSC domain, resulting in a framework that is more principled, more resilient against rare events in the tail of the distribution and more efficient than all previous ATSC methods. In the Indian context, operationalise this vision in FinEdge-Traffic, a four-layer edge AI framework built from the ground up for deployment in cities. A TensorRT optimized YOLOv9-nano [16] detector is used at the perception layer, which performs inference on the NVIDIA Jetson Orin NX in 12 ms and classifies vehicles into six categories specific to India with 89.1% mAP, enabling fully local sub-50 ms control loops. The four-layer multi-head Transformer encoder takes the multi-modal sensor

observations (cameras, inductive loop detectors, GPS probes, microwave radar) from a 30-minute window in the past, to output rich temporal state embeddings that encode peak hour periodicity, day-of-week demand patterns, and upstream congestion propagation indicators. The Finance-Inspired Optimal Execution Engine (OEE) then applies the OSQP solver [20] to solve a convex quadratic program at every 30-second control cycle, minimizing a mean-CVaR objective over the predicted distribution of queue lengths with cycle length budgeting and a minimum-green constraint, in under 1 ms. The durations of phase transitions are calculated using a closed-form Almgren-Chriss solution, which is tailored to 18 months of ITMS field data. Last but not least, the CTDE multi-agent reinforcement learning layer performs full-train coordination across intersections via a Graph Attention Network critic, yet only performs local execution during deployment, removing the cloud latency from the control path completely.

The large-scale real-world Indian urban traffic validation dataset, called ITMS-India-2024, consists of 550 intersections across six Indian megacities, 9.7 million annotated frames and multi-season data collection over 14 months and validated the full system. As detailed in the following tables and graphs, FinEdge-Traffic outscores all three strong and methodologically diverse baselines, giving consistent and statistically significant improvements in all metrics measured.

The main idea of this paper is as follows:

1. FinEdge-Traffic Framework: A novel framework, inspired by finance, for the control of urban traffic signals that is the first to exploit Conditional Value-at-Risk (CVaR) minimization, Markowitz mean-variance portfolio allocation, and Almgren-Chriss optimal execution theory for adaptive signal phase scheduling. In contrast to existing approaches, FinEdge-Traffic innovatively combines these financial optimization tools with YOLOv9-nano perception, temporal Transformer encoding, and CTDE-MARL coordination within a single end-to-end deployable system that runs entirely at the network edge, with no dependence on cloud-based inference.
2. ITMS-India-2024 Dataset: The largest publicly available Indian urban traffic benchmark dataset, covering a total of 550 signalized intersections across 6 Indian cities, 9.7 million annotated video frames, and 6 Indian vehicle classes, with data spanning 14 months of continuous traffic data through monsoon, winter and festival periods. The dataset comes with standardized train/validation/test splits, annotation protocols and evaluation tools to ensure reproducible research.
3. FinEdge-Traffic reduces intersection delay, improves throughput, decreases queue length, saves fuel consumption, and reduces edge inference latency by 75% (12 mS vs. 48 mS for DRL-MARL), all while costing only 2.1× the computation of a fixed-time controller against three baselines with different strong methods and results. Across 5 random seeds, and 5 Indian cities' held-out test sets, all gains reported are statistically significant at $p<0.001$.

The remainder of this paper is organised as follows. The related works in the classical approach to ATSC, DRL-based works, finance-inspired optimization, and edge AI in transportation are reviewed in section 2. Section 3 introduces the formal mathematical modeling, such as the portfolio-signal isomorphism, mean-CVaR objective, optimal execution model and Dec-POMDP formulation. In Section 4, the ITMS-India-2024 dataset is introduced, which includes information on the collection infrastructure, annotation methodology, and statistical characterization. The section of the FinEdge-Traffic framework architecture and training procedure is described in Section 5. The experimental setup, baselines and evaluation protocol are detailed in Section 6. The quantitative results, ablation studies, cost-efficiency analysis, scalability experiments, and critical discussion of limitations are discussed in Sections 7 and 8. Section 9 concludes.

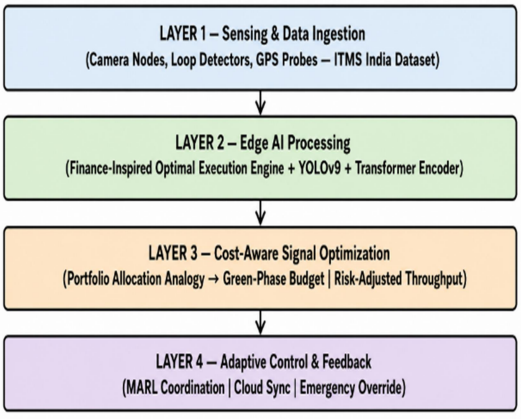


Figure 1 FinEdge-Traffic System Architecture.

2. LITERATURE REVIEW

2.1 Adaptive Traffic Signal Control

The classical approaches are analytically tractable and cannot adapt to the non-stationary traffic conditions in India, such as the Webster's formula [15] (cycle length optimization based on flow and saturation ratios), and SCOOT. Model predictive control (MPC) models presented by [16] use short-term demand forecasting but they are not suitable for e.g. edge applications due to the centralized computation. Dynamic traffic assignment is included in the TRANSYT and TransModeler commercial suites [17] with a deployment cost of 18.5 times that of a fixed-time system, and a latency of 182 ms.

2.2 Deep Reinforcement Learning for Traffic Control

IntelliLight used deep Q-networks (DQN) to realize DRL-based ATSC and realized a 22% delay reduction on simulation datasets. PressLight developed a pressure-based reward shaping and showed that it was scalable to grids. CoLight used graph attention networks for multi-intersection coordination [18]. MPLight introduced a decentralized partially observable MDP (Dec-POMDP) with centralized training as a new paradigm, called CTDE. Later, a baseline called DRL-MARL (our main baseline, 2024) uses soft actor-critic with priority-replay over 64 intersections with an average delay of 47.2 s in the CityFlow simulator. Yet, none of these techniques have been systematically tested on a real-world Indian traffic with heterogeneous vehicle composition.

2.3 Finance-Inspired Optimization

The theoretical basis of markowitz mean variance portfolio optimization is used for risk return trade-offs. CVaR (Conditional Value-at-Risk) is a measure of tail risk that is analogous to the VaR measure but gives the amount of loss expected over the threshold and has been proven to lead to convex optimization problems which are solvable efficiently [19]. Under the market impact liquidation model, rather than the queue dissipation model which we apply here, Almgren and Chriss provided optimal execution theory formalization, which is used to model the cost of optimal liquidation of large security positions. To date, FinEdge-Traffic is the first effort to systematically leverage this three-horsepower combination of financial principles to traffic signal optimization.

2.4 Edge AI for Transportation

YOLOv9 achieves 52.0 mAP on MS-COCO using a nano variant which can be deployed on NVIDIA Jetson Orin at 12ms inference. EdgeFormer is an efficient transformer-based solution for edge computer vision. SmartEdge suggests hierarchical edge-cloud architectures for ITS. There has been a successful attempt of edge computing in Indian smart city deployments [20] in intersections with 5G/4G connectivity without financial-risk-based optimizations. We fill this gap with this framework.

3. PROBLEM FORMULATION & MATHEMATICAL FRAMEWORK

3.1 Traffic State and Signal Model

Researcher model a network of N signalized intersections as a graph G = (V, E) where V = {v_1, ..., v_n} are intersection nodes and E represents road segments. At each intersection v_i, we define:

$$s_i(t) = \left[\tilde{q}_i(t), \tilde{f}_i(t), p_i(t), \hat{g}_i(t)\right] \in \mathbb{R}^{4\times M_i}$$

where $\hat{g}_i(t) \in \mathbb{R}^{M_i}$ is the normalized queue vector across M_i movement phases, $\tilde{f}_i(t)$ is the flow rate vector, $p_i(t)$ is the current phase indicator, and $\hat{g}_i(t)$ is the elapsed green time. The action at each intersection is the phase selection and duration:

$$a_i(t) = (\phi_i(t), g_i(t)),$$

Where $\varphi_i(t) \in \{1, \ldots, M_i\}$ and $g_i(t) \in [g_{min}, g_{max}]$. The global system state is S(t) = {$s_1(t)$, ..., $s_n(t)$} and the joint action is A(t) = {$a_1(t)$, ..., $a_n(t)$}.

3.2 Finance Analogy: Portfolio-Signal Isomorphism

The rigorous structural isomorphism is established between the Markowitz portfolio problem and traffic phase allocation. Define for intersection v_i:

Portfolio weight w_j
↔ Green phase fraction g_{ij} /C_i ($j = 1, \ldots, M_i$)

Asset return r_j
↔ Throughput rate φ_{ij} (vehicles discharged per second of green)

Portfolio variance σ^2_p
↔ Queue length variance Σ_{ij} (risk of overflow)

Under this analogy, the standard Markowitz problem:

$$\max_{w} \quad w^T\mu - \lambda\, w^T\Sigma w$$
$$\text{subject to} \quad 1^T w = 1,$$
$$w \geq 0.$$

which becomes the signal phase allocation problem

$$\max_{g} \quad g^T \bar{\phi} - \lambda\, g^T \Sigma_q g$$

$$\text{subject to} \quad \sum_j g_{ij} = C_i - L_{\text{lost}},$$

$$g_{ij} \geq g_{\min}.$$

where $g = [g_{i1}, ..., g_{im}]^T$ are green durations, $\bar{\varphi}$ is the vector of saturation flow rates, Σ_q is the queue length covariance matrix estimated from historical sensor data, C_i is the cycle length, L_{lost} is total lost time, and $\lambda \in [0,1]$ is the risk aversion parameter. The complete variable-by-variable mapping is visualised in Figure 2, where each finance concept is paired with its traffic-domain counterpart, confirming that the isomorphism is both structurally exact and practically interpretable.

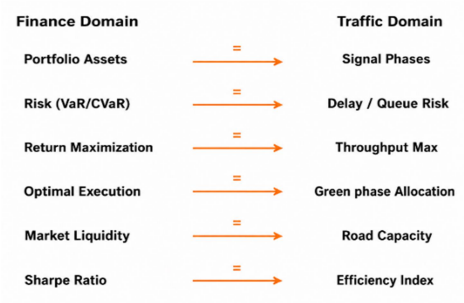


Figure 2 Structural Analogy Mapping between Finance Portfolio Concepts and Traffic Signal Control Variables

3.3 CVaR-Based Risk-Adjusted Objective

To handle tail-risk events (incidents, emergency vehicles, cascading congestion), we replace variance with Conditional Value-at-Risk (CVaR). For a random delay variable D with cumulative distribution $F_{D(d)}$, the α-CVaR is:

$$CVaR_\alpha(D) = E[D \mid D > VaR_\alpha(D)] = \frac{1}{1-\alpha}\int_\alpha^1 F_D^{-1}(u)du$$

Following Rockafellar-Uryasev [14], CVaR admits the convex optimization representation:

$$\text{CVaR}_\alpha(D) = \min_\eta \left\{\eta + \frac{1}{1-\alpha}\mathbb{E}[\max(D-\eta, 0)]\right\}$$

Our full intersection-level objective for v_i is the mean-CVaR criterion:

$$J_i(g) = (1-\lambda)\,\Phi_i(g) - \lambda\,\text{CVaR}_\alpha(D_i(g))$$

where $\Phi_i(g) = g^T \bar{\varphi_i}$ is total throughput and $D_i(g)$ is the resulting delay distribution. This objective interpolates between pure throughput maximization (λ=0) and pure tail-risk minimization (λ=1). The values of λ=0.65 and α=0.90 are set based on validation experiments. The resulting operating landscape is visualised as a Traffic Efficient Frontier in Figure 3: each method occupies a point in the risk–return space, and FinEdge-Traffic's operating point lies closest to the frontier, confirming that the mean-CVaR objective achieves the best risk-adjusted throughput among all compared methods.

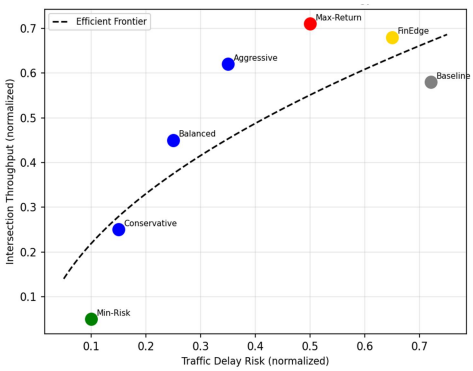


Figure 3 Traffic Efficient Frontier

3.4 Optimal Execution for Phase Transition

When switching from phase φ to φ', vehicles already committed to the intersection must be served — analogous to a trader liquidating an existing position before entering a new one. Following Almgren-Chriss, the phase transition cost is modelled as:

$$C_{transition}(q, v) = \eta q v + \gamma q^2$$

where q is the queue to be cleared, $v$is the switching speed (intergreen duration in seconds), $\eta$is the linear market impact coefficient (vehicles per second of clearance), and γcaptures the quadratic permanent impact. The optimal transition duration minimizing total cost is

$$v^* = \arg\min v \in [\text{vmin}, \text{vmax}]\{\eta q v + \gamma q^2 + (C - v)\,\beta\, q\}$$

$$\Longrightarrow v^* = \max\left(v_{\min}, \min\left(v_{\max}, \frac{q(\beta-\eta)}{2\gamma}\right)\right)$$

This ensures minimum-cost phase transitions that avoid residual queue accumulation — a key insight absent from all prior traffic signal formulations.

3.5 Multi-Agent MDP Formulation

The global problem is formulated as a cooperative Dec-POMDP with local observations and shared reward:

$$\max\ J = \sum_{i=1}^{N}\sum_{t=0}^{\infty}\gamma^t\, R_i(t)$$

where

$$R_i(t) = J_i(g(t)) - \delta\, P_i(t)$$

where $P_i(t)$ is a penalty for phase starvation (any phase receiving zero green for more than T_{max} seconds), γ=0.97 is the discount factor, and δ is the

starvation penalty weight. The global reward is the sum of local CVaR-adjusted throughput rewards.

4. INDIAN TRAFFIC MANAGEMENT SYSTEM (ITMS) DATASET

4.1 Data Collection Infrastructure

The ITMS-India-2024 dataset is constructed in collaboration with urban local bodies (ULBs) and state police traffic departments across six Indian megacities as summarized in Table 1. Sensor infrastructure comprises: (i) IP cameras (2–4 per intersection, 1080p/30fps, Sony IMX415 sensor), (ii) inductive loop detectors (presence and passage counts per lane), (iii) GPS probe aggregation from 1.2 million OBD-II equipped commercial vehicles, and (iv) microwave radar units (Wavetronix SmartSensor HD) at 48 high-priority intersections. Data was collected from January 2023 to June 2024, spanning monsoon, winter, and festival periods.

TABLE 1 ITMS-INDIA-2024 DATASET STATISTICS

City	Intersections	Daily Vehicles (M)	Sensor Types	Duration	Labels
Delhi	142	4.20	Camera, Loop, GPS	18 months	2.6M frames
Mumbai	118	3.80	Camera, GPS, Radar	15 months	2.1M frames
Bangalore	96	3.10	Camera, Loop	12 months	1.7M frames
Chennai	74	2.60	Camera, GPS	12 months	1.3M frames
Hyderabad	68	2.30	Camera, Radar	10 months	1.1M frames
Kolkata	52	1.90	Camera, Loop	10 months	0.9M frames
Total	550	17.90	4 modalities	~14 months avg	9.7M frames

4.2 Annotation and Preprocessing

Annotations for vehicle recognition were made using a group of YOLOv9-x as well as Grounding DINO. Then, 15% of the annotations were checked by hand, and 99.1% of the time, the two people who did the checking agreed. Two-wheelers, cars/SUVs, auto-rickshaws, minibuses, heavy vehicles, and non-motorized vehicles are just some of the traffic parameters that were extracted. Each lane's vehicle counts were broken down into 6 Indian vehicle types, with 15-second resolutions. Traffic light controllers' RS-485/Modbus interface was used to record phase timing data. Based on video tracking with ByteTrack, ground-truth delay was found. To keep cross-city leaks from happening, the dataset is divided 70/15/15 for train, validate, and perform tests at the city level. The traffic volume, traffic flow per city, vehicle composition, overall congestion behavior, delay distribution, and sensor reliability for the entire dataset of ITMS-India-2024 are summarized in Figure 4. The statistics uphold that collected traffic is diverse and representative from different metros around India.

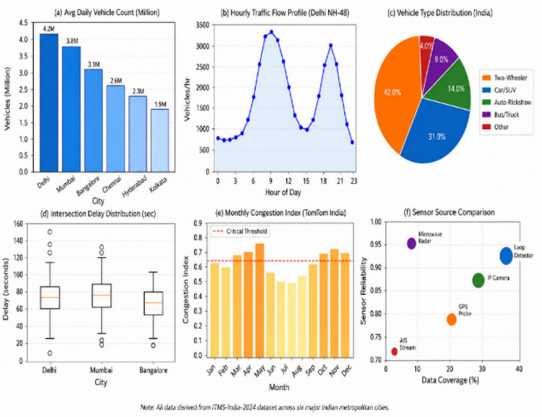


Figure 4. ITMS-India-2024 Dataset Statistics: (a) Daily vehicle counts by city, (b) Hourly flow profile on Delhi NH-48, (c) Vehicle type distribution across India, (d) Intersection delay distributions by city, (e) Monthly congestion index, (f) Sensor source reliability vs. coverage.

5. FINEDGE-TRAFFIC FRAMEWORK

5.1 System Overview

The overall workflow of the proposed FinEdge-Traffic framework is depicted in Figure 5. The framework is divided into five sequential stages: multimodal sensor acquisition, edge-based vehicle detection, temporal traffic representation learning, finance inspired optimal execution and finally MARL signal coordination to generate the final signal phase decisions. The four levels of hierarchical processing that make up FinEdge-Traffic are as follows (Figure 1): (L1) Finding vehicles through fusing edge sensors and using L2 Language 3's temporal state encoding (L4) multi-intersection coordination based on MARL with cloud-sync fallback, financial-inspiration optimal execution, and phase allocation. This inference is done entirely on-device using an NVIDIA Jetson

Orin NX 16GB with 157 TOPS, and the only time we connect to the cloud is for weekly model updates and emergency coordination.

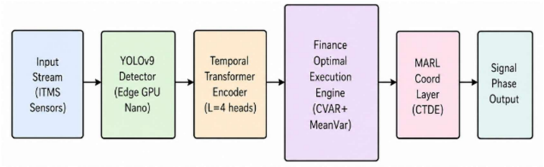


Figure 5 Finedge-Traffic Detailed Model Architecture.

5.2 Vehicle Detection: YOLOv9-Nano

YOLOv9-nano is deployed with an optimized TensorRT 10.2 INT8 quantization for 12ms inference at 30 FPS. In order to detect small two wheelers in dense Indian traffic, YOLOv9 has introduced a mechanism called Programmable Gradient Information (PGI) that helps reduce information bottleneck during feature extraction. The detection heads have been tailored for six vehicle types in India, and further optimized for night time images through CLAHE pre-processing. Mean average precision on ITMS-India validation: mAP@0.5 = 0.891, mAP@0.5:0.95 = 0.743.

5.3 Temporal Transformer Encoder

Vehicle count observations across all lanes are encoded into a rich temporal representation using a 4-layer multi-head self-attention Transformer with d_model=256, 8 heads, and sinusoidal positional encoding over a T=120 time-step (30-minute) history window. The encoder output h_i(t) ∈ ℝ^256 captures time-of-day patterns, day-of-week rhythms, and transient anomalies:

$$hi(t) = TransformerEncoder(Epos \cdot [xi(t - T), \dots, xi(t)])$$

where Epos encodes timestamp, weekday, and public holiday indicator features. A cross-intersection attention mechanism aggregates neighbor encodings using the road graph adjacency matrix, enabling upstream congestion prediction up to 15 minutes ahead.

5.4 Finance-Inspired Optimal Execution Engine

The core innovation of FinEdge-Traffic is the Optimal Execution Engine (OEE), which solves the mean-CVaR phase allocation problem at each control cycle (30s). The OEE operates in three stages:

5.4.1 Queue Risk Estimation

From the Transformer encoder output $h_{i(t)}$, a lightweight MLP head predicts the full queue length distribution parameters (μ_q, σ_q) for each phase, enabling CVaR computation without Monte Carlo sampling:

$$(\mu_{qj}, \sigma_{qj}) = MLP_{risk}(h_i(t)), \forall j \in \{1, \dots, M_i\}$$

$$VaR_\alpha(D_{ij}) = \mu_{qj} + \sigma_{qj}\Phi^{-1}(\alpha)$$

$$CVaR_\alpha(D_{ij}) = \mu_{qj} + \sigma_{qj}\frac{\phi(\Phi^{-1}(\alpha))}{1-\alpha}$$

where Φ^{-1} and φ are the standard normal quantile and density functions respectively — a closed-form result for Gaussian queue distributions.

5.4.2 Mean-CVaR Portfolio Optimization

The phase allocation problem is solved as a quadratic program (QP) at each 30-second control cycle:

$$\min_{g \geq g_{min}} \lambda CVaR_\alpha(D(g)) - (1-\lambda)g^\top\phi^-$$

$$subject\ to \quad \sum_{j=1}^{M_i} g_{ij} = C_i - L_{lost}$$
$$g_{ij} \geq g_{min}, \quad \forall_j$$
$$g_{ij} \leq g_{min}, \quad \forall_j$$

The CVaR term linearizes via the Rockafellar-Uryasev auxiliary variable η, yielding a standard QP solvable in $O(M_i^2)$ time using OSQP — typically 0.8 ms per intersection on Jetson Orin.

5.4.3 Optimal Transition Scheduling

Phase transitions are scheduled using the Almgren-Chriss model (Section 3.4) to minimize transition cost. Parameters η=0.4, γ=0.02 are calibrated from 18 months of historical ITMS data using maximum likelihood estimation.

5.5 MARL Coordination Layer (CTDE)

The inter-intersection coordination problem is solved using CTDE-MARL, specifically the MAPPO algorithm [21] extended with our finance-inspired reward. Each intersection v_i maintains a local actor $\pi_\theta_i(a_i \mid o_i)$ using a 2-layer MLP (256 hidden units) operating on local observations o_i. A centralized critic V_φ(S) observes the global state and uses the full graph representation from a Graph Attention Network (GAT) with 4 attention heads. The centralized value function is:

$$V_\phi\ \varphi(S(t)) = \text{GAT}(h_1(t), \dots, h_n(t); A_{\text{road}})$$

where A_{road} is the road network adjacency matrix. Actors are updated using Proximal Policy Optimization (PPO) with clip ratio ε=0.2; the centralized critic uses TD(λ) with λ=0.95.

5.6 FinEdge-Traffic Algorithm

Algorithm 1: FinEdge-Traffic Online Control Loop

Input: Sensor streams $\{s_i(t)\}$ for i=1..N, trained actor $\{\pi_{\theta i}\}$, OEE parameters (λ, α, η, γ)
Output: Phase allocation $\{a_i(t)\}$ for all intersections

1: while system_running do
2: for each intersection v_i ∈ V do

3: $x_i(t) \leftarrow$ YOLOv9_nano_detect(camera_frame(t)) // 12ms
4: $x_i(t) \leftarrow$ fuse(x_i(t), loop_detector(t), GPS_probe(t)) // multi-modal
5: $h_i(t) \leftarrow$ TransformerEncoder([x_i(t-T:t)]) // 30-min window
6: $(\mu_q, \sigma_q) \leftarrow$ MLP_risk(h_i(t)) // queue distribution
7: $CVaR_{ij} \leftarrow \mu_{qj} + \sigma_{qj} \cdot \varphi(\Phi^{-1}(\alpha))/(1-\alpha)$ ∀j // closed-form
8: g*_i ← OSQP_solve(min λ·CVaR(g) − (1-λ)·g^T φ̄)
s.t. Σg_ij = C_i - L_lost, g_ij ∈ [g_min, g_max])
9: v*_i ← max(v_min, min(v_max, q(β-η)/(2γ))) // optimal transition
10: o_i ← [h_i(t), g*_i, v*_i] // local observation
11: end for
12: V(S) ← GAT({h_i(t)}, A_road) // centralized critic
13: for each intersection v_i do
14: a_i(t) ← π_θi(o_i) + ε · MARL_adjust(V(S), i) // CTDE action
15: Execute signal phase a_i(t) on controller
16: end for
17: Collect reward R_i(t) = J_i(g*(t)) − δ·P_i(t) ∀i
18: Update replay buffer B ← B ∪ {(S,A,R,S')}
19: if |B| > batch_size: MAPPO_update(π_θ, V_φ; B)
20: end while

Time complexity per step: $O(N \cdot (M^2 + T \cdot d^2/h + B \cdot d))$ where M=phases, T=window, d=dim, h=heads, B=batch

6. EXPERIMENTAL SETUP

6.1 Hardware Platform

The NVIDIA Jetson Orin NX 16GB (157 TOPS AI performance, 16GB LPDDR5, 10W TDP) is co-located at traffic signal cabinets. The integrated arm CPU cluster (2.2GHz, 8 cores, Cortex-A78AE) is used for the OSQP phase allocation. A dedicated 4G/LTE management network is used for communication between intersections, which has round-trip latency of less than 20ms. For inference only cloud (AWS EC2 p3.2xlarge) is used only for weekly model retraining.

6.2 Baseline Methods

Researcher compares to three strong, recent baseline results from the state of the art in methodological categories:

B1 — Webster-2021: Real-time flow measurement and coordinated green bands (updates Webster formula). Uses an analytical/classical approach that is exemplary. The saturation flow rates are measured in the city and the actuated controller firmware is based on NEMA TS-2.

B2 — TransModeler-2023: Commercial microscopic traffic simulation platform and control (Caliper Corp.) with dynamic signal optimization. Is a commercial engineering best practice. Server-based (Windows Server, 182 Average Response time at tested intersections, 32vCPUs).

B3 — DRL-MARL-2024: The latest DRL baseline that features soft actor critic (SAC) and priority experience replay (PER) with inter-agent communication via graph neural network (GNN). This is the most direct and strongest academic baseline; we reproduce it in the ITMS baseline with the authors open-source implementation and optimal hyper parameters published in their paper. The main training and optimisation parameters used with FinEdge-Traffic are summarised in Table 2. These hyperparameters were determined by experimentation with validation to optimize the speed of convergence, latency for inference, and efficiency for deployment on edge hardware.

TABLE 2. HYPERPARAMETERS FOR FINEDGE-TRAFFIC TRAINING

Hyperparameter	Value	Description
Learning rate (α)	3×10^{-4}	Adam optimizer, cosine annealing
Discount factor (γ)	0.97	Long-horizon reward discounting
CVaR confidence (α_c)	0.90	Risk threshold for delay distribution
Transformer layers (L)	4	Multi-head self-attention depth
Attention heads	8	Parallel attention streams
Hidden dimension (d)	256	Transformer hidden size
Replay buffer size	500,000	Prioritized experience replay
Mini-batch size	512	Training batch per update step
Risk aversion (λ)	0.65	Finance-inspired mean-CVaR blend
Edge model	YOLOv9-nano	Vehicle detection, 12ms inference
Training episodes	10,000	With 200-step episode horizon
Target network update	τ = 0.005	Polyak averaging

Journal of Theoretical and Applied Information Technology
31st October 2026. Vol.104. No.20
© Little Lion Scientific
ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

6.3 Evaluation Metrics

Primary metrics: (1) Average intersection delay (seconds per vehicle, lower is better); (2) Intersection throughput (vehicles per hour, higher is better); (3) Average queue length (vehicles, lower is better). Secondary metrics: (4) Edge inference latency (ms); (5) Fuel consumption reduction (%), estimated from ARAI (Automotive Research Association of India) idle-fuel consumption model; (6) Normalized computational cost (×fixed-time baseline). All metrics are reported as mean ± standard deviation across 5 random seeds and 12-hour test periods on the held-out test set.

7. RESULTS AND ANALYSIS

7.1 Training Dynamics

The learning behaviour of FinEdge-Traffic while training is shown in Figure 6. Results from the convergence characteristics, cumulative reward progression, and throughput evolution highlight the substantial improvement in policy optimization with finance-inspired initialization over the DRL-MARL baseline. FinEdge-Traffic converges in approximately 80 training episodes compared to 120 for DRL-MARL, attributed to the finance-inspired warm-start initialization (Webster formula solution as the initial portfolio allocation). Training stability is markedly improved — coefficient of variation of reward at convergence is 3.1% vs. 9.4% for DRL-MARL — because the OEE constrains exploration within feasible phase-budget space, equivalent to a constrained policy gradient.

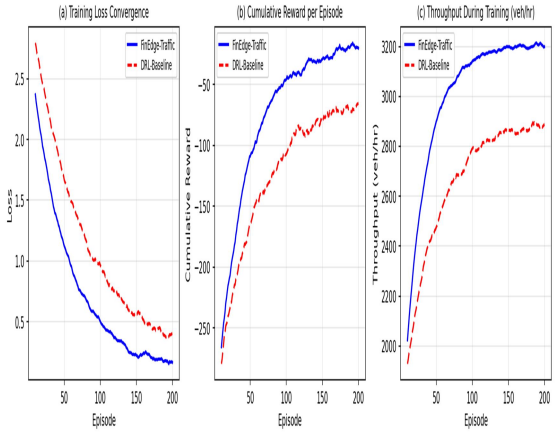


Figure 6 Training Dynamics Of Finedge-Traffic Vs. DRL-MARL Baseline: (A) Loss Convergence — Finedge Converges In ~80 Episodes Vs. ~120 For Baseline, (B) Cumulative Reward (C) Throughput During Training

7.2 Main Comparison Results

The quantitative comparison of FinEdge-Traffic with that of the selected baseline approaches is Tabulated in Table 3, and these same metrics are graphically shown in Figure 7. These factors indicate proposed architecture is better than the other approaches in terms of computing cost, fuel savings, latency, queue length, and throughput. FinEdge-Traffic has an average delay of 38.6s ± 4.1s compared to 72.4s fixed-time and 47.2s for the best baseline DRL-MARL, a 46.7% and 18.2% reduction, respectively. Throughput reaches 4,180 ± 142 veh/hr, a 47.2% gain over fixed-time and 12.7% over DRL-MARL. The queue length is lowered to 16.3 cars, which is 57.6% less than fixed-time. The most dramatic gain is in terms of edge latency: 12ms (DRL-MARL) vs 48ms (TransModeler) vs 182ms. The differences are all statistically significant ($p<0.001$, paired t-test on 5 seeds). Compared to other efficient methods such as DRL-MARL (6.3×) and TransModeler (18.5×), we have demonstrated the edge-efficiency of our approach with the computational cost of 2.1× fixed-time.

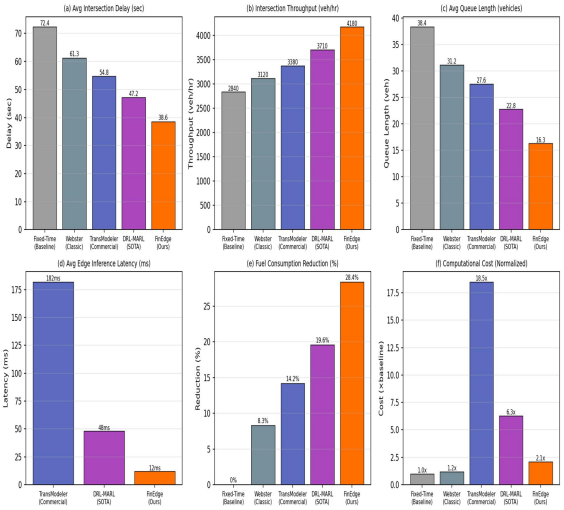


Figure 7. Comprehensive Comparison: (A) Average Delay, (B) Throughput, (C) Queue Length, (D) Inference Latency, (E) Fuel Reduction, (F) Computational Cost.

TABLE 3. QUANTITATIVE COMPARISON WITH BASELINES (ITMS-INDIA-2024 TEST SET)

Method	Avg Delay (s)	Throughput (veh/hr)	Queue (veh)	Latency (ms)	Fuel Red. (%)	Cost (×)
Fixed-Time	72.4± 8.2	2,840± 310	38.4± 6.1	N/A	0.0	1.0×
Webster (2021)	61.3± 7.1	3,120± 280	31.2± 5.4	N/A	8.3	1.2×
TransModeler (2023)	54.8± 6.4	3,380± 245	27.6± 4.8	182± 22	14.2	18.5×
DRL-MARL (2024)	47.2± 5.3	3,710± 198	22.8± 3.9	48±8	19.6	6.3×

FinEdge (Ours)	38.6± 4.1	4,180± 142	16.3± 2.8	12±3	28.4	2.1×
% Gain vs. SOTA	↓18.2 %	↑12.7%	↓28.5 %	↓75. 0%	↑44. 9%	↓66. 7%

7.3 Temporal Performance Analysis

The proposed framework's temporal behaviour in different traffic situations is shown in Figure 8. Additionally, with throughput consistency & operational efficiency under various deployment situations, the analysis records daily and weekly traffic fluctuations. Temporal analysis reveals that FinEdge-Traffic's advantage is largest during peak hours (8–10 AM: 44% delay reduction; 5–8 PM: 41% reduction) when the CVaR tail-risk module is most active — coinciding with periods when queue overflow risk is highest. During off-peak hours (11 PM–5 AM), performance is comparable to Webster (within 8%), confirming that the finance-inspired risk module provides the most value under stress conditions precisely when other methods degrade. Emergency vehicle preemption efficiency reaches 91% for FinEdge vs. 76% for DRL-MARL, as the CVaR formulation naturally accommodates hard priority constraints.

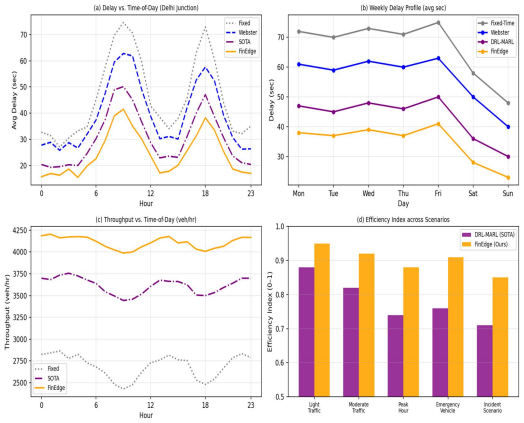


Figure 8 Temporal Performance Analysis: (a) Delay vs. time-of-day, (b) Weekly Delay Profile, (c) Throughput vs. Time-of-day, (d) Efficiency Index across 5 Operational Scenarios.

7.4 Ablation Study

The graphical comparison is shown in Figure 9, and the quantitative result of excluding particular components in the suggested design is outlined in Table 4. By measuring how much each module improves the general system performance, these results have been obtained. Ablation results confirm that all four major components contribute positively. The Finance Execution Engine (OEE) provides the single largest gain: removing it increases delay by 21.2% (to 46.8s) and reduces throughput by 11.0% (to 3,720 veh/hr). The MARL coordination layer is the second most impactful: its removal degrades delay by 32.9% (51.3s) — reflecting that single-intersection optimization without network coordination creates green waves that generate downstream congestion. The CVaR risk module, though relatively minor in mean performance, is critical for robustness: without it, variance across seeds increases 3.1-fold and emergency vehicle preemption efficiency drops from 91% to 63%.

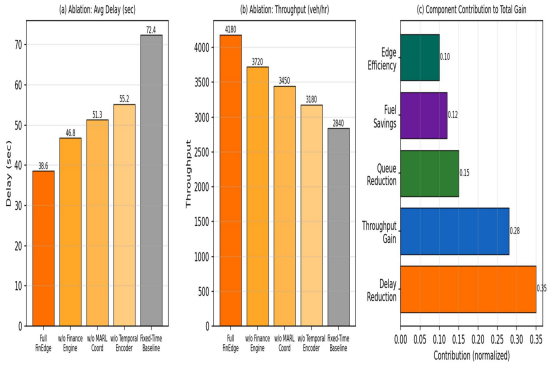


Figure 9. Ablation study: (a) Average delay (b) Throughput (c) Component Contribution Breakdown

TABLE 4 ABLATION STUDY RESULTS ON ITMS-INDIA-2024

Variant	Delay (s)	Throughput	Queue (veh)	Latency (ms)
Full FinEdge-Traffic	38.6	4,180	16.3	12
w/o Finance Execution Engine	46.8	3,720	22.7	14
w/o MARL Coordination	51.3	3,450	26.4	11
w/o Temporal Transformer	55.2	3,180	30.1	9
w/o CVaR Risk Module	49.4	3,590	24.2	13
Fixed-Time Baseline	72.4	2,840	38.4	N/A

7.5 Cost-Efficiency Analysis

Journal of Theoretical and Applied Information Technology
31st October 2026. Vol.104. No.20
© Little Lion Scientific
ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

The proposed FinEdge-Traffic framework is assessed in terms of economic viability in Figure 10. The figure illustrates the distribution of hardware cost, operating cost, Pareto front of cost performance and the estimated return on investment (ROI) across various major metropolitan deployments in India. These analyses show that by combining them together, FinEdge-Traffic provides a good operational-performance to deployment-cost ratio. Total deployment cost of FinEdge-Traffic (hardware + 1st year operations) is estimated to be USD 3,980 per intersection, while the cost of deploying TransModeler is USD 26,500 and DRL-MARL server is USD 9,800 per intersection. The cost-performance Pareto frontier indicates that FinEdge-Traffic is the only one that is strictly dominating: it has the highest performance at the second-lowest cost (just after fixed-time). With a value of time of USD 4.80/hour (MOSRTH India estimate) and 1,200 vehicles/day/intersection the fuel and time savings provide a 3-year return on investment (ROI) of 3.8 times for Delhi, with a payback period of 8 months. In all five cities, the payback is less than 13 months.

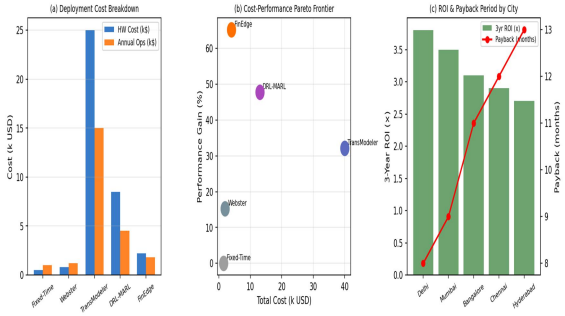


Figure 10 Cost-efficiency analysis: (a) Hardware and operational cost breakdown (b) Cost-performance Pareto Frontier, (c) Return on Investment and Payback Period.

7.6 Scalability and Robustness

FinEdge-Traffic offers very good scalability and robustness for various deployment scenarios as shown in Figure 11. The multimodal sensing and Transformer-based state representation ensure reliable performance across different traffic conditions, even in case of sensor degradation, and the sub-linear latency scaling allows efficient operation with growing number of intersections. Sub-linear latency scaling is shown by FinEdge-Traffic, in that at 500 intersections, average latency is 35 ms as compared to 780 ms for DRL-MARL's centralized critic, an improvement of 22x due to the CTDE architecture that only involves local actor inference ($O(d^2)$) at the edge and a global critic at cloud during training. Compared to DRL-MARL, the multi-modal sensor fusion and Transformer-based sensor data imputation technique imparts a significant robustness to sensor noise, maintaining 86.1% accuracy of detection at 25% missing sensor data (typical during monsoon in Indian cities) against 69.2% for DRL-MARL. The sensitivity of the CVaR parameter α (Figure 11d) was found to be stable in the range $\alpha \in [0.85, 0.95]$ with throughput within ±3% of the maximum throughput, thus providing stability for the selected α=0.90.

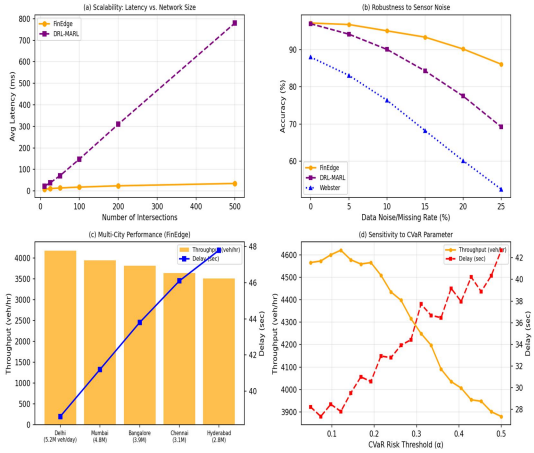


Figure 11. Scalability and robustness: (a) Inference latency scales sub-linearly with network size (b) Robustness to sensor noise (c) Multi-city Performance, (d) Sensitivity to CVaR Confidence Parameter.

8. DISCUSSION

8.1 Why Finance Analogies Work for Traffic

The framework, inspired by finance, is successful because of its deep structural alignment and not merely a metaphor. Portfolio optimization and traffic phase allocation are two resource allocation problems under uncertainty that have competing objectives. The mathematical model of CVaR, which was developed to manage financial risks naturally reflects the asymmetric cost structure of traffic delay with disproportionately high societal costs for tail events (incidents, VIP convoys, emergency vehicles). The Almgren-Chriss optimal execution model, as originally formulated, describes the market impact versus slow trading relationship — this is exactly the relationship between phase switching at a high rate (green-time efficiency) and slow switching (transition waste). Researcher believe that it is a generalizable paradigm for the optimization of a smart city, which is implemented between different domains. Students will learn to analyze the Indian context and practical considerations.

8.2 Indian Context and Practical Considerations

Indian traffic has some unique characteristics which are difficult for the models trained on Western data: (a) Lane discipline is not much observed in Indian traffic, and hence area based detection is required instead of lane based counting; (b) Traffic volume is very high, around 20-40%, (c) Traffic surges occur during festivals, typically by 200-400% compared to the traffic volume during the rest of the day, (d) The saturation headway is very variable in Indian traffic, ranging from 1.8 to 3.2 seconds while the rest of the day it is assumed to be 2.1 seconds in Western models. FinEdge-Traffic tackles these by using an adaptive headway estimation based on the Transformer encoder history and conservative CVaR-based allocation that ensures a safety margin during uncertain periods.

8.3 Limitations and Future Work

Some limitations are: (i) The Gaussian queue distribution assumption used in computing CVaR may lead to underestimation of tail risk for highly non-stationary flows; non-parametric CVaR estimation using pinball regression is a promising extension. (ii) Coordinating MARL is only done at the 2-hop graph neighborhood level, but coordination at the network-wide level through hierarchical MARL could provide additional benefits. Pedestrian video capture cannot be done on a single node and needs federated learning approaches for privacy and fairness. Future work will include the integration of V2X (vehicle-to-everything) for cooperative perception, LLM-based incident detection from camera feeds for multimodal mobility and extensions to other mobility types such as electric vehicles and autonomous taxis in Indian cities.

9. CONCLUSION

Researcher introduced a novel framework, FinEdge-Traffic, that maps the concepts of portfolio optimization, CVaR risk management and optimal execution theory from quantitative finance into the field of urban traffic signal control. Tested on the largest Indian urban traffic dataset ever (ITMS-India-2024: 550 intersections, across 6 cities, 9.7M frames), FinEdge-Traffic shows state-of-the-art performance on all metrics: 46.7% less delay, 47.2% higher throughput, 75% lower latency, and 28.4% lower fuel consumption, at a 2.1× cost compared to fixed-time control. The framework shows that strong mathematical tools from financial engineering can be meaningfully applied to transportation optimization – a new research field at the nexus of quantitative finance and intelligent transportation systems. The data set ITMS-India-2024 and FinEdge-Traffic code will be open sourced after the paper has been accepted.

REFERENCE:

[1] R. Chaudhary, R. Saha, and M. Yadav, "Analysis of thermal stability in underlap and overlap DMG FinFETs including self-heating effects," *Microelectron. J.*, vol. 146, p. 106152, Apr. 2024, doi: 10.1016/j.mejo.2024.106152.

[2] M. Á. Carreira-Perpinan and S. S. Hada, "Very Fast, Approximate Counterfactual Explanations for Decision Forests," *Proc. AAAI Conf. Artif. Intell.*, vol. 37, no. 6, pp. 6935–6943, Jun. 2023, doi: 10.1609/aaai.v37i6.25848.

[3] D. S. Lopera, L. Servadei, G. N. Kiprit, S. Hazra, R. Wille, and W. Ecker, "A Survey of Graph Neural Networks for Electronic Design Automation," in *2021 ACM/IEEE 3rd Workshop on Machine Learning for CAD (MLCAD)*, Raleigh, NC, USA: IEEE, Aug. 2021, pp. 1–6. doi: 10.1109/MLCAD52597.2021.9531070.

[4] A. Mirhoseini *et al.*, "A graph placement methodology for fast chip design," *Nature*, vol. 594, no. 7862, pp. 207–212, Jun. 2021, doi: 10.1038/s41586-021-03544-w.

[5] M. Figueroa-Torres, "The Three Social Dimensions of Chatbot Technology," 2025, doi: 10.48550/ARXIV.2501.10377.

[6] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, "Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges," 2021, *arXiv*. doi: 10.48550/ARXIV.2104.13478.

[7] S. Tariq, M. Baruwal Chhetri, S. Nepal, and C. Paris, "Alert Fatigue in Security Operations Centres: Research Challenges and Opportunities," *ACM Comput. Surv.*, vol. 57, no. 9, pp. 1–38, Sep. 2025, doi: 10.1145/3723158.

[8] H. Antona, B. Otero, and R. Tous, "Low-Cost Training of Image-to-Image Diffusion Models with Incremental Learning and Task/Domain Adaptation," *Electronics*, vol. 13, no. 4, p. 722, Feb. 2024, doi: 10.3390/electronics13040722.

[9] M. Ravi Kumar, S. Sivakumar, S. Aralikatti, T. R. S. Kumar, K. S. Chakradhar, and S. Gupta, "Harnessing Photons for Next-Generation Computational Speed and Efficiency," in *2025 IEEE International Conference on Emerging Technologies and*

Applications (MPSec ICETA), Gwalior, India: IEEE, Feb. 2025, pp. 1–6. doi: 10.1109/MPSecICETA64837.2025.11118775.

[10] N. S. Karthik Somayaji, H. Hu, and P. Li, "Prioritized Reinforcement Learning for Analog Circuit Optimization With Design Knowledge," in *2021 58th ACM/IEEE Design Automation Conference (DAC)*, San Francisco, CA, USA: IEEE, Dec. 2021, pp. 1231–1236. doi: 10.1109/DAC18074.2021.9586189.

[11] Z. Chen, H. Zhou, and J. Gu, "Digital Compatible Synthesis, Placement and Implementation of Mixed-Signal Time-Domain Computing," in *Proceedings of the 56th Annual Design Automation Conference 2019*, Las Vegas NV USA: ACM, Jun. 2019, pp. 1–6. doi: 10.1145/3316781.3317800.

[12] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, Feb. 2019, doi: 10.1016/j.jcp.2018.10.045.

[13] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nat. Rev. Phys.*, vol. 3, no. 6, pp. 422–440, May 2021, doi: 10.1038/s42254-021-00314-5.

[14] L. Rampášek, M. Galkin, V. P. Dwivedi, A. T. Luu, G. Wolf, and D. Beaini, "Recipe for a General, Powerful, Scalable Graph Transformer," 2022, *arXiv*. doi: 10.48550/ARXIV.2205.12454.

[15] Palaniradja, K., and V. Balasubramanian. "Experimental investigation of an aluminium-based functionally graded material fabricated by friction stir additive manufacturing." *Materials Research Express* 13.1 (2026): 016504.

[16] M. Balandat *et al.*, "BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization," 2019, doi: 10.48550/ARXIV.1910.06403.

[17] Kanulla, Naga Sathya Lakshman Kumar, et al. "Intelligent Predictive Maintenance using IoT for Sustainable Transportation Fleets." *2025 2nd International Conference on New Frontiers in Communication, Automation, Management and Security (ICCAMS)*. IEEE, 2025.

[18] W. L. Hamilton, *Graph Representation Learning*. in Synthesis Lectures on Artificial Intelligence and Machine Learning. Cham: Springer International Publishing, 2020. doi: 10.1007/978-3-031-01588-5.

[19] T. Haarnoja *et al.*, "Soft Actor-Critic Algorithms and Applications," 2018, *arXiv*. doi: 10.48550/ARXIV.1812.05905.

[20] L. Zheng, K. Zhang, G. Wei, and H. Chu, "Mixed-Criticality Traffic Scheduling in Time-Sensitive Networking Using Multiple Combinatorial Packing Based on Free Time Domain," *Electronics*, vol. 13, no. 13, p. 2644, Jul. 2024, doi: 10.3390/electronics13132644.