entropy

MDPI

Article

A Freeway Dual-Ramp Cooperative Control Method Driven by Entropy and Risk Feedback

Ken Chen [1,2,3], Keyi Zhang [2,3], Yuwei Jiao [2,3,*], Haomin Li [2,3], Zilin Luo [2,3], Yufeng Tan [2,3] and Xiaowei Zhu [1]

[1] State Key Laboratory of Millimeter Waves, School of Information Science and Engineering, Southeast University, Nanjing 210096, China
[2] Southeast University–Sichuan Digital Joint Research Center for Intelligent Transportation, Southeast University and Sichuan Digital Transportation Technology Co., Ltd., Chengdu 610000, China
[3] Sichuan Shudao Intelligent Transportation Group Co., Ltd., Chengdu 610000, China
* Correspondence: jiaoyuwei@scsz.com

Abstract

Dual-ramp merging control requires coordination of safety risk, traffic operations, and command smoothness. We propose EA-QD-CVaR, which uses traffic-state entropy, risk entropy, and event proportions to update control objectives and local search parameters. Conditional value-at-risk (CVaR) evaluates the upper tail of time-to-collision (TTC) loss. In a SUMO dual-ramp merging scenario with demand scaled to real traffic conditions, each stochastic method undergoes 25 independent search trials. The main joint-risk measure is based on the union of events with TTC below 7 s or with net headway below 3 s. Compared with Uniform, SA, Greedy-global, and MPC-CVaR, EA-QD-CVaR reduces mean TTC risk by 6.22–24.95% and mean joint risk by 1.54–8.65%, while increasing mean downstream arrivals by 0.82–2.91%. Relative to the single observations for the deterministic references Greedy-local, No-Control, and ALINEA, mean TTC risk decreases by 13.23%, 28.56%, and 31.79%, respectively. The corresponding reductions in mean joint risk are 5.28%, 10.15%, and 7.89%, and mean downstream arrivals increase by 0.94%, 1.69%, and 1.69%. Relative to fixed-weight QD-CVaR, mean joint risk decreases by 4.71% at the 3 s net-headway threshold. At the 1.5 s threshold used for control feedback, the decrease is only 0.15%. Control variation F_u increases. The method achieves lower mean TTC and joint risk and slightly higher mean downstream arrivals. It also supports the design and simulation-based evaluation of coordinated dual-ramp metering and variable speed limits under a downstream freeway bottleneck.

Keywords: dual-ramp merging control; entropy and risk feedback; finite-budget search; variable speed limits; multi-objective trade-offs

Check for updates

Academic Editor: Xinyang Deng

Received: 22 September 2026
Revised: 7 November 2026
Accepted: 11 November 2026
Published: 19 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Highway traffic control is a complex systems optimization problem jointly driven by vehicle interactions, demand fluctuations, bottleneck capacity, and external disturbances. Car-following, lane changing, and ramp merging can amplify local speed disturbances into upstream congestion waves or attenuate them through traffic dispersion and control. Kinematic-wave theory, the cell transmission model, second-order traffic-flow models, and microscopic car-following models explain traffic evolution from the complementary perspectives of conservation, discrete propagation, speed dynamics, and vehicle interaction [1–7]. Cellular automata, three-phase traffic theory, and ring-road experiments further demonstrate nonlinear state transitions and self-organized congestion [8–13]. Highway

control must therefore coordinate safety risk, operating efficiency, and control stability under time-varying conditions rather than optimize travel time or throughput alone.

Existing highway-control approaches primarily rely on feedback regulation, predictive optimization, and data-driven learning. Feedback methods such as ALINEA update ramp-metering rates from occupancy or flow deviations and are attractive because of their simple structure and online applicability [14,15]. Model predictive control (MPC) can coordinate ramp metering and variable speed limits (VSLs) over a finite horizon while explicitly handling density, queue, and actuator constraints [16–22]. Reinforcement learning and hybrid control have further extended decision making in complex traffic environments [23–28]. These methods have clear theoretical and engineering value, but their performance in multi-objective and multi-constraint settings often depends on predefined scalarization weights, predictive models, or training distributions. When traffic moves from stable operation to congestion formation or a risk disturbance, fixed weights cannot directly express changing priorities between efficiency and safety, while expected loss or mean risk can underrepresent rare but severe outcomes.

An informative representation of safety is a prerequisite for adaptive control. Time to collision (TTC), headway, relative speed, required deceleration, and trajectory-conflict measures are widely used for proactive safety analysis [29–36]. These surrogate measures characterize risk exposure through relative vehicle motion, following gaps, and conflict intensity. The same mean TTC can arise from a compact distribution or a long-tailed distribution containing very low-TTC observations. Likewise, similar mean speeds and densities can correspond to different lane occupancies, lane-changing competition, and spacing dispersion. Means, variances, and threshold counts consequently provide an incomplete description of how risk is distributed and migrates across vehicles, lanes, and time windows.

Shannon entropy supplies a unified measure of this distributional information. For a discrete traffic-state distribution, entropy quantifies the dispersion of probability mass among state categories. Permutation entropy, approximate entropy, multiscale entropy, and complexity–entropy methods extend this idea to temporal ordering and scale-dependent structure [37–43]. Information entropy has accordingly been used to analyze speed distributions, traffic-flow complexity, congestion labels, and scenario complexity [44–49]. In most such studies, however, entropy remains a state descriptor, prediction input, or post hoc complexity metric rather than an endogenous feedback signal that modifies the optimization objective. Moreover, low risk entropy does not necessarily imply low risk: observations concentrated in a hazardous category can have low entropy but high absolute risk. Entropy therefore provides structural information about the risk distribution, whereas risk probability and tail loss preserve the safety semantics. These quantities are complementary rather than interchangeable.

Quantum-inspired optimization provides a mathematical mapping among candidate strategies, objective costs, and probability-based search. Quantum annealing, adiabatic optimization, the quantum approximate optimization algorithm, and variational quantum algorithms commonly encode objectives and constraints in a Hamiltonian and redistribute probability mass toward low-energy regions using parameterized operations [50–57]. We adopt this computational structure to represent candidate control sequences as strategy states and map composite costs to energies. The Hamiltonian defines the energy ordering of candidate strategies. Wang et al. linked an objective function, potential energy, and candidate-solution probability in a computable optimization process, while Jiao and Wang used a multiscale quantum harmonic-oscillator mechanism to describe probability sampling and migration across search regions [58,59]. These studies provide a direct methodological basis for the energy-landscape and probability-search mapping used here.

Quantum and quantum-inspired transport studies have considered traffic assignment, signal timing, and trajectory deconfliction using QUBO, quantum annealing, or parameterized algorithms [60–63]. Current evidence nevertheless remains concentrated on small or synthetic problems, with limited operational evaluation and few standardized benchmarks [64,65]. We examine whether feedback based on traffic-distribution information reduces risk using the same prediction model and candidate-evaluation budget and how it affects arrivals and command variation. Using synthetic dual-ramp control, we compare complete feedback, fixed parameters, feedback ablations, and classical search and control methods to quantify these trade-offs.

Recent studies have advanced adaptive regulation and risk optimization in coordinated traffic control. Yue et al. dynamically allocate ramp-metering and VSL control weights using pressure feedback from the mainline and on-ramp [66]. Zhang et al. incorporate CVaR into perimeter control with boundary queues to characterize congestion risk arising from demand and model uncertainty [67]. Salloum et al. decompose quadratic unconstrained binary optimization (QUBO) problems and use quantum annealing for vehicle route assignment [68]. These studies address control-flow allocation, macroscopic congestion risk, and route optimization, respectively. Here, we examine how distributional feedback and finite-budget search jointly affect safety, traffic output, and command smoothness in coupled dual-ramp control. Traffic-state entropy, risk entropy, and event proportions jointly update objective weights and search parameters, while CVaR characterizes upper-tail TTC loss. Quantum-inspired circuit proposals and MH state updates organize rolling search. Weight scans, feedback ablations, and search comparisons assess this closed-loop design.

We develop EA-QD-CVaR to incorporate entropy and risk feedback into joint ramp metering and variable speed limit control. Traffic-state entropy H_T and risk entropy H_R provide category-distribution information, while the risk-event proportion p_R contributes to safety-weight updates. CVaR evaluates upper-tail loss within the safety cost. A local circuit generates candidate proposals, and a Metropolis–Hastings (MH) acceptance rule [69] updates the search position using energy differences and proposal ratios. Once the evaluation budget is exhausted, the controller selects the lowest-cost evaluated feasible sequence, executes its first action, and replans in the next control cycle.

The contributions to cooperative control of two freeway on-ramps are threefold:

1. We provide evidence of the trade-offs among risk, traffic efficiency, and command smoothness in coupled dual-ramp control. Fixed-safety-weight scans, feedback ablations, and comparisons of 13 methods quantify the effects of adaptive feedback under alternating entrance demands and a shared downstream bottleneck. The outcomes include TTC events, joint risk, downstream arrivals, and command variation, revealing the measured relationship between risk reduction and greater control activity. The main comparison comprises 25 independent search trials for each of ten stochastic methods and single observations for three deterministic references.
2. We design an entropy-and-risk-feedback mechanism for joint dual-ramp decisions. Traffic-state entropy H_T, risk entropy H_R, and the risk-event proportion p_R jointly update objective weights and local search parameters. CVaR evaluates upper-tail safety loss. This mechanism coordinates both VSL and ramp-metering commands as traffic distributions and risk exposure change.
3. We develop local probability search and rolling control under a finite evaluation budget. Nine-state local candidate blocks and circuit proposals organize high-dimensional control-sequence search within 64 complete-sequence evaluations per cycle. MH transitions update subsequent search centers, and the controller executes the first action of the lowest-cost evaluated sequence. Fixed-local-kernel convergence analysis and

separate 16-step diagnostics from five independent trials provide theoretical and experimental support for the local probability search used in rolling dual-ramp control.

Section 2 presents the traffic model, feedback objectives, and local search. Section 3 analyzes fixed-local-kernel convergence and presents the 16-step diagnostic. Section 4 reports the comparison of 25 independent search trials in the dual-ramp scenario. Section 5 discusses the control trade-offs and applicability of the proposed method and presents the conclusions.

2. Adaptive Control Driven by Entropy and Risk Feedback

2.1. Method Overview

At each control cycle, EA-QD-CVaR updates control preferences using traffic-state entropy H_T, risk entropy H_R, and the event proportion p_R. It jointly optimizes ramp metering and variable speed limits. The composite objective includes efficiency, safety, command smoothness, and operational protection, with CVaR evaluating the upper tail of TTC loss.

Figure 1 links observation, feedback, candidate search, and rolling execution. A local four-qubit circuit generates proposals for nine VSL–ramp combinations, and MH acceptance updates the search position according to candidate energies. Once the evaluation budget is exhausted, the controller executes the first action of the lowest-cost evaluated sequence. New observations update the feedback in the next cycle.

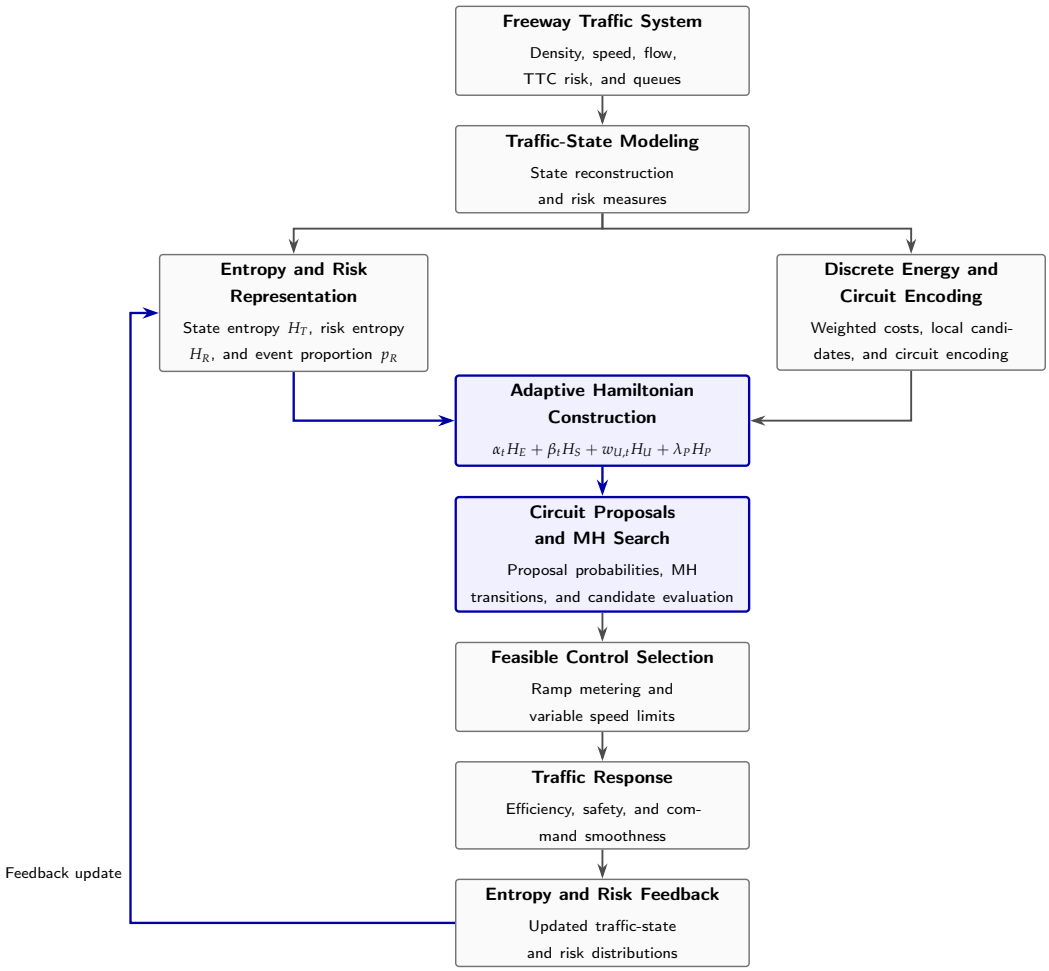


Figure 1. Entropy and risk feedback, candidate search, and rolling execution in cooperative dual-ramp control.

2.2. Discrete Energy and Probability Search

Let $\mathcal{Q}_t$ denote the set of evaluated feasible complete sequences in control cycle t, and let $E_t(\mathbf{z})$ denote the composite cost of candidate $\mathbf{z}$. The selected control sequence is

$$\mathbf{z}_t^* = \arg\min_{\mathbf{z}\in\mathcal{Q}_t} E_t(\mathbf{z}). \tag{1}$$

Feedback updates the weights so that candidate energy rankings respond to current traffic observations. Each local proposal changes one VSL–ramp variable pair at one prediction position, holding all other sequence components fixed. Four qubits represent the two three-level variables, yielding nine valid combinations.

Circuit probabilities and candidate energies jointly determine MH transitions. The resulting sequence becomes the center of the next local block. Complete sequences use classical data structures, and a traffic predictor computes their costs. Rolling control executes the first action selected by Equation (1). Section 3 analyzes the stationary distribution and convergence of fixed local kernels.

2.3. Freeway Traffic Control and Risk Modeling

Consider a freeway control region comprising the mainline, a merging area, and entrance ramps. With t denoting the control cycle, the traffic state is

$$\mathbf{x}_t = (\mathcal{X}_t, \boldsymbol{\ell}_t, \mathbf{u}_{t-1}), \tag{2}$$

where $\mathcal{X}_t$ contains the vehicle positions, speeds, leader–follower relationships, and other operational information. The vector $\boldsymbol{\ell}_t$ describes the queues on the ramps and vehicles waiting to enter the network, and $\mathbf{u}_{t-1}$ is the command executed in the preceding cycle. For the synthetic SUMO dual-ramp system evaluated here, the control input is

$$\mathbf{u}_t = \left[v_{1,t}^{\mathrm{VSL}}, r_{1,t}, v_{2,t}^{\mathrm{VSL}}, r_{2,t}\right]^{\mathsf{T}}, \tag{3}$$

where $v_{a,t}^{\mathrm{VSL}}$ and $r_{a,t}$ are the variable speed limit and ramp release rate in control region a, respectively. The traffic state evolves according to the classical dynamical model

$$\mathbf{x}_{t+1} = F_\theta(\mathbf{x}_t, \mathbf{u}_t, \mathbf{d}_t), \tag{4}$$

where $\mathbf{d}_t$ represents traffic demand and other external disturbances.

F_θ describes the mainline vehicle motion, merging interactions, and ramp-queue evolution. Given demand estimates and a candidate control sequence, the predictor recursively generates a state trajectory. It determines vehicle accumulation, speed, downstream arrivals, and following risk at each stage.

Over a prediction horizon of N_p steps, a candidate control sequence is

$$\mathbf{z}_t = \left(\mathbf{u}_{t|t}, \mathbf{u}_{t+1|t}, \ldots, \mathbf{u}_{t+N_p-1|t}\right). \tag{5}$$

The predictor evaluates each candidate sequence from the current state. The controller executes only the first action of the selected sequence, then updates predictions and decisions using new observations.

For vehicle i with valid leader information, let s_i be its net following gap used in risk calculation, v_i its speed, and $\Delta v_i = v_i - v_i^{\mathrm{lead}}$ its relative closing speed. Time to collision and net time headway are defined as

$$\mathrm{TTC}_i = \begin{cases} s_i/\Delta v_i, & \Delta v_i > \varepsilon_{\Delta v}, \\ +\infty, & \Delta v_i \le \varepsilon_{\Delta v}, \end{cases} \qquad h_i = \frac{s_i}{\max(v_i, \varepsilon_v)}. \tag{6}$$

The gap is adjusted according to the simulator interface's minimum-gap convention. The constants $\varepsilon_{\Delta v}$ and ε_v prevent numerical instability. TTC measures proximity to a conflict under relative motion, whereas net headway measures the current following interval.

We use $T_{\mathrm{TTC}} = 7$ s to identify potential conflicts early in the specified freeway-merging scenario, rather than as a universal severe-conflict threshold. Relevant engineering considerations include truck acceleration capabilities and acceleration-lane requirements at metered on-ramps [70], as well as merging lane-change conflicts between cars and trucks [71]. These studies inform the selection of risk-identification timescales in relation to vehicle characteristics and merging interactions.

A freeway-merging study reported critical forward and backward TTC values of 5 s and 6 s, respectively, for merging-vehicle gap acceptance or rejection [72]. In freeway safety analysis with varying truck proportions, Zhang et al. used 6 s as the upper threshold for general traffic conflicts [73]. They classified 2 s $<$ TTC $<$ 6 s as general conflicts and TTC $<$ 2 s as severe conflicts. Heavy-vehicle forward-collision-warning tests by the US National Highway Traffic Safety Administration (NHTSA) reported mean warning-trigger TTC values of 6.6–7.0 s across several scenarios. The prescribed range for those tests was 5.5–7.5 s [74]. These studies describe timescales for merging gap acceptance, conflict classification, and collision warning, respectively. Based on this scenario-specific evidence, we use 7 s to characterize potential merging-risk exposure and consistently calculate TTC-risk proportions and normalized TTC loss.

Let $\mathcal{P}_t$ be the current set of valid leader–follower pairs, with $n_t^{\mathrm{pair}} = |\mathcal{P}_t|$. The TTC-risk proportion is

$$p_{\mathrm{TTC}}(t) = \frac{1}{n_t^{\mathrm{pair}}} \sum_{i \in \mathcal{P}_t} \mathbb{I}\{\mathrm{TTC}_i < T_{\mathrm{TTC}}\}.$$

The joint-risk proportion used for control feedback is

$$p_R(t) = \frac{1}{n_t^{\mathrm{pair}}} \sum_{i \in \mathcal{P}_t} \mathbb{I}\Big\{\mathrm{TTC}_i < T_{\mathrm{TTC}} \ \vee \ h_i < T_h^{\mathrm{fb}}\Big\}, \tag{7}$$

where $\mathbb{I}\{\cdot\}$ is the indicator function and T_h^{fb} is the feedback threshold for net headway. A pair satisfying both conditions is counted only once. The corresponding risk proportion is set to zero when no valid pairs are available. Risk proportions over a stage or evaluation window are obtained by equally averaging the proportions across observation frames.

Offline joint-risk evaluation uses the same union of events, with a separately specified evaluation threshold T_h^{eval}:

$$p_J(t) = \frac{1}{n_t^{\mathrm{pair}}} \sum_{i \in \mathcal{P}_t} \mathbb{I}\Big\{\mathrm{TTC}_i < T_{\mathrm{TTC}} \ \vee \ h_i < T_h^{\mathrm{eval}}\Big\}.$$

To characterize both risk-event frequency and tail severity, we define the TTC loss as

$$L_i = \max\left\{0, \, 1 - \frac{\mathrm{TTC}_i}{T_{\mathrm{TTC}}}\right\}.$$

Let $L_{t,\ell}(\mathbf{z})$ denote the empirical loss distribution generated by candidate sequence $\mathbf{z}$ at prediction stage ℓ. All valid pairs are retained, including zero-loss samples. We use upper-tail conditional value-at-risk [75]

$$\mathrm{CVaR}^{\uparrow}_{\alpha_R}(L_{t,\ell}) = \frac{1}{1-\alpha_R}\int_{\alpha_R}^{1} F^{-1}_{L_{t,\ell}}(q)\,\mathrm{d}q, \tag{8}$$

where $F^{-1}_{L_{t,\ell}}$ is the empirical quantile function and α_R is the tail confidence level.

The safety cost is defined as

$$J_S(\mathbf{z};t) = \frac{c_S}{s_R}\sum_{\ell=0}^{N_p-1}\delta^{\ell}\left[p_{\mathrm{TTC},t,\ell}(\mathbf{z}) + \lambda_R\,\mathrm{CVaR}^{\uparrow}_{\alpha_R}(L_{t,\ell}(\mathbf{z}))\right], \tag{9}$$

where δ is the temporal discount factor, λ_R adjusts the relative contribution of tail loss, and c_S and s_R are fixed scaling coefficients.

The efficiency cost combines system vehicle accumulation, speed loss, and the benefit of downstream arrivals:

$$J_E(\mathbf{z};t) = \sum_{\ell=0}^{N_p-1}\delta^{\ell}\left[\frac{n^{\mathrm{sys}}_{t,\ell}}{n_{\mathrm{ref}}} + a_v\left(1-\frac{\bar{v}_{t,\ell}}{v_{\mathrm{ref}}}\right) - a_A\frac{A_{t,\ell}}{A_{\mathrm{ref}}}\right]. \tag{10}$$

Here, $n^{\mathrm{sys}}_{t,\ell}$, $\bar{v}_{t,\ell}$, and $A_{t,\ell}$ are the mean system vehicle count, mean mainline speed, and downstream arrivals at the prediction stage, respectively. The quantities n_{ref}, v_{ref}, and A_{ref} are fixed reference scales, and a_v and a_A are fixed combination coefficients.

The command-smoothness cost uses squared normalized command changes:

$$J_U(\mathbf{z};t) = c_U\sum_{\ell=0}^{N_p-1}\frac{1}{4}\sum_{j=1}^{4}\left(\frac{u_{t+\ell|t,j} - u_{t+\ell-1|t,j}}{D_j}\right)^2. \tag{11}$$

Here, D_j is the physical reference range for the corresponding control component. The initial change is calculated relative to the command actually executed in the preceding cycle. This quadratic smoothness cost sums equally over prediction stages. The outcome metric F_u instead uses the normalized absolute changes in executed commands.

These scalings convert different physical quantities into dimensionless costs. The contribution of each term to the composite objective depends jointly on its fixed scaling coefficients and objective weight. The fixed-weight control problem is written as

$$\min_{\mathbf{z}_t\in\mathcal{Z}_t}\; w_E J_E(\mathbf{z}_t;t) + w_S J_S(\mathbf{z}_t;t) + w_U J_U(\mathbf{z}_t;t) + \lambda_P J_P(\mathbf{z}_t;t), \tag{12}$$

where $\mathcal{Z}_t$ is the set of admissible discrete control sequences and J_P is a nonnegative operational-protection penalty relative to a hold-action reference. The reference $\mathbf{z}^{\mathrm{h}}_t = (\mathbf{u}_{t-1},\ldots,\mathbf{u}_{t-1})$ retains the previously executed command throughout the prediction horizon. Its prediction starts from the same current state as each candidate sequence. Let its stage-ℓ mean system vehicle count and arrivals be $n^{\mathrm{h}}_{t,\ell}$ and $A^{\mathrm{h}}_{t,\ell}$, respectively. Then,

$$\begin{aligned}
\Delta n_{t,\ell} &= n^{\mathrm{sys}}_{t,\ell} - n^{\mathrm{h}}_{t,\ell}, \qquad \Delta A_{t,\ell} = A^{\mathrm{h}}_{t,\ell} - A_{t,\ell},\\
J_P(\mathbf{z};t) &= \sum_{\ell=0}^{N_p-1}\delta^{\ell}\left[\left(\frac{[\Delta n_{t,\ell} - 0.05n^{\mathrm{h}}_{t,\ell} - 2]_+}{150}\right)^2 + \left(\frac{[\Delta A_{t,\ell} - 1]_+}{5}\right)^2\right].
\end{aligned} \tag{13}$$

Here, $[x]_+ = \max(0,x)$. The system count includes vehicles in the network, vehicles awaiting insertion in SUMO, and external waiting demand. Penalties apply when additional accumulation exceeds 5% of the reference count plus 2 vehicles or when arrivals decrease by more than 1 vehicle. The two excesses are normalized by 150 and 5 vehicles, respectively,

and squared. The composite objective uses $\lambda_P = 10$. The experimental settings list all reference scales, fixed coefficients, and thresholds.

2.4. Traffic-State Entropy and Risk Uncertainty

Traffic-state entropy H_T, risk entropy H_R, and the risk proportion p_R describe speed-category dispersion, the distribution of risk levels, and event frequency, respectively. Together, they provide feedback inputs to the adaptive Hamiltonian.

Traffic-state entropy is calculated from the speed distribution of mainline vehicles. Speeds are divided into K_T fixed bins, $\{\mathcal{B}_k^T\}_{k=1}^{K_T}$, and $\mathcal{V}_t$ denotes the set of observed mainline vehicles. The bin probabilities and normalized Shannon entropy are

$$p_k^T(t) = \frac{\sum_{i\in\mathcal{V}_t} \mathbb{I}\{v_i(t) \in \mathcal{B}_k^T\}}{|\mathcal{V}_t|}, \qquad H_T(t) = -\frac{1}{\log K_T} \sum_{k=1}^{K_T} p_k^T(t) \log p_k^T(t), \tag{14}$$

with the convention $0 \log 0 = 0$. Speed-bin boundaries are fixed in advance. A larger H_T indicates a more dispersed distribution across speed bins, reflecting heterogeneity in operating speeds.

Risk entropy is calculated from the TTC and net headway of valid leader–follower pairs. Prespecified thresholds map each measure to four risk levels. The higher of the two levels determines the category assigned to each pair. Let the category of pair i be $c_i(t) \in \{0,1,2,3\}$. The category probabilities and normalized risk entropy are

$$p_k^R(t) = \frac{\sum_{i\in\mathcal{P}_t} \mathbb{I}\{c_i(t) = k\}}{n_t^{\text{pair}}}, \qquad H_R(t) = -\frac{1}{\log K_R} \sum_{k=0}^{K_R-1} p_k^R(t) \log p_k^R(t), \tag{15}$$

where $K_R = 4$. The speed-bin boundaries and risk-category thresholds are specified in advance.

H_R measures dispersion across risk levels, whereas p_R in Equation (7) measures the observed proportion of risk events. For example, when all observations occupy the highest-risk category, H_R can be zero while p_R remains high. The safety weight therefore uses both the distributional information in risk entropy and the frequency information in the risk proportion. Traffic-state entropy describes the dispersion in operating conditions. Together with risk feedback, it adjusts the trade-offs among efficiency, safety, and smoothness.

At the start of each control cycle, $H_T(t)$, $H_R(t)$, and $p_R(t)$ are calculated from current observations. They update the objective weights defined in the next subsection. Within a control cycle, the weights and reference scales used for candidate evaluation remain fixed. After executing the first control action and obtaining new observations, the controller updates feedback for the next cycle.

2.5. Adaptive Hamiltonian and Local Probability Search

Let the three valid levels of each VSL or ramp control variable be $\mathcal{L} = \{0,1,2\}$. The corresponding local action set is

$$\mathcal{A} = \{(v_a^{\text{VSL}}, r_b) : a, b \in \mathcal{L}\}. \tag{16}$$

Each control variable is encoded using two qubits, with

$$|0\rangle \equiv |00\rangle, \qquad |1\rangle \equiv |01\rangle, \qquad |2\rangle \equiv |10\rangle. \tag{17}$$

The binary code $|11\rangle$ is reserved. A local VSL–ramp pair therefore uses four qubits and contains nine valid combinations. The logical state $|a,b\rangle_{\text{loc}}$ denotes VSL level a and ramp

level b. The VSL occupies the two low-order bits and the ramp the two high-order bits, giving the computational-basis index $a + 4b$.

Each local proposal updates one VSL–ramp block in a complete candidate control sequence while holding the other control variables and prediction positions fixed. The complete candidate sequence obtained from each local combination is denoted by $\mathbf{z}_t^{(a,b)}$, and the candidate set is $\mathcal{C}_t = \{\mathbf{z}_t^{(a,b)} : a, b \in \mathcal{L}\}$. The local block index is omitted from this notation. As the search position changes, the local candidate set and its conditional energies are updated together. The local circuit state is written as

$$|\psi_t(\tau)\rangle = \sum_{a,b\in\mathcal{L}} c_{ab}(t,\tau)|a,b\rangle_{\mathrm{loc}}, \qquad \sum_{a,b\in\mathcal{L}} |c_{ab}(t,\tau)|^2 = 1. \tag{18}$$

$|c_{ab}|^2$ is the proposal probability of a local combination. The local proposal circuit is coupled to rolling optimization of the complete control sequence, while traffic evolves according to Equation (4).

Table 1 maps the optimization representation to traffic-control quantities.

Table 1. Mapping between the optimization representation and traffic-control quantities.

Concept	Representation	Traffic-Control Quantity	Role
Local candidate state	$\|a,b\rangle_{\mathrm{loc}}$	VSL–ramp combination at one prediction position	One of nine local replacements in a complete sequence.
Proposal probability	$\|c_{ab}\|^2$	Proposal distribution over local controls	Circuit measurement probability before MH acceptance.
Local state space	$\mathcal{H}_{\mathrm{loc}}$	Nine valid combinations	Valid subspace of the four-qubit register.
Energy	$J(\mathbf{z};t)$	Efficiency, safety, and smoothness costs	Ranks complete candidate sequences.
Hamiltonian	$H(t)$	Weighted objectives and operational protection	Represents candidate energies as diagonal entries.
Efficiency term	$H_E(t)$	Accumulation, speed loss, and arrivals	Evaluates traffic operating costs.
Safety term	$H_S(t)$	TTC-event proportion and upper-tail loss CVaR	Evaluates risk frequency and tail severity.
Smoothness term	$H_U(t)$	Speed-limit and ramp-command changes	Penalizes abrupt changes between adjacent actions.
Operational protection	$H_P(t)$	Extra accumulation and lost arrivals relative to holding	Penalizes degradation beyond tolerances.
Circuit operation	$U_t(\boldsymbol{\theta})$	Amplitude updates for local combinations	Generates the candidate proposal distribution.
Executed plan	$\mathbf{z}_t^*$	Rolling control sequence	Uses the first action of the lowest-cost evaluated sequence.
Feedback input	H_T, H_R, p_R	Distribution entropies and risk-event proportion	Updates objective weights and associated search parameters.

Traffic costs define the diagonal component Hamiltonians:

$$H_j(t) = \sum_{\mathbf{z}\in\mathcal{Z}} J_j(\mathbf{z};t)|\mathbf{z}\rangle\langle\mathbf{z}|, \qquad j \in \{E,S,U,P\}. \tag{19}$$

Here, J_P is the nonnegative operational-protection penalty. The feedback-weighted objective Hamiltonian is

$$H_{\text{prob}}(t) = \alpha_t H_E(t) + \beta_t H_S(t) + w_{U,t} H_U(t) + \lambda_P H_P(t). \tag{20}$$

Equation (20) maps current traffic feedback to an energy ordering of candidate sequences. Each diagonal entry is the composite cost of the corresponding sequence. The weights follow these feedback rules:

$$\begin{aligned} \alpha_t &= \alpha_0 + k_T H_T(t), \\ \beta_t &= \text{clip}(\beta_0 + k_H H_R(t) + k_p p_R(t), \beta_{\min}, \beta_{\max}), \\ w_{U,t} &= w_{U,0} + k_{UT} H_T(t) + k_{UR} H_R(t), \end{aligned} \tag{21}$$

where $(\alpha_0, \beta_0, w_{U,0}) = (0.30, 0.45, 0.35)$ and $(k_T, k_H, k_p, k_{UT}, k_{UR}) = (0.45, 1.80, 4.50, 0.25, 0.25)$. Only β_t is explicitly clipped, with bounds $[0.45, 8.00]$; no additional clipping is applied to α_t or $w_{U,t}$. For $H_T, H_R, p_R \in [0, 1]$, the attainable weight ranges are $[0.30, 0.75]$, $[0.45, 6.75]$, and $[0.35, 0.85]$, respectively. These ranges follow jointly from the update equations and normalized inputs.

The signs of the feedback gains specify the control preferences associated with different observations. Higher traffic-state entropy indicates greater dispersion across speed categories. Positive k_T increases the efficiency weight, placing greater emphasis on system vehicle count, speed, and downstream arrivals in candidate ranking. Higher risk entropy indicates a more balanced coexistence of risk levels. Positive k_H increases the safety weight in this condition, whereas k_p makes the safety weight directly responsive to the risk-event proportion. Risk entropy is invariant to permutations of category probabilities; the event definitions and CVaR loss preserve risk severity. Positive k_{UT} and k_{UR} increase the command-change penalty as the two entropies rise, coordinating efficiency and safety preferences with command smoothness.

These prespecified empirical gains determine the adjustment magnitude of normalized feedback. Increasing H_T from 0 to 1 adds 0.45 to the efficiency weight. The maximum contributions of H_R and p_R to the safety weight are 1.80 and 4.50, respectively. Each entropy contributes at most 0.25 to the smoothness weight. Both entropies also update the circuit migration scale, which controls the migration angle of local proposal operations. Section 4.2 specifies the feedback channels retained in each ablation.

The composite energy of candidate $\mathbf{z}$ is

$$\mathcal{E}_t(\mathbf{z}) = \alpha_t J_E(\mathbf{z}; t) + \beta_t J_S(\mathbf{z}; t) + w_{U,t} J_U(\mathbf{z}; t) + \lambda_P J_P(\mathbf{z}; t). \tag{22}$$

Entropy adjusts candidate energy rankings through objective-weight updates. The implemented local circuit directly uses the two-qubit three-level embedding in Equation (17). It applies energy phases, adjacent-level mixing, and coordinated migration to the nine valid combinations.

For control variable $j \in \{v, r\}$, define the mixing generators between adjacent valid levels as

$$X_a^{(j)} = |a+1\rangle_j\langle a| + |a\rangle_j\langle a+1|, \qquad a \in \{0, 1\}. \tag{23}$$

Here, $j = v$ denotes the local VSL and $j = r$ the local ramp variable. The operator acts as the identity on the other control variable and as zero on the reserved code $|11\rangle_j$. $X_a^{(j)}$ mixes amplitudes between adjacent valid levels of the same variable. For local VSL–ramp combinations, define coordinated-migration generators as

$$T_a = |a+1, a+1\rangle_{\mathrm{loc}}\langle a, a|_{\mathrm{loc}} + |a, a\rangle_{\mathrm{loc}}\langle a+1, a+1|_{\mathrm{loc}}, \qquad a \in \{0, 1\}. \tag{24}$$

T_0 couples logical combinations $(0,0)$ and $(1,1)$, and T_1 couples $(1,1)$ and $(2,2)$. They act as zero on all other basis states. Coordinated migration acts on the VSL and ramp variables within the same local block. The mixing and coordinated-migration generators are, respectively,

$$H_{\mathrm{mix}} = \sum_{j \in \{v,r\}} \sum_{a=0}^{1} X_a^{(j)}, \qquad H_{\mathrm{tun}} = T_0 + T_1. \tag{25}$$

These generators are dimensionless Hermitian operators whose exponential maps define the rotation gates in the local proposal circuit.

Each proposal layer sequentially applies the elementary rotations for $a = 0$ and $a = 1$. The mixing and coordinated-migration gates in layer l are (with the rightmost operator acting first)

$$\begin{aligned} U_{\mathrm{mix}}^{(j)}(\theta_{j,l}) &= e^{-\mathrm{i}\theta_{j,l} X_1^{(j)}} e^{-\mathrm{i}\theta_{j,l} X_0^{(j)}}, \qquad j \in \{v, r\}, \\ U_{\mathrm{mix}}(\boldsymbol{\theta}_l) &= U_{\mathrm{mix}}^{(r)}(\theta_{r,l}) U_{\mathrm{mix}}^{(v)}(\theta_{v,l}), \\ U_{\mathrm{tun}}(\tau_l) &= e^{-\mathrm{i}\tau_l T_1} e^{-\mathrm{i}\tau_l T_0}, \qquad \boldsymbol{\theta}_l = (\theta_{v,l}, \theta_{r,l}). \end{aligned} \tag{26}$$

For any generator A above, within its coupled two-dimensional subspace, the diagonal entries of $e^{-\mathrm{i}\omega A}$ are $\cos\omega$, and the off-diagonal entries are $-\mathrm{i}\sin\omega$. The rotation angles use this two-dimensional phase convention. Adjacent generators generally do not commute. Mixing and coordinated migration follow the order in Equation (26).

For the current fixed local candidate set, let

$$\begin{aligned} E_{\min,t} &= \min_{a,b \in \mathcal{L}} \mathcal{E}_t(\mathbf{z}_t^{(a,b)}), \\ E_{\max,t} &= \max_{a,b \in \mathcal{L}} \mathcal{E}_t(\mathbf{z}_t^{(a,b)}). \end{aligned}$$

The normalized energies and local diagonal matrix used by the phase gate are

$$\begin{aligned} \widetilde{\mathcal{E}}_t(a,b) &= \frac{\mathcal{E}_t(\mathbf{z}_t^{(a,b)}) - E_{\min,t}}{\max(E_{\max,t} - E_{\min,t}, 10^{-12})}, \\ \widetilde{H}_{\mathrm{prob}}(t) &= \sum_{a,b \in \mathcal{L}} \widetilde{\mathcal{E}}_t(a,b) |a,b\rangle_{\mathrm{loc}}\langle a,b|_{\mathrm{loc}}. \end{aligned} \tag{27}$$

The matrix is zero on the other seven four-qubit basis states, so the corresponding phase gate acts as the identity on those states. One proposal layer evolves the state as

$$|\psi_{l+1}\rangle = U_{\mathrm{tun}}(\tau_l) U_{\mathrm{mix}}(\boldsymbol{\theta}_l) e^{-\mathrm{i}\varphi_l \widetilde{H}_{\mathrm{prob}}(t)} |\psi_l\rangle. \tag{28}$$

where φ_l, $\theta_{v,l}$, $\theta_{r,l}$, and τ_l are the angles for phase, VSL mixing, ramp mixing, and coordinated migration, respectively. The nine valid combinations are initialized in an equal-amplitude superposition, $|\psi_0\rangle = \frac{1}{3}\sum_{a,b \in \mathcal{L}} |a,b\rangle_{\mathrm{loc}}$. The phase, mixing, and migration gates all preserve this valid subspace. Each local proposal uses one ordered layer as defined above; l indexes circuit layers.

After normalization within the valid subspace, measurement yields proposal probabilities $q_t^{\mathrm{circ}}(\mathbf{z}_t^{(a,b)}) = |\langle a,b|_{\mathrm{loc}}\psi_1\rangle|^2$. The local candidate set for an MH update is $\mathcal{Z} = \mathcal{C}_t$, with $|\mathcal{Z}| = 9$. Mixing the measurement distribution with a uniform distribution gives a full-support proposal:

$$q_t(\mathbf{z}) = (1-\varepsilon_q)q_t^{\mathrm{circ}}(\mathbf{z}) + \frac{\varepsilon_q}{|\mathcal{Z}|}, \qquad \mathbf{z} \in \mathcal{Z}, \quad 0 < \varepsilon_q < 1. \tag{29}$$

The circuit proposes candidates, and candidate energies and the Hastings ratio determine their MH acceptance probabilities. The proposal distribution is reconstructed as the local candidate set and conditional energies change with the search position. Section 3 analyzes probability evolution for a fixed candidate set and transition kernel.

For each complete local block, multiple sets of circuit angles are generated, and their nine-state proposal probabilities are calculated. Parameters are selected using the probability-weighted lower-tail mean energy:

$$L_\eta^{\downarrow}(\mathcal{E}_t) = \frac{1}{\eta}\int_0^\eta F_{\mathcal{E}_t}^{-1}(q)\,\mathrm{d}q. \tag{30}$$

Here, $\eta \in (0,1)$ is the lower-tail probability mass. The quantile function weights the nine candidate energies by the corresponding circuit probabilities [76]. Parameter comparisons reuse the evaluated candidate energies, and the selected angles determine the proposal distribution for that local block. Lower-tail energy statistics select circuit parameters; traffic safety uses the upper-tail CVaR in Equation (8).

Given the current local state $\mathbf{z}_i$, an independent proposal $\mathbf{z}_j$ is drawn using Equation (29). Its acceptance probability is

$$a_t(i,j) = \min\left\{1, \exp[-\vartheta_t(\mathcal{E}_t(\mathbf{z}_j) - \mathcal{E}_t(\mathbf{z}_i))]\frac{q_t(\mathbf{z}_i)}{q_t(\mathbf{z}_j)}\right\}. \tag{31}$$

Here, $\vartheta_t > 0$ is the inverse temperature, and acceptance uses the composite energy $\mathcal{E}_t$. The normalized energies in Equation (27) are used by the phase gate. The Hastings ratio corrects for nonuniform proposals. Each transition proposes a local combination, calculates acceptance, and draws an acceptance random number. Rejection retains the current complete sequence. A self-proposal also counts as one transition. The conditional target distribution on a fixed local set is proportional to $\exp[-\vartheta_t\mathcal{E}_t(\mathbf{z})]$.

Each cycle initializes a hold-action sequence and performs its first evaluation. The first cycle uses preset commands; subsequent cycles repeat the preceding executed command throughout the horizon. Local blocks are visited cyclically in prediction-step order, with region 1 preceding region 2 at each step. The nine local combinations replace the corresponding two components, and complete-trajectory energies are queried. Circuit angles are then selected, followed by a preset number of MH transitions on the fixed local set. The resulting sequence becomes the center of the next block.

Each cycle limits the number of unique complete-sequence evaluations, including the hold-action reference. All evaluated candidates are cached; repeated queries read the cache without increasing the count. Search stops when the budget is exhausted. Evaluated candidates in an incomplete final block remain eligible for selection. The controller selects the lowest-energy feasible sequence $\mathbf{z}_t^*$ among all evaluated candidates and executes its first action. Algorithm 1 presents the complete EA-QD-CVaR procedure, from feedback updates through local exploration to rolling execution.

Algorithm 1 Rolling cooperative control with EA-QD-CVaR

Require: Initial traffic observations, previous command $\mathbf{u}_{t-1}$, prediction steps $N_p = 16$, and at most 64 new evaluations per cycle
Ensure: Speed-limit and ramp-metering commands executed each cycle
while control is ongoing **do**
  Compute H_T, H_R, p_R; update weights and κ_t; fix this cycle's prediction inputs
  $\mathbf{z} \leftarrow (\mathbf{u}_{t-1}, \ldots, \mathbf{u}_{t-1})$; initialize evaluation cache $\mathcal{Q}_t$
  Predict hold sequence $\mathbf{z}$; save its reference trajectory; count evaluation 1
  $b \leftarrow 0$
  while the stopping condition for new evaluation queries is not met **do**
    $s \leftarrow b \bmod (2N_p)$; select prediction step $\lfloor s/2 \rfloor$ and region $(s \bmod 2) + 1$
    Substitute nine VSL–ramp pairs in this block; hold all other components fixed
    Score full sequences; reuse cached energies; store new candidates in $\mathcal{Q}_t$
    if the 64-evaluation limit interrupts queries before the local set is complete **then**
      break
    end if
    Compare 8 angle sets; fix the full-support proposal $\mathbf{q}_t$ using Equation (29)
    for $k = 1, \ldots, 8$ **do**
      Propose a combination from $\mathbf{q}_t$; compute acceptance using Equation (31)
      Draw an acceptance variate; update $\mathbf{z}$ if accepted; otherwise, retain it
    end for
    $b \leftarrow b + 1$; use the transitioned sequence as the next block's center
  end while
  $\mathbf{z}_t^* \leftarrow \arg\min_{\mathbf{z} \in \mathcal{Q}_t} \mathcal{E}_t(\mathbf{z})$
  Execute the first four-dimensional action of $\mathbf{z}_t^*$ for 15 s; update observations
  $t \leftarrow t + 1$; start the next control cycle
end while

3. Stationarity and Finite-Step Convergence of MH Search

For the fixed local candidate sets in Section 2, candidate energies, proposal probabilities, and acceptance rules jointly determine the independence Metropolis–Hastings (MH) process [69]. Drawing on stationarity and ergodicity theory [77,78], we derive the stationary distribution and finite-step error for a fixed kernel. We then specify the conditions for distribution tracking when targets change across cycles. Section 3.5 presents separate 16-step diagnostics from five independent trials in the same dual-ramp scenario, examining nine-state local kernels encountered during search.

$$\begin{aligned} \text{Candidate energies and proposals} &\longrightarrow \text{Hastings acceptance} \\ &\longrightarrow \text{MH kernel} \\ &\longrightarrow \text{Stationarity and finite-step error.} \end{aligned} \tag{32}$$

3.1. Proposal–MH Process on a Fixed Candidate Set

For a fixed local block b in control cycle t, denote its candidate set by $\mathcal{C}_{t,b}$. Omitting the block index below, write $\mathcal{Z}_t = \{\mathbf{z}_1, \ldots, \mathbf{z}_N\}$. The inner index k counts MH transitions on this set. The current objective Hamiltonian is

$$H_t = \alpha_t H_{E,t} + \beta_t H_{S,t} + w_{U,t} H_{U,t} + \lambda_P H_{P,t}. \tag{33}$$

Here, α_t, β_t, and $w_{U,t}$ weight efficiency, safety, and command smoothness, respectively, following Equation (21); λ_P is the constraint-penalty coefficient. The candidate energy supplied to MH acceptance is

$$E_i(t) = \mathcal{E}_t^{\mathrm{MH}}(\mathbf{z}_i). \tag{34}$$

The energy scale and inverse temperature jointly determine the target probabilities.

Let $\widetilde{q}_{i,t}$ be the circuit measurement probability conditioned and normalized on the feasible candidate set. The full-support independent proposal is

$$q_{i,t} = (1-\varepsilon_q)\widetilde{q}_{i,t} + \frac{\varepsilon_q}{N}, \qquad 0 < \varepsilon_q < 1. \tag{35}$$

Thus, $q_{i,t} > 0$, and the probability of proposing $\mathbf{z}_j$ from any state is $q_{ij,t} = q_{j,t}$. Within a local block, $\mathcal{Z}_t$, $E_i(t)$, $\mathbf{q}_t$, and the inverse temperature ϑ_t remain fixed, defining a homogeneous transition kernel.

Within a fixed local block, proposal probabilities, candidate energies, and MH acceptance determine distribution evolution. The following matrix representation of the kernel and stationary distribution supports convergence analysis and numerical diagnostics.

3.2. Markov Representation of Candidate-Strategy Probabilities

Define the row probability vector after transition k as

$$\mathbf{P}_t^{(k)} = [p_{1,t}^{(k)}, \ldots, p_{N,t}^{(k)}], \qquad p_{i,t}^{(k)} = \Pr\{Z_t^{(k)} = \mathbf{z}_i\}. \tag{36}$$

Conditioning the circuit probabilities on the feasible subspace gives

$$\widetilde{q}_{i,t} = \frac{|\langle \mathbf{z}_i | U_t(\boldsymbol{\theta}_t^*) | \psi_{0,t} \rangle|^2}{\sum_{\ell=1}^{N} |\langle \mathbf{z}_\ell | U_t(\boldsymbol{\theta}_t^*) | \psi_{0,t} \rangle|^2}. \tag{37}$$

The denominator is assumed positive. Here, $\widetilde{q}_{i,t}$ proposes candidates, whereas $p_{i,t}^{(k)}$ describes the classical state probabilities after MH transitions.

For $\vartheta_t > 0$, the target distribution is

$$\pi_{i,t} = \frac{\exp[-\vartheta_t E_i(t)]}{\sum_{\ell=1}^{N} \exp[-\vartheta_t E_\ell(t)]}. \tag{38}$$

The independence-proposal acceptance probability and transition kernel are

$$a_{ij,t} = \min\left\{1, \exp[-\vartheta_t(E_j(t) - E_i(t))]\frac{q_{i,t}}{q_{j,t}}\right\}, \tag{39}$$

$$T_{ij,t} = q_{j,t} a_{ij,t}, \qquad i \neq j, \tag{40}$$

$$T_{ii,t} = 1 - \sum_{j \neq i} T_{ij,t}. \tag{41}$$

Diagonal entries include self-loops caused by self-proposals and rejections. The distribution evolves according to

$$\mathbf{P}_t^{(k+1)} = \mathbf{P}_t^{(k)} T_t. \tag{42}$$

Theorem 1 (Markov property of fixed-kernel MH transitions)**.** *When $\mathcal{Z}_t$, candidate energies, $\mathbf{q}_t$, and ϑ_t are fixed,*

$$\Pr(Z_t^{(k+1)} = \mathbf{z}_j \mid Z_t^{(k)}, \ldots, Z_t^{(0)}) = \Pr(Z_t^{(k+1)} = \mathbf{z}_j \mid Z_t^{(k)}). \tag{43}$$

The inner search is a finite-state homogeneous Markov chain.

Proof. Given $Z_t^{(k)} = \mathbf{z}_i$, the next step depends only on the fixed proposal, current and candidate energies, and a new independent uniform random number. Its conditional transition probability is $T_{ij,t}$, independently of earlier chain states. □

3.3. Stationary Distribution and Asymptotic Convergence for a Fixed Target

MH acceptance preserves the target through detailed balance. Irreducibility and aperiodicity on a finite state space further ensure convergence in distribution [77].

Theorem 2 (Existence and uniqueness of the stationary distribution)**.** *If $N < \infty$, all candidate energies are finite real numbers, $\vartheta_t \in (0, \infty)$, and $0 < \varepsilon_q < 1$ in Equation (35), then T_t is irreducible and aperiodic. Its unique stationary distribution is given by Equation (38).*

Proof. Full support gives $q_{i,t} \geq \varepsilon_q / N > 0$. Finite energies ensure $a_{ij,t} > 0$, so $T_{ij,t} > 0$ between any two distinct states. Detailed balance follows from

$$\begin{aligned} \pi_{i,t} T_{ij,t} &= \pi_{i,t} q_{j,t} \min\left\{1, \frac{\pi_{j,t} q_{i,t}}{\pi_{i,t} q_{j,t}}\right\} \\ &= \min\{\pi_{i,t} q_{j,t}, \pi_{j,t} q_{i,t}\} = \pi_{j,t} T_{ji,t}. \end{aligned} \tag{44}$$

Hence, $\boldsymbol{\pi}_t T_t = \boldsymbol{\pi}_t$. Since $T_{ii,t} \geq q_{i,t} > 0$, the chain is aperiodic. A finite irreducible chain has a unique stationary distribution. □

Theorem 3 (Asymptotic convergence for a fixed target)**.** *Under the conditions of Theorem 2, every initial distribution satisfies*

$$\mathbf{P}_t^{(k)} = \mathbf{P}_t^{(0)} T_t^k, \qquad \lim_{k \to \infty} \mathbf{P}_t^{(k)} = \boldsymbol{\pi}_t. \tag{45}$$

Proof. Equation (42) gives the matrix-power representation. The finite, irreducible, aperiodic matrix T_t is a primitive stochastic matrix. Its eigenvalue 1 is simple, and all other eigenvalues have modulus strictly below 1. Thus, $T_t^k \to \mathbf{1}\boldsymbol{\pi}_t$. Consequently, there exist $C_t < \infty$ and $0 < \rho_t < 1$ such that

$$\|\mathbf{P}_t^{(k)} - \boldsymbol{\pi}_t\|_{\mathrm{TV}} \leq C_t \rho_t^k. \quad \square \tag{46}$$

Energy and inverse temperature determine stationary probabilities. The proposal distribution affects the rate of approach through the transition kernel.

3.4. Finite-Step Error and Target Changes Across Cycles

Uniform ergodicity of independence MH can be characterized by a lower bound on the proposal relative to the target [78]. For a local set with fixed energies and proposal, define

$$m_t = \min_i \frac{q_{i,t}}{\pi_{i,t}}, \qquad 0 < m_t \leq 1. \tag{47}$$

For $i \neq j$, $T_{ij,t} = \min\{q_{j,t}, \pi_{j,t} q_{i,t} / \pi_{i,t}\} \geq m_t \pi_{j,t}$. Diagonal entries satisfy the same lower bound. For $m_t < 1$, $T_t = m_t \mathbf{1}\boldsymbol{\pi}_t + (1 - m_t) R_t$, where R_t is stochastic. Non-expansion of total variation under stochastic matrices gives

$$\|\mu T_t - \nu T_t\|_{\mathrm{TV}} \leq (1 - m_t) \|\mu - \nu\|_{\mathrm{TV}}. \tag{48}$$

Starting from distribution ν_t and applying K transitions gives

$$e_t(K) := \|\nu_t T_t^K - \boldsymbol{\pi}_t\|_{\mathrm{TV}} \leq (1 - m_t)^K \|\nu_t - \boldsymbol{\pi}_t\|_{\mathrm{TV}}. \tag{49}$$

When $m_t = 1$, $q_t = \boldsymbol{\pi}_t$, so one transition reaches the target distribution. Small m_t yields a slower convergence bound.

Cross-cycle analysis additionally assumes a common state space and continuous propagation of the chain distribution. Suppose all cycles share a finite state space and a uniform contraction coefficient $\rho < 1$. For the continuously propagated process $\mu_{t+1} = \mu_t T_t^K$, let $e_t = \|\mu_t - \boldsymbol{\pi}_t\|_{\mathrm{TV}}$, and $\Delta_t = \|\boldsymbol{\pi}_{t+1} - \boldsymbol{\pi}_t\|_{\mathrm{TV}}$. The triangle inequality and $\boldsymbol{\pi}_t T_t^K = \boldsymbol{\pi}_t$ give

$$\begin{aligned} e_{t+1} &\le \|\mu_t T_t^K - \boldsymbol{\pi}_t T_t^K\|_{\mathrm{TV}} + \Delta_t \\ &\le \rho^K e_t + \Delta_t. \end{aligned} \tag{50}$$

Recursion yields

$$e_t \le \rho^{Kt} e_0 + \sum_{j=0}^{t-1} \rho^{K(t-1-j)} \Delta_j. \tag{51}$$

If $\Delta_j \to 0$, then $e_t \to 0$. Furthermore, if $\boldsymbol{\pi}_t \to \boldsymbol{\pi}_\infty$,

$$\mu_t \to \boldsymbol{\pi}_\infty. \tag{52}$$

Equation (51) separates tracking error into decay of the initial discrepancy and accumulated target changes. It applies to continuously propagated chain distributions on a common state space with a uniform contraction coefficient. Rolling search reconstructs candidate sets block by block. The following diagnostic therefore calculates the distribution distance in Equation (49) from each block's actual fixed energies and proposal.

3.5. Finite-Step Distribution Diagnostics for Local MH Search

The 25-trial closed-loop main comparison in Section 4 uses eight MH iterations per complete local block. Five separate diagnostic trials increase this number to 16 to examine distributional changes with additional iterations. The road network, traffic and prediction random streams, 16-step horizon, 15 s control interval, and 64 complete-sequence evaluations per cycle remain unchanged. The eight circuit-angle sets and EA feedback rules are also retained. Each diagnostic trial contains 16 control cycles, with seven complete nine-state local blocks per cycle. Together, the five trials contain 80 cycles, 560 complete local blocks, 8960 actual MH transitions, and 5120 complete-sequence evaluations.

For each block, nine raw candidate energies, full-support proposal probabilities, the actual initial state, and 16 proposal/acceptance records reconstruct the target and 9×9 transition matrix. The inverse temperature is $\vartheta = 0.5$, and the uniform mixing mass is $\varepsilon_q = 10^{-3}$. All block proposal probabilities are positive, with a minimum $q_i = 1.1353 \times 10^{-4}$. Reconstructed-kernel residuals for row sums, detailed balance, and stationarity are all below 10^{-12}.

Starting from the actual initial state $i_{0,t,b}$, we calculate each fixed kernel's marginal distributions and total variation (TV) distance:

$$\begin{aligned} \mathbf{P}_{t,b}^{(0)} &= \mathbf{e}_{i_{0,t,b}}, \qquad \mathbf{P}_{t,b}^{(k)} = \mathbf{e}_{i_{0,t,b}} T_{t,b}^k, \\ d_{t,b}(k) &= \frac{1}{2} \sum_{i=1}^{9} |p_{i,t,b}^{(k)} - \pi_{i,t,b}|, \qquad k = 0, \ldots, 16. \end{aligned} \tag{53}$$

Each block is compared with its own target, and results are averaged equally through the block, cycle, and run levels. Figure 2 and Table 2 report the mean, sample SD, and range of the five run means. Independent trials are the statistical units; blocks and cycles are nested observations within runs. The results at Steps 8 and 16 are obtained by matrix propagation of the same fixed kernel.

Matrix propagation from the actual initial states of the five diagnostic trials gives a mean TV distance at step 8 of 0.088509 ± 0.018014. At this step, 132 local blocks satisfy

TV $\leq$ 0.01, corresponding to 23.57%. The step-16 results are 0.045174 ± 0.010453 and 51.79% (290/560), respectively. Mean TV distance decreases by another 48.96% relative to step 8, with all five trial means continuing to decline. The initial mean TV distance is 0.862212. The maximum step-16 block TV distance is 0.827885; this block has $m = 0.003146$, corresponding to a slower contraction bound. These results quantify finite-step distributional error and differences among local blocks. Additional iterations reduce the error, while most diagnostic blocks remain above the TV $\leq$ 0.01 reference level under the eight-step budget.

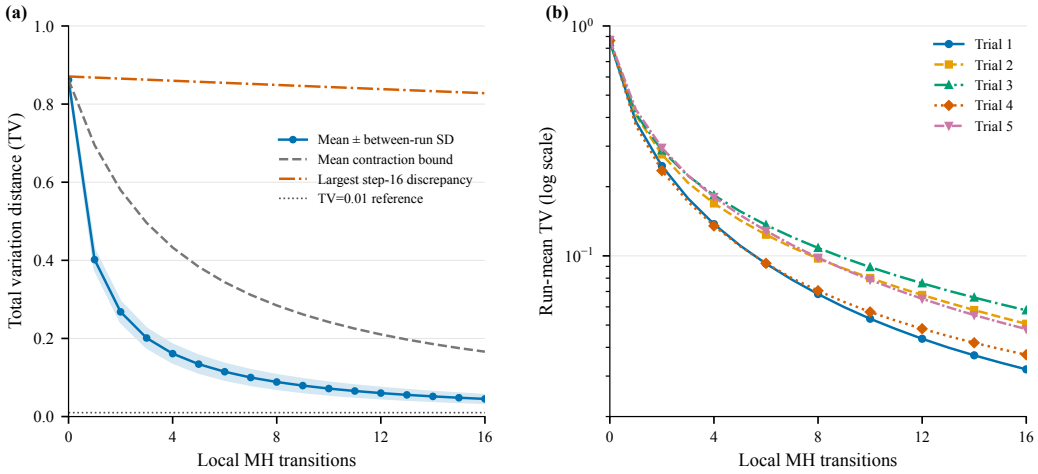


Figure 2. Sixteen -step local MH diagnostics from five independent trials in the dual-ramp scenario. (**a**) Mean and between-run sample SD of run means, the mean contraction bound in Equation (49), and the block with the largest step-16 discrepancy. TV= 0.01 is a descriptive reference. (**b**) Mean TV for each trial, aggregated across blocks and cycles, on a logarithmic vertical axis. Each trial contains 16 cycles and 112 complete blocks; $n = 5$.

Table 2. Finite-step TV distances in the 16-step diagnostic. Mean, SD, and range are based on the five run means. Maximum block TV and reference-level counts cover all 560 complete blocks.

Step k	Mean TV $\pm$ SD	Range of Run Means	Maximum Block TV	Blocks with TV $\leq$ 0.01
0	0.862212 ± 0.007976	$[0.851758, 0.870437]$	0.968288	0/560
1	0.401795 ± 0.025215	$[0.369059, 0.432485]$	0.900137	0/560
2	0.268047 ± 0.026438	$[0.234834, 0.295315]$	0.884166	0/560
4	0.161114 ± 0.023340	$[0.135025, 0.183568]$	0.859790	17/560
8	0.088509 ± 0.018014	$[0.068203, 0.108276]$	0.849021	132/560
16	0.045174 ± 0.010453	$[0.032091, 0.057952]$	0.827885	290/560

Theorem 3 establishes the unique stationary distribution and asymptotic convergence of a fixed local MH kernel under full-support proposals and finite candidate energies. Finite-step diagnostics characterize distributional error under a specified iteration budget. The controller uses eight MH transitions to update search centers for subsequent blocks. It selects the lowest-energy sequence among all candidates evaluated in the current cycle and executes its first action. The 25-trial main comparison in Section 4 separately evaluates closed-loop performance under this finite budget. Fixed-kernel asymptotic properties do not imply that every local block has mixed sufficiently within eight steps.

4. Closed-Loop Experiments and Comparisons in a Dual-Ramp Merging Scenario

The dual-ramp merging scenario represents freeway interchanges and adjacent entrance segments, where interacting traffic streams create operational complexity. Compared with a single-ramp scenario, the dual-ramp scenario involves interactions across space among multiple entrance demands, mainline vehicles, and merging maneuvers. Speed fluctuations, merging conflicts, vehicle queues, and disturbance propagation become intertwined. Multiple ramp release rates and mainline VSLs jointly affect mainline conditions, coupling the control inputs and requiring coordinated decisions. This scenario therefore provides a representative simulation environment for testing adaptation to complex traffic and examining efficiency, safety risk, and command smoothness.

The two ramps alternate between high and low demand and share a downstream bottleneck, linking speed-limit and release decisions at different locations. Two speed-limit variables and two ramp-flow variables are jointly optimized over a multistep horizon. This setting tests adaptive feedback and finite-budget search for coupled multi-input control.

Inputs are scaled with reference to real traffic demand. SUMO simulates vehicle motion under dual-ramp merging, alternating demand, and a downstream bottleneck. We compare risk, traffic operations, and command variation across EA-QD-CVaR, feedback ablations, and classical controls. Table 3 specifies the vehicle composition and traffic demand settings. All methods use the same vehicle composition and traffic demand levels in the 25-trial main comparison.

Table 3. Driving composition and demand in the dual-ramp scenario. The upper panel lists vehicle composition. The lower panel describes the 240 s evaluation window, with high and low demands of 1100 and 300 veh/h, respectively.

Driving Class	**Normal**	**Slow**	**Faster**
Proportion (%)	25	25	50
High/Low Demand Time Share (%)	**Ramp 1 High/Low Duration (s)**	**Ramp 2 High/Low Duration (s)**	**Initial Demand Ramp 1/Ramp 2**
50/50	120/120	120/120	Low/high

4.1. Experimental Environment and Traffic Composition

SUMO 1.27.1 exchanges vehicle states and commands with a Python 3.12.7 controller through TraCI. Qiskit Aer's `AerSimulator` computes ideal local four-qubit statevectors on a classical computer in `statevector` mode, extracting proposal probabilities for nine valid candidates. The mainline is 4 km long: two lanes over the upstream 3 km and a single-lane bottleneck over the downstream 1 km, with a fixed 24 m/s speed limit. Entrance ramps merge 1.0 and 2.5 km from the mainline entrance. The VSL regions span 0–1.5 and 1.5–3.0 km, respectively (Figure 3). Ramp control assigns release budgets to demand-side vehicle insertion; unreleased demand remains external waiting demand.

During evaluation, prescribed mainline demand is 2800 veh/h. The two ramps alternate high and low demands of 1100 and 300 veh/h every 45 s, maintaining total ramp demand at 1400 veh/h. During the first 300 s, mainline and ramp demands are 1600, 200, and 200 veh/h, followed by alternating demand. Warm-up continues to 720 s, followed by 240 s of closed-loop evaluation, for a total of 960 s. The evaluation contains 16 control decisions at 15 s intervals and 80 observation frames sampled every 3 s.

These durations are simulation times. The experiments use an offline synchronous closed loop. Candidate search finishes before SUMO executes the selected command and

advances simulation time by 15 s; the simulation clock remains fixed during computation. Thus, 15 s denotes the simulation interval for traffic-state updates and command holding. Search budgets count candidate evaluations, while wall-clock computation time is recorded separately.

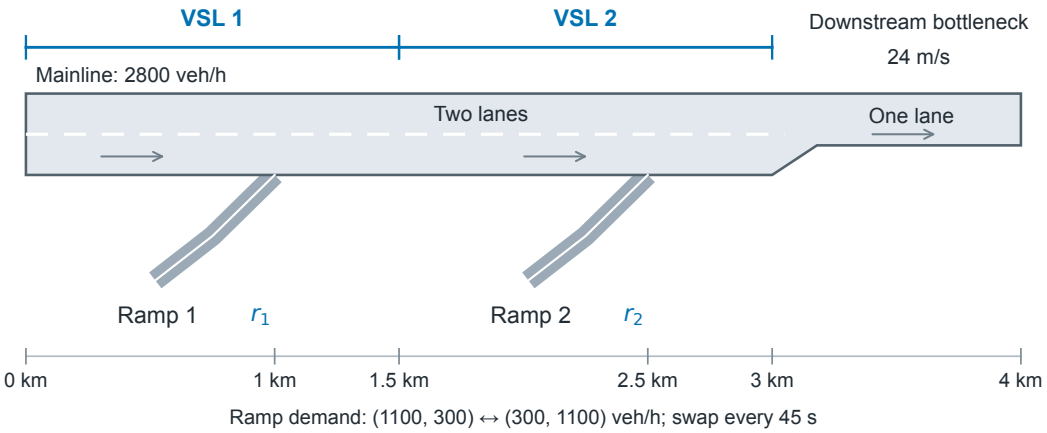


Figure 3. Synthetic network with two entrance ramps and two VSL regions. Both speed limits and both ramp release rates are independently controlled. Ramp demands swap every 45 s. Input demand and controlled release are defined separately; ramp curvature and lane widths are schematic.

The alternating-demand phase starts at simulation time 300 s, with a full period of 90 s. The evaluation window $[720, 960)$ s starts at phase 60 s, when ramp 1 has low demand and ramp 2 high demand. Demand switches occur at 750, 795, 840, 885, and 930 s within the window. Second-by-second accumulation of the demand function gives each ramp 120 s at high demand and 120 s at low demand, or 50% each (Table 3).

All three driving classes use the Krauss passenger-car following model. Acceleration is 2.6 m/s^2, deceleration is 4.5 m/s^2, the time-gap parameter is 1 s, the minimum gap is 2.5 m, the vehicle length is 5 m, the stochasticity is 0.35, and the maximum vehicle speed is 33.33 m/s. Slow and faster drivers use truncated desired-speed-factor distributions with means 0.75 and 1.10, respectively, and SD 0.05. Their truncation ranges are $[0.6, 0.9]$ and $[1.0, 1.2]$, respectively. The parameters above are shared by all classes.

Normal, slow, and faster drivers account for 25%, 25%, and 50% of vehicles, respectively. This configuration emphasizes faster drivers while retaining mixed desired speeds to test merging interactions, risk exposure, and cooperative control under speed heterogeneity. The proportions are simulation test settings, not estimates from NGSIM. NGSIM I–80 informs demand and operating scales; SUMO generates vehicle trajectories and control responses. The dual-ramp network is an engineering-representative synthetic scenario with adjacent entrances, alternating demand, and a common bottleneck.

4.2. Control Scale, Parameters, and Comparison Methods

The command at each control instant is $\mathbf{u} = (v_1, r_1, v_2, r_2)$. Each VSL has levels 20, 25, and 30 m/s; each ramp has levels 300, 600, and 900 veh/h. The horizon has $N_p = 16$ steps of 15 s, totaling 240 s. There are $M = 3^4 = 81$ joint actions per step. A complete sequence contains 64 ternary components, with nominal size $N_{\text{global}} = 3^{64} \approx 3.43 \times 10^{30}$. A local proposal changes one VSL–ramp pair at one prediction position, yielding $N_{\text{local}} = 9$ candidates. Two qubits embed each ternary variable, giving four qubits per local register.

SUMO microsimulation predicts candidate trajectories. Each evaluation replays executed commands from the initial traffic state to restore vehicle positions, speeds, driving types, waiting demand, and remaining ramp-release allowances at the decision instant. The restored state is checked. Future demand uses the mean prescribed demand over

the preceding 60 s. The bottleneck speed limit retains its observed value, and predicted new vehicles use the specified composition and fixed prediction random stream. The 16 candidate actions are then executed for 15 s each. Vehicle counts, mainline speeds, and valid following pairs are sampled every 3 s, and downstream arrivals are accumulated.

Each cycle evaluates 64 distinct complete 16-step candidate sequences, including the hold-action reference; repeated queries use the cycle's cache. Once all nine local energies are available, eight angle sets are compared, followed by eight MH transitions. Angles minimize the circuit-probability-weighted lower 25% mean of raw energies; migration angles follow $\tau = \kappa_t U(0.75, 1.25)$. The first action of the lowest-energy evaluated sequence is executed. Table 4 and Algorithm 1 specify the parameters and procedure.

Table 4. Algorithm and metric parameters for the dual-ramp experiment. Weight offsets, explicit clipping bounds, and attainable ranges under normalized inputs are listed separately.

Parameter	Value and Meaning
N_p, M, and N	16 steps; per-step $M = 81$, local $N_{\text{local}} = 9$, full-sequence $N_{\text{global}} = 3^{64}$
Initial sequence	Initial command $(30, 900, 30, 900)$; repeat the preceding executed command over 16 steps each cycle
Weight offsets	$(\alpha_0, \beta_0, w_{U,0}) = (0.30, 0.45, 0.35)$, the constant terms in the update equations
Feedback gains	$(k_T, k_H, k_p, k_{UT}, k_{UR}) = (0.45, 1.80, 4.50, 0.25, 0.25)$
Efficiency/smoothness bounds	No explicit clipping of $\alpha_t, w_{U,t}$; attainable ranges $[0.30, 0.75]$ and $[0.35, 0.85]$, respectively
Safety-weight bounds	Explicit clipping of β_t to $[0.45, 8.00]$; attainable range $[0.45, 6.75]$ under default gains
Circuit migration scale	$\kappa_t = 0.50 + 0.35(H_T - 0.5) + 0.35(H_R - 0.5)$; no explicit clipping; attainable range $[0.15, 0.85]$
Evaluation budget	64 new complete sequences per cycle, including the hold reference; duplicates cached, with reevaluation across cycles
Local circuit	Four qubits, nine valid combinations, one ordered proposal layer; 8 angle sets compared per block
Angle generation	Independent uniform draws: $\varphi \in [0, 2\pi)$, $\theta_v, \theta_r \in [0, \pi)$, $\tau/\kappa_t \in [0.75, 1.25)$
Parameter selection	Minimize the circuit-weighted lower 25% mean of raw energies, reusing evaluated energies
Energy normalization	Local min–max normalization for phase gates, denominator floor 10^{-12}; raw composite energy for MH
MH settings	Inverse temperature 0.5, proposal mixing mass 10^{-3}, 8 transitions per complete local block
Entropy categories	Three H_T bins with boundaries $(-1, 18, 24, 35)$ m/s; four H_R categories
Risk classification	TTC thresholds $(6, 3, 1.5)$ s and net-headway thresholds $(2, 1.5, 0.8)$ s; retain the higher risk level
Evaluation thresholds	TTC: 7 s; joint-risk net headway: 3 s
Tail/numerical protection	$\alpha_R = 0.75$, $\lambda_R = 0.25$; $\varepsilon_{\Delta v} = 10^{-6}$ m/s, $\varepsilon_v = 0.1$ m/s
Cost scales	$n_{\text{ref}} = 150$ vehicles, $v_{\text{ref}} = 30$ m/s, $A_{\text{ref}} = 5$ vehicles; $a_v = 0.5$, $a_A = 0.25$
Cost coefficients	$\delta = 0.9$, $c_S = 0.5$, $s_R = 0.1$, $c_U = 0.25$, $\lambda_P = 10$
Operational protection	Accumulation tolerance $0.05n^{\text{h}} + 2$ vehicles, arrival tolerance 1 vehicle; scales 150 and 5 vehicles
Computation time	Per-cycle wall time for search and candidate-log archiving; statistics in Section 4.2

The 240 s horizon evaluates each candidate's predicted consequences; the search budget counts newly evaluated complete sequences. The 16 positions contain two VSL–ramp blocks each, giving 32 possible blocks. EA-QD-CVaR and QD ablations visit positions from earliest to latest, with region 1 before region 2. The 64-evaluation budget completes seven nine-candidate blocks per cycle. One evaluation scores the hold reference; each complete block reuses its evaluated center and adds eight sequences, totaling $1 + 7 \times 8 = 57$ evaluations. The remaining seven evaluate part of the eighth block. The first seven blocks each perform eight MH transitions. Evaluated candidates in block eight enter the cycle-wide minimum-cost selection, without an MH update on incomplete energies.

For positions $\ell = 0, \ldots, 15$, complete search covers both variable pairs at $\ell = 0, 1, 2$ and region 1 at $\ell = 3$. Region 2 at $\ell = 3$ receives partial evaluation only. Positions $\ell = 4, \ldots, 15$ retain the preceding executed command from the cycle's initial sequence and undergo the full 240 s simulation with each candidate. Consequently, full-horizon costs include the downstream effects of early changes on queues, mainline speed, risk, and arrivals. Only the first 15 s action is executed, followed by reinitialization and search using new observations. The 32 possible blocks therefore receive unequal search coverage and evaluation budgets.

Table 5 lists 13 methods. QD-CVaR is the fixed-parameter baseline for comparison with complete feedback. Ablations compare modules comprising each signal and its associated parameter updates. Uniform, SA, and Greedy-global compare search strategies under finite candidate budgets. Greedy-local provides a local greedy reference; MPC-CVaR provides risk-sensitive predictive control. No-Control and ALINEA represent no active control and conventional ramp-feedback control, respectively. Entropies use mainline speed categories and valid following-pair risk categories at each control instant, with a 15 s control interval.

Table 5. Comparison methods, enabled feedback, and sample sizes in the dual-ramp scenario. A value of 1 indicates enabled feedback; a dash indicates that the feedback is not used.

Method	n	(H_T, H_R, p_R)	Search or Control Mechanism
EA-QD-CVaR	25	$(1, 1, 1)$	Complete feedback, local circuit proposals, and MH
QD-CVaR	25	$(0, 0, 0)$	Same local search; fixed objective weights and migration parameters
QD-HT	25	$(1, 0, 0)$	Traffic-state entropy feedback only
QD-HR-only	25	$(0, 1, 0)$	Risk-entropy feedback only
QD-pR-only	25	$(0, 0, 1)$	Joint-risk-proportion feedback only
QD-HR-pR	25	$(0, 1, 1)$	Risk-entropy and risk-proportion feedback
Uniform	25	$(1, 1, 1)$	Same local blocks and adaptive objective; uniform proposals, all accepted
SA	25	—	Fixed weights, single-coordinate proposals, annealing acceptance
Greedy-global	25	—	Fixed weights; budgeted single-coordinate neighbors of complete sequences
Greedy-local	1	—	Fixed weights; greedy selection among nine local candidates
MPC-CVaR	25	—	Fixed-weight categorical cross-entropy shooting MPC with CVaR
No-Control	1	—	No active VSL or ramp metering
ALINEA	1	—	Occupancy feedback at both ramps; VSLs held at 30 m/s

The fixed-weight QD-CVaR baseline uses $(\alpha, \beta, w_U) = (0.4679, 0.7063, 0.5856)$ and migration scale $\kappa = 0.50$. Safety weight $\beta = 0.7063$ was selected from preliminary parameter

exploration because it gave the highest downstream arrivals among the configurations tested. It remains fixed across all 25 search-seed trials. This efficiency-oriented baseline tests whether EA-QD-CVaR can further improve safety through entropy-guided adaptive weighting. Relative changes in the main comparison use this fixed-parameter method's mean as the reference.

Each QD configuration starts from fixed parameters and replaces only the updates associated with enabled signals (Table 6). H_T updates the efficiency weight, smoothness weight, and circuit migration scale. H_R updates the safety weight, smoothness weight, and migration scale. p_R updates only the safety weight. With risk feedback enabled, the safety weight uses offset 0.45, adds enabled contributions $1.80H_R$ or $4.50p_R$, and clips to $[0.45, 8.00]$. Without risk feedback, it retains the fixed baseline value. All groups retain the same TTC-event and CVaR components in the safety cost.

Table 6. QD feedback configurations and associated parameter updates. Listed weights update with observations; unlisted weights retain fixed baseline values. κ_t is the circuit migration scale.

Method	Enabled Signals	Updated Weights	Migration Scale κ_t
EA-QD-CVaR	H_T, H_R, p_R	$\alpha_t, \beta_t, w_{U,t}$	Updated by H_T, H_R
QD-CVaR	None	None; all fixed	Fixed
QD-HT	H_T	$\alpha_t, w_{U,t}$	Updated by H_T
QD-HR-only	H_R	$\beta_t, w_{U,t}$	Updated by H_R
QD-pR-only	p_R	β_t	Fixed
QD-HR-pR	H_R, p_R	$\beta_t, w_{U,t}$	Updated by H_R

Starting from the fixed baseline, the smoothness weight adds $0.25(H_T - 0.4712)$ or $0.25(H_R - 0.4712)$ for each enabled entropy module. The common decomposition reference 0.4712 recovers the complete smoothness weight in Equation (21) when both entropies are enabled. The migration scale starts at 0.50 and adds $0.35(H_T - 0.5)$ or $0.35(H_R - 0.5)$ for each enabled entropy module. Disabling a signal retains the fixed values or enabled updates of the remaining parameters. Comparisons therefore concern the combined feedback and parameter paths in Table 6. Offsets and gains use the presets in Table 4, unchanged across search repetitions in the dual-ramp scenario.

SA, Greedy-global, Greedy-local, and MPC-CVaR share QD-CVaR's predictor, CVaR loss, and fixed weights $(0.4679, 0.7063, 0.5856)$. Each search starts from a 16-step sequence repeating the previous command. The reference counts toward 64 new evaluations, and the first action of the lowest-cost evaluated sequence is executed. All budgets count unique complete sequences; cache hits do not increase the count. Internal transitions, neighborhood visits, and shooting generations depend on each search mechanism. Table 7 gives the classical settings.

Table 7. Implementation and budget allocation of classical comparisons. Search methods use fixed objective weights and 64 unique complete-sequence evaluations per cycle.

Method	Implementation Settings
SA	Uniformly select one of 64 ternary coordinates, then uniformly propose one of its two alternative levels. Temperature follows $T(s) = \exp[\log(0.1)\min\{1, (s-1)/31\}]$ for cumulative new-evaluation count s, decreasing from 1 to 0.1 and then remaining fixed. Acceptance is $\min\{1, \exp[-\Delta\mathcal{E}/T(s)]\}$, with no restart.
Greedy-global	The hold sequence has $64 \times 2 = 128$ single-coordinate neighbors. Each round fixes the neighborhood center, visits the first action's four coordinates first, and randomly orders the remaining 60. Two alternative levels per coordinate are randomly ordered and visited in two passes. The 64 evaluations cover the reference and 63 new neighbors. Improvement requires an energy decrease exceeding 10^{-12}.

Table 7. *Cont.*

Method	Implementation Settings
Greedy-local	Evaluate nine local candidates in prediction-step and region order; use the lowest-energy sequence as the next center. Ties favor the current center, then lexicographic order of the complete level sequence. Evaluated candidates in an incomplete final block remain eligible.
MPC-CVaR	Categorical cross-entropy shooting with 16 candidates and 4 elites per generation, update rate 0.7, and uniform mixing 0.05. Initialization assigns an additional mass of 0.8 to the hold level and distributes the remaining 0.2 equally across three levels. Evaluated candidates in the final partial generation remain eligible.
No-Control	Both VSL regions remain at 30 m/s. Ramp vehicles enter according to demand without active release quotas of 300–900 veh/h, subject to SUMO insertion conditions. The downstream bottleneck speed limit is unchanged.
ALINEA	Each ramp uses occupancy from its downstream detectors, with a target of 20% and a gain of 70 veh/h per occupancy percentage point. Internal release starts at 900 veh/h, updates every 15 s, and is quantized to 300, 600, or 900 veh/h. VSLs remain at 30 m/s; queue override is disabled.

ALINEA first averages occupancy across the two downstream lane detectors for each ramp, then over the preceding 15 s, giving $o_{j,t}$ in percent. The internal continuous release rate and executed level are

$$\widehat{r}_{j,t} = \mathrm{clip}_{[300,900]}\left[\widehat{r}_{j,t-1} + 70(20 - o_{j,t})\right], \qquad r_{j,t} = \mathcal{Q}_{\{300,600,900\}}(\widehat{r}_{j,t}). \tag{54}$$

Here, $\widehat{r}_{j,-1} = 900$ veh/h, and $\mathcal{Q}$ selects the nearest level, choosing the lower when the distances are equal. Recursion uses the preceding unquantized internal release rate. Both VSL commands independently remain at 30 m/s.

Table 8 summarizes wall times for 251 search-based runs and 4016 control cycles. Timing consistently starts at evaluator initialization and ends after search and compressed candidate-log archiving. It includes history replay, state restoration, 240 s prediction for 64 candidates, selection, and applicable circuit and MH operations. It excludes observation preparation before evaluator creation, the 15 s simulated execution of the selected action, and actual trajectory-file writing. No-Control and ALINEA invoke no candidate search; their control computation was not separately timed. Statistics are grouped by hardware. The Mac mini uses an Apple M6, 12 logical cores, and 16 GiB of memory; the local computer uses an Apple M1 Pro, 10 logical cores, and 32 GiB of memory. Each group is timed under its experimental batch's concurrent load. Mean, 95th-percentile, and maximum times are calculated across control cycles.

Table 8. Per-cycle wall times for search and candidate-log archiving, grouped by hardware. Each run has 16 cycles, with 64 complete-sequence evaluations per cycle under the experimental batch's concurrent load.

Method	Runs	Cycles	Mean (s)	95th Percentile (s)	Maximum (s)
Mac mini: Apple M6, 12 logical cores, 16 GiB memory					
EA-QD-CVaR	23	368	110.71	125.17	128.10
QD-CVaR	25	400	107.45	123.43	129.30
QD-HT	23	368	110.19	123.79	125.48
QD-HR-only	25	400	109.91	124.82	128.55
QD-pR-only	25	400	109.22	124.46	127.54

Table 8. *Cont.*

Method	Runs	Cycles	Mean (s)	95th Percentile (s)	Maximum (s)
QD-HR-pR	23	368	110.77	124.78	127.83
Uniform	23	368	111.22	124.31	126.79
SA	23	368	107.09	123.63	126.48
Greedy-global	23	368	108.87	124.07	133.00
MPC-CVaR	23	368	106.19	122.92	126.85
Local computer: Apple M1 Pro, 10 logical cores, 32 GiB memory					
EA-QD-CVaR	2	32	154.45	170.17	170.88
QD-HT	2	32	151.81	169.09	171.72
QD-HR-pR	2	32	154.55	171.34	173.16
Uniform	2	32	155.70	171.13	172.47
SA	2	32	154.60	168.99	171.62
Greedy-global	2	32	152.75	168.74	169.52
Greedy-local	1	16	151.95	167.63	168.88
MPC-CVaR	2	32	152.13	165.20	167.42

This study develops the control method and evaluates it through offline closed-loop simulation. The current software implementation has not undergone dedicated computational optimization for online deployment. The experiments use an offline synchronous loop: simulation time remains paused during candidate evaluation and search. Once computation finishes, the selected command is executed and simulation time advances by 15 s. Thus, 15 s denotes the simulated interval for traffic-state updates and command holding. Under the specified hardware and experimental load, EA-QD-CVaR averages 110.71 s across 368 Mac mini cycles, with a 95th percentile of 125.17 s. The maximum is 128.10 s, while 32 local-computer cycles average 154.45 s. Current computation exceeds 15 s and does not meet the corresponding online update deadline. The reported control performance therefore applies to offline synchronous execution. Deployment requires further implementation optimization and evaluation of computational delay effects on closed-loop control.

4.3. Evaluation Metrics and Paired Statistics

The TTC metric is the frame-wise proportion of valid following pairs with TTC below 7 s, averaged equally over 80 frames per run. The 7 s threshold identifies potential conflicts early; its literature basis is given in Section 2.3. CVaR uses loss $L = \max(0, 1 - \mathrm{TTC}/7)$ and averages the worst 25% of losses across all valid pairs in the evaluation window. Zero losses are retained, and fractional probability mass is used at the tail boundary. Joint risk follows the event union in Equation (7), with a separate evaluation threshold:

$$p_J(t) = \frac{1}{n_t^{\mathrm{pair}}} \sum_{i \in \mathcal{P}_t} \mathbb{I}\{\mathrm{TTC}_i < 7\,\mathrm{s} \ \vee \ h_i < 3\,\mathrm{s}\}. \tag{55}$$

Joint risk counts the union of events; a pair satisfying both conditions is counted once.

Downstream arrivals count vehicles reaching the network exit during the 240 s evaluation window. Mean speed divides the sum of mainline vehicle speeds across 80 frames by the total number of vehicle observations. Control variation is the mean normalized absolute change in executed commands:

$$F_u = \frac{1}{16} \sum_{k=1}^{16} \frac{1}{4} \sum_{j=1}^{4} \frac{|u_{k,j} - u_{k-1,j}|}{D_j}, \quad \mathbf{D} = (10\,\mathrm{m/s}, 600\,\mathrm{veh/h}, 10\,\mathrm{m/s}, 600\,\mathrm{veh/h}). \tag{56}$$

The first change is relative to the initial command. Dimensionless F_u differs from the squared smoothness cost used in optimization in Equation (11). Higher arrivals and speed indicate better operations within the evaluation window. Lower TTC risk, CVaR, joint risk, and F_u indicate better corresponding outcomes. All risk measures are surrogate indicators.

For each stochastic method, the main comparison reports the mean, sample SD (ddof = 1), and pointwise 95% BCa confidence interval of the mean across 25 runs. EA-QD-CVaR and stochastic comparisons are paired by search seed. Differences are EA-QD-CVaR minus the comparison, and relative changes are $(\bar{x}_{\mathrm{EA}}/\bar{x}_{\mathrm{control}} - 1) \times 100\%$. Differences in risk proportions use percentage points; relative changes use percent. Paired mean-difference intervals use 200000 paired-seed BCa bootstrap resamples with a fixed random stream, preserving seed correspondence across methods within each resample. Deterministic references contribute single observations only, without SDs, intervals, or paired tests.

Paired differences use two-sided Wilcoxon signed-rank tests, excluding zeros and assigning average ranks to tied absolute differences. The conditional sign distribution gives p values. Ranks are determined after rounding differences to 12 decimal places. Under symmetric paired differences, Wilcoxon tests assess distribution location; mean-difference intervals describe average effects. The six primary metrics are TTC risk, CVaR, joint risk at the 3 s threshold, arrivals, mean speed, and F_u. Each compares EA-QD-CVaR with nine stochastic controls. In the 25-pair main comparison, p denotes the raw two-sided Wilcoxon signed-rank value. Holm p denotes adjustment across all 54 tests involving six metrics and nine stochastic comparisons. Raw and adjusted values are each evaluated against a threshold of 0.05. Intervals are estimated separately for each metric. The 1.5 s net-headway analysis is reported separately as an offline sensitivity result.

4.4. Fixed-Safety-Weight Sensitivity and Safety–Efficiency Trade-Offs

We scan fixed safety weights $\beta \in \{0.45, 2, 4, 6, 8\}$ in the same dual-ramp scenario, with 25 independent search trials per setting. Only the safety weight changes. The efficiency weight $\alpha = 0.4679$, smoothness weight $w_{lat} = 0.5856$, migration scale $\kappa = 0.50$, and objective components remain fixed. The 16-step prediction horizon and budget of 64 complete-candidate-sequence evaluations per cycle are also unchanged. The EA-QD-CVaR reference uses the same 25-trial main-comparison results. Table 9 reports group means and sample SDs.

Table 9. Risk, arrivals, and command variation in the fixed-safety-weight scan and EA-QD-CVaR. Each group has 25 trials; values are mean±sample SD. EA-QD-CVaR uses the 25-trial main-comparison results. Dimensionless command variation F_u is defined in Equation (56); lower values indicate smaller changes between consecutive commands.

Method	β	TTC Risk (%)	CVaR	Arrivals (Vehicles)	F_u
QD-CVaR	0.45	9.1625 ±0.6804	0.13935 ±0.01054	136.84 ±2.15	0.15563 ±0.03314
QD-CVaR	2	8.8194 ±0.8727	0.13354 ±0.01224	137.04 ±1.77	0.27344 ±0.06123
QD-CVaR	4	8.6445 ±0.7247	0.13288 ±0.01107	136.96 ±2.61	0.29531 ±0.05399
QD-CVaR	6	8.5267 ±0.7913	0.12961 ±0.01286	135.92 ±2.31	0.30656 ±0.05648
QD-CVaR	8	8.5367 ±0.6433	0.13001 ±0.01056	136.24 ±1.39	0.30000 ±0.04688
EA-QD-CVaR	Adaptive	8.0689 ±0.5715	0.12462 ±0.00799	137.28 ±1.59	0.28125 ±0.05980

Increasing the safety weight from 0.45 to 6 reduces mean TTC risk from 9.1625% to 8.5267% and CVaR from 0.13935 to 0.12961. At 8, both means increase slightly to 8.5367% and

0.13001, respectively. Mean arrivals across the five settings are 136.84, 137.04, 136.96, 135.92, and 136.24 vehicles, showing a non-monotonic response. Fixed-safety-weight changes affect both risk and traffic operations, requiring joint examination of their mean trade-offs.

For each of the 25 EA-QD-CVaR trials, weights are first time-averaged over the 240 s evaluation window, then averaged equally across trials. Each trial contains 16 control cycles lasting 15 s each, so its time average equals the arithmetic mean of the 16 applied cycle weights. The resulting mean efficiency, safety, and smoothness weights are $\bar{\alpha} = 0.5196$, $\bar{\beta} = 4.1012$, and $\bar{w}_U = 0.6563$, respectively. These values summarize overall weight levels during evaluation. Cycle-specific weights continue to follow Equation (21), using the current traffic-state entropy, risk entropy, and risk-event proportion.

The fixed-weight scan shows that increasing the safety weight can reduce TTC risk and CVaR within a limited range, while increasing command variation. Raising β from 0.45 to 6 reduces mean TTC risk by approximately 6.94%. Meanwhile, F_u rises from 0.15563 to 0.30656, an increase of approximately 97%. At 8, both risk means increase slightly, indicating saturation of the risk improvement. Compared with the high-safety-weight settings $\beta = 6$ and $\beta = 8$, EA-QD-CVaR achieves lower TTC risk and CVaR with 8.26% and 6.25% lower F_u, respectively. It also yields higher downstream arrivals. Across the tested mean outcomes, entropy-guided adaptive weighting in the complete method provides a favorable trade-off among safety, traffic output, and command smoothness.

Lower risk and higher arrivals are preferred. A mean point dominates another if it is no worse on both metrics and strictly better on at least one. Among fixed-weight points, $\beta = 6$ has the lowest mean TTC risk and CVaR, while $\beta = 2$ has the highest mean arrivals. In both the TTC-risk–arrivals and CVaR–arrivals planes, $\beta = 2, 4, 6, 8$ form the empirical frontier of tested fixed-weight points. The $\beta = 0.45$ point is dominated by $\beta = 2$ and $\beta = 4$. Figures 4 and 5 show the means and SDs for risks and arrivals; the gray dashed line connects the points in weight-scan order.

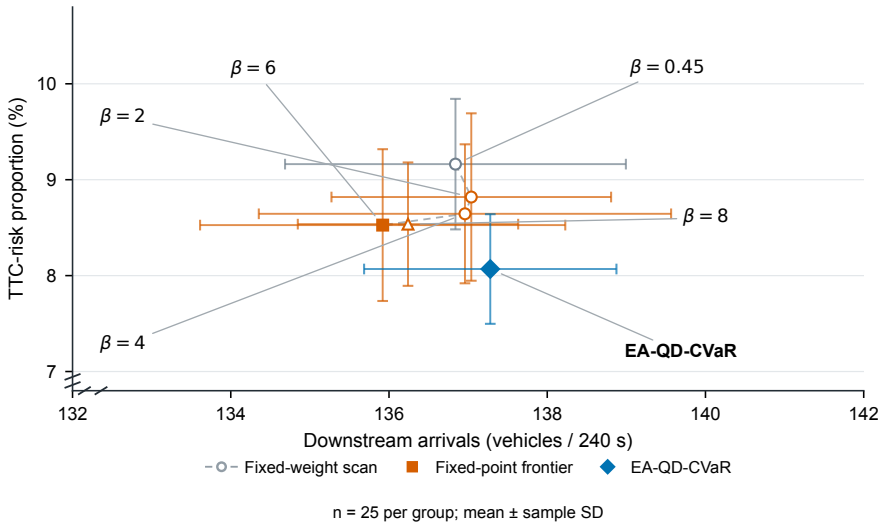


Figure 4. TTC risk and downstream arrivals at different fixed safety weights. Each group has 25 trials; horizontal and vertical error bars show sample SDs around the means. Orange markers indicate the empirical frontier of tested fixed-weight points, and the blue diamond denotes EA-QD-CVaR. Squares and triangles identify $\beta = 6$ and $\beta = 8$, respectively. The gray dashed line follows increasing β; the empirical frontier is defined by the tested mean points. Axis-break symbols indicate nonzero display origins.

EA-QD-CVaR has mean TTC risk 8.0689%, CVaR 0.12462, and arrivals 137.28 vehicles. Relative to $\beta = 6$, the fixed setting with the lowest risk means, TTC risk and CVaR decrease by 5.37% and 3.85%, respectively. Arrivals increase by 1.36 vehicles. Relative to $\beta = 2$, the highest-arrival fixed setting, the two risks decrease by 8.51% and 6.68%, respectively, while

arrivals increase by 0.24 vehicles. Relative to $\beta = 8$, the reductions are 5.48% and 4.15%, respectively, with 1.04 additional arrivals. Across the tested weight range and 25-run group means, EA-QD-CVaR dominates all five fixed-weight points in both risk–arrivals planes. It therefore combines lower mean risks with higher mean arrivals in this comparison.

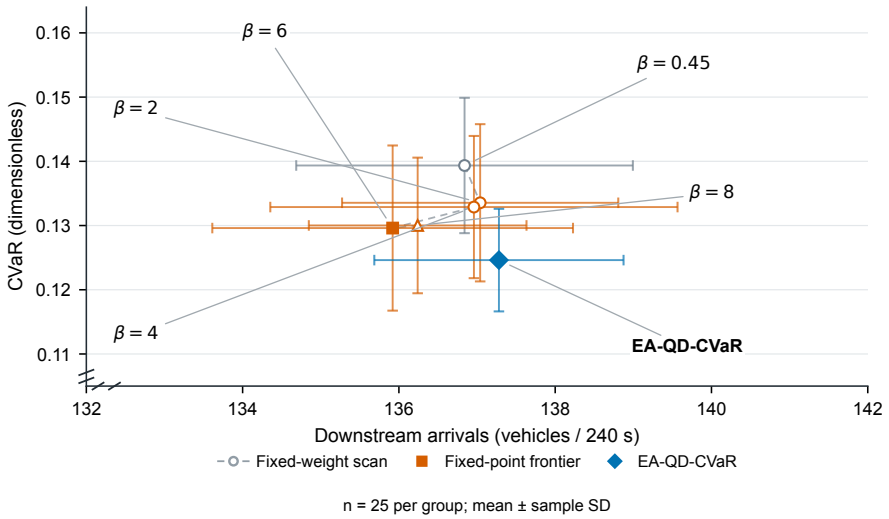


Figure 5. CVaR and downstream arrivals at different fixed safety weights. Each group has 25 trials, and error bars show sample SDs. Colors, markers, the empirical frontier, and the scan-order line follow Figure 4. Axis-break symbols indicate nonzero display origins.

4.5. Risk Results and Main Fixed-Weight Comparison

Table 10 lists the risk means, sample SDs, mean intervals, and sample sizes for 13 methods. Figures 6 and 7 show every run. Boxes span the 25th–75th percentiles, center lines mark medians, whiskers use 1.5 times the interquartile range, and diamonds mark 25-run means. Each deterministic reference is a single square.

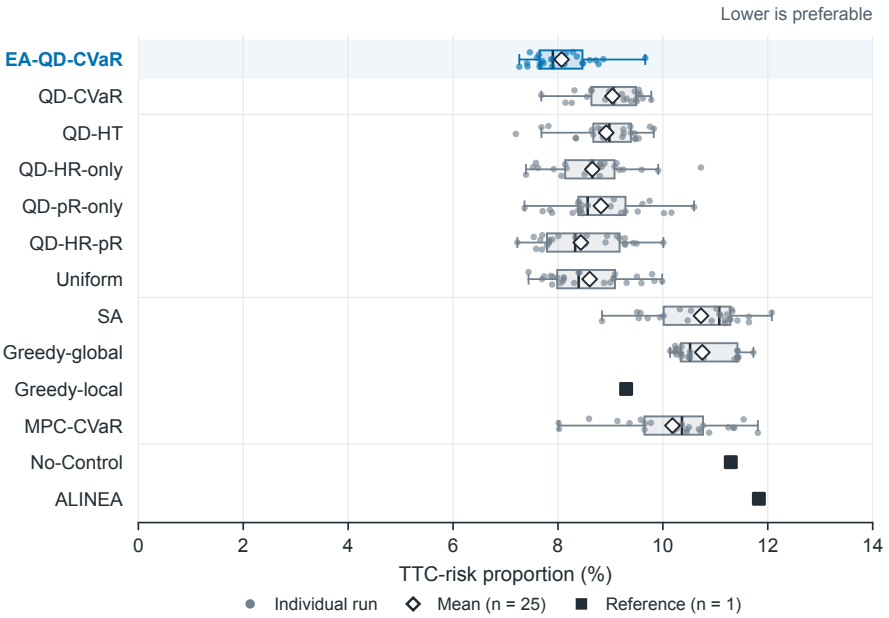


Figure 6. TTC-risk proportions of 13 methods in the dual-ramp scenario, with a 7 s threshold; lower is preferable. The vertical axis lists control methods, and the horizontal axis shows risk. Each of ten stochastic methods has 25 points and a horizontal box plot; each deterministic reference has one point. Diamonds mark stochastic-method means. EA-QD-CVaR is blue, other stochastic methods are gray, and deterministic references are black squares.

Table 10. Risk metrics for 13 methods in the dual-ramp scenario. For stochastic methods ($n = 25$ each), upper entries are the mean±sample SD, and lower entries are the mean 95% BCa intervals. The three deterministic references ($n = 1$ each) report single observations.

Method	*n*	TTC (%)	CVaR	Joint Risk (%)
EA-QD-CVaR	25	8.069 ± 0.572 [7.875, 8.319]	0.12462 ± 0.00799 [0.12184, 0.12801]	76.712 ± 3.944 [75.399, 78.482]
QD-CVaR	25	9.039 ± 0.536 [8.808, 9.222]	0.13771 ± 0.00950 [0.13377, 0.14107]	80.504 ± 2.990 [79.266, 81.573]
QD-HT	25	8.920 ± 0.650 [8.628, 9.136]	0.13449 ± 0.01055 [0.13026, 0.13835]	80.397 ± 3.568 [78.984, 81.726]
QD-HR-only	25	8.655 ± 0.776 [8.379, 8.977]	0.13314 ± 0.01241 [0.12861, 0.13816]	78.369 ± 4.347 [76.798, 80.145]
QD-pR-only	25	8.819 ± 0.804 [8.527, 9.144]	0.13397 ± 0.01200 [0.12989, 0.13919]	78.821 ± 3.996 [77.318, 80.392]
QD-HR-pR	25	8.436 ± 0.789 [8.144, 8.751]	0.12949 ± 0.01210 [0.12506, 0.13437]	77.443 ± 4.172 [75.936, 79.148]
Uniform	25	8.604 ± 0.767 [8.324, 8.915]	0.13265 ± 0.01410 [0.12789, 0.13890]	77.913 ± 3.662 [76.583, 79.400]
SA	25	10.727 ± 0.830 [10.379, 11.019]	0.16788 ± 0.01345 [0.16231, 0.17266]	82.868 ± 1.979 [81.978, 83.523]
Greedy-global	25	10.752 ± 0.523 [10.567, 10.971]	0.16771 ± 0.00824 [0.16472, 0.17107]	83.981 ± 1.230 [83.124, 84.287]
Greedy-local	1	9.299	0.13867	80.990
MPC-CVaR	25	10.183 ± 1.017 [9.758, 10.541]	0.15836 ± 0.01618 [0.15184, 0.16424]	82.253 ± 3.242 [80.739, 83.300]
No-Control	1	11.295	0.17728	85.378
ALINEA	1	11.830	0.18227	83.283

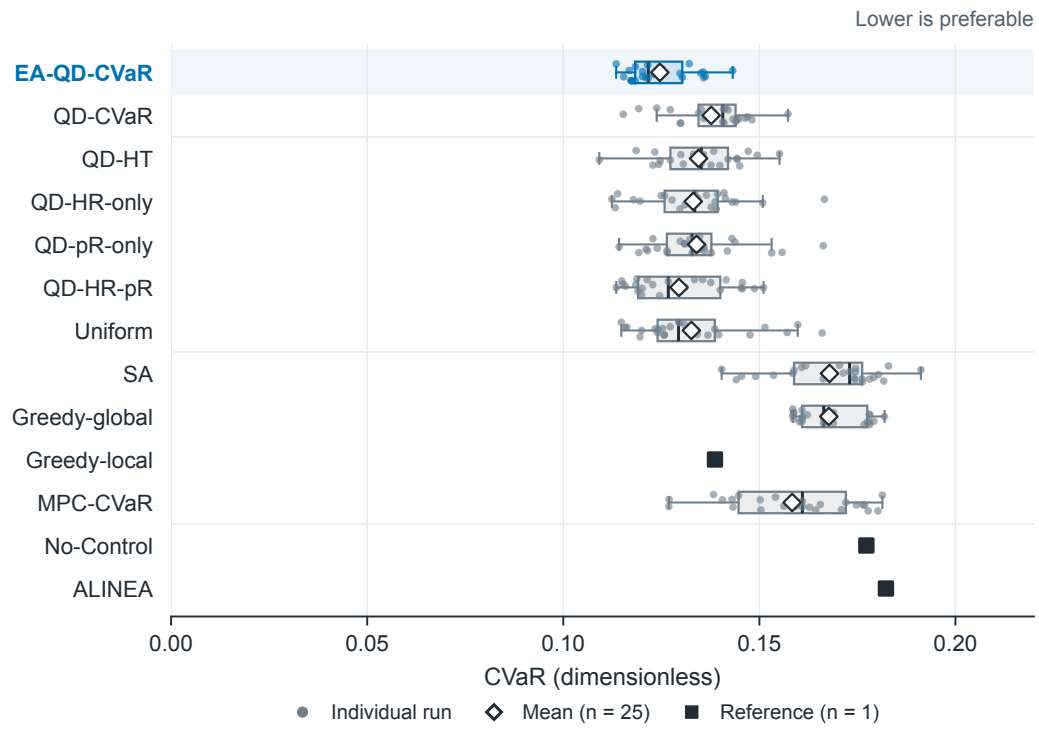


Figure 7. Upper-tail CVaR of TTC loss in the dual-ramp scenario at confidence level 0.75; lower is preferable. Observations, sample sizes, and box-plot definitions follow Figure 6.

EA-QD-CVaR achieves TTC risk $8.069\% \pm 0.572\%$ and CVaR achieves 0.12462 ± 0.00799, the lowest means among the ten stochastic methods. QD-CVaR has corresponding means of 9.039% and 0.13771, giving relative reductions of 10.73% and 9.51%. In Table 11, the paired mean TTC difference is -0.970 percentage points, with a 95% BCa interval of

$[-1.240, -0.650]$. The CVaR difference is -0.01310, with an interval of $[-0.01762, -0.00749]$. The respective p values are 1.007×10^{-5} and 0.000103. EA-QD-CVaR achieves lower TTC risk in 21 pairs and lower CVaR in 20 of the 25 pairs.

Table 11. Risk comparison of EA-QD-CVaR with QD-CVaR, with 25 pairs per metric. Proportion differences are percentage points. Δ is EA-QD-CVaR minus the comparison; relative changes use the comparison mean as the denominator. Intervals are the paired mean-difference 95% BCa intervals; p denotes two-sided Wilcoxon signed-rank p values. Holm p denotes adjustment across all 54 main-comparison tests, as specified in Section 4.3.

Metric	Difference Δ 95% CI	Change (%)	p	Holm p
TTC	-0.970 $[-1.240, -0.650]$	-10.73	1.01×10^{-5}	0.000403
CVaR	-0.01310 $[-0.01762, -0.00749]$	-9.51	0.000103	0.003811
Joint risk	-3.792 $[-5.627, -1.418]$	-4.71	0.002255	0.069913

Relative to QD-CVaR, the Holm p values for TTC risk and CVaR are 0.000403 and 0.003811, respectively, both below 0.05. The Holm p value for joint risk is 0.069913.

4.6. Feedback-Path Ablations and Uniform-Proposal Comparison

The ablation unit is a complete feedback path comprising an observed signal and its associated parameter updates. The H_T path links efficiency weight, smoothness weight, and migration scale; H_R links safety weight, smoothness weight, and migration scale; p_R updates safety weight. Table 6 specifies the mapping. Statistical conclusions therefore assess the closed-loop effects of these paths and their combinations rather than an independent causal effect of any single entropy variable.

Table 12 compares four feedback ablations and Uniform, with 25 pairs per comparison. Relative to QD-HT, EA-QD-CVaR reduces mean TTC risk and CVaR by 9.54% and 7.34%, with $p = 0.000103$ and 0.001155, respectively. Relative to QD-pR-only, the reductions are 8.51% and 6.98%, with $p = 0.000250$ and 0.000715. Complete feedback achieves lower TTC risk and CVaR than these two single-signal configurations.

Relative to QD-HR-only, mean TTC risk and CVaR decrease by 6.77% and 6.40%, with $p = 0.004605$ and 0.008069. Both raw p values are below 0.05; the corresponding Holm p values are 0.138139 and 0.225926. Relative to QD-HR-pR, reductions are 4.35% and 3.76%, with $p = 0.027479$ and 0.101397. The raw p value is below 0.05 for TTC risk and above 0.05 for CVaR. The corresponding Holm p values are 0.714465 and 1.000000, respectively. Thus, complete feedback reduces both risk means relative to the single-path configurations and further reduces TTC risk relative to the two-risk-path configuration. These effects combine objective-weight adjustment, command-smoothness regulation, and changes in local proposal scale.

Uniform retains the adaptive objective, local candidate blocks, and 64-evaluation budget but proposes uniformly and accepts every proposal. Relative to Uniform, EA-QD-CVaR reduces mean TTC risk and CVaR by 6.22% and 6.06%, with $p = 0.006129$ and 0.031808. Both raw p values are below 0.05, with Holm p values of 0.177747 and 0.763388, respectively. Both risk comparisons against QD-HT and QD-pR-only remain below 0.05 after Holm adjustment.

Tables 10 and 12 show that complete EA-QD-CVaR achieves the best mean risk performance, with differing contributions from feedback signals and associated parameter updates. Among ablations, combined risk-entropy and risk-probability feedback in QD-HR-pR performs relatively well. Its mean TTC risk and CVaR are 8.436% and 0.12949, both lower than the corresponding means for configurations using risk entropy or risk

probability alone. Integrating traffic-state and risk information in the complete method further reduces these means to 8.069% and 0.12462. This result characterizes the mean-risk improvement associated with the complete combination of feedback paths.

Table 12. Feedback-ablation and uniform-proposal comparisons, with 25 pairs each. Δ is EA-QD-CVaR minus the comparison; relative changes use the comparison mean as the denominator. Intervals are the paired mean-difference 95% BCa intervals; p denotes two-sided Wilcoxon signed-rank p values. TTC-risk differences are percentage points. Holm p denotes adjustment across all 54 main-comparison tests.

Comparison Metric	Difference Δ 95% CI	Change (%)	p	Holm p
QD-HT TTC	−0.851 [−1.127, −0.480]	−9.54	0.000103	0.003811
QD-HT CVaR	−0.00988 [−0.01454, −0.00464]	−7.34	0.001155	0.036955
QD-HR-only TTC	−0.586 [−0.988, −0.249]	−6.77	0.004605	0.138139
QD-HR-only CVaR	−0.00852 [−0.01466, −0.00315]	−6.40	0.008069	0.225926
QD-pR-only TTC	−0.750 [−1.097, −0.428]	−8.51	0.000250	0.008493
QD-pR-only CVaR	−0.00936 [−0.01440, −0.00480]	−6.98	0.000715	0.023594
QD-HR-pR TTC	−0.367 [−0.693, −0.110]	−4.35	0.027479	0.714465
QD-HR-pR CVaR	−0.00487 [−0.00965, −0.00090]	−3.76	0.101397	1.000000
Uniform TTC	−0.535 [−0.909, −0.198]	−6.22	0.006129	0.177747
Uniform CVaR	−0.00804 [−0.01480, −0.00214]	−6.06	0.031808	0.763388

4.7. Classical Stochastic Optimization and Predictive-Control Comparisons

Table 13 reports the 25-pair risk comparisons against SA, Greedy-global, and MPC-CVaR. Under the same predictor, horizon, and evaluation budget, EA-QD-CVaR reduces mean TTC risk by 24.78%, 24.95%, and 20.76%, respectively. The corresponding CVaR reductions are 25.77%, 25.69%, and 21.31%. All six paired-difference intervals lie below zero, with $p = 5.960 \times 10^{-8}$ for each comparison.

Table 13. Risk comparisons with classical stochastic optimization and MPC, with 25 pairs each. Δ is EA-QD-CVaR minus the comparison; relative changes use the comparison mean as the denominator. Intervals are the paired mean-difference 95% BCa intervals; p denotes two-sided Wilcoxon signed-rank p values. TTC-risk differences are percentage points. Holm p denotes adjustment across all 54 main-comparison tests.

Comparison Metric	Difference Δ 95% CI	Change (%)	p	Holm p
SA TTC	−2.658 [−3.012, −2.269]	−24.78	5.96×10^{-8}	3.22×10^{-6}
SA CVaR	−0.04327 [−0.04870, −0.03682]	−25.77	5.96×10^{-8}	3.22×10^{-6}
Greedy-global TTC	−2.683 [−2.961, −2.438]	−24.95	5.96×10^{-8}	3.22×10^{-6}
Greedy-global CVaR	−0.04309 [−0.04696, −0.03932]	−25.69	5.96×10^{-8}	3.22×10^{-6}
MPC-CVaR TTC	−2.114 [−2.477, −1.707]	−20.76	5.96×10^{-8}	3.22×10^{-6}
MPC-CVaR CVaR	−0.03374 [−0.03958, −0.02769]	−21.31	5.96×10^{-8}	3.22×10^{-6}

All six risk comparisons have Holm $p = 3.22 \times 10^{-6}$ and remain below 0.05 after adjustment.

Greedy-local, No-Control, and ALINEA have reference TTC risks of 9.299%, 11.295%, and 11.830%; their CVaRs are 0.13867, 0.17728, and 0.18227, respectively. Each has an actual sample size of one.

4.8. Classical Search with Shared Adaptive Weighting and Proposal-Mechanism Comparisons

To compare closed-loop search performance under a shared adaptive objective, we extend SA, Greedy-global, and MPC-CVaR with the feedback used in EA-QD-CVaR. The three extensions use H_T, H_R, p_R, and the weight-update rule in Equation (21). They are denoted as EA-SA, EA-Greedy, and EA-MPC-CVaR, respectively. Each retains its original proposal, acceptance, or candidate-distribution update mechanism; EA-Greedy uses Greedy-global. Each method computes feedback and updates weights from its own closed-loop traffic state. The traffic scenario, demand and prediction random streams, 240 s horizon, constraints, objective components, and budget of 64 complete-sequence evaluations per cycle are shared. Weight offsets, feedback gains, and clipping bounds follow Table 4.

We additionally include Uniform-MH and Boltzmann, two local-search methods without quantum circuits. Both retain the nine-state local candidate set, Hamiltonian energy calculation, block-visitation order, and eight samples per complete local block. They also execute the first action of the lowest-energy candidate evaluated within each cycle. Table 14 summarizes the risk metrics, arrivals, and mean ramp waiting vehicles for the five additional methods and EA-QD-CVaR.

Table 14. Closed-loop performance of the additional search comparisons and the original EA-QD-CVaR. Each method has 25 trials; values are mean±sample SD. Lower risks and waiting counts are preferred, whereas higher arrivals are preferred.

Method	TTC Risk (%)	CVaR	Joint Risk (%)	Arrivals (Vehicles)	Mean Ramp Waiting Vehicles
EA-QD-CVaR	8.0689 ±0.5715	0.12462 ±0.00799	76.7123 ±3.9440	137.28 ±1.59	10.30 ±2.90
EA-SA	9.1172 ±0.9817	0.14082 ±0.01556	79.2592 ±4.2467	136.36 ±2.27	9.38 ±4.89
EA-Greedy	10.5407 ±0.6625	0.16513 ±0.00984	83.8201 ±2.1576	133.28 ±0.68	2.06 ±1.63
EA-MPC-CVaR	8.7076 ±0.5600	0.13586 ±0.00993	78.3417 ±2.6995	136.04 ±2.21	9.37 ±4.16
Uniform-MH	9.1025 ±0.7580	0.13913 ±0.01227	80.7446 ±2.8625	136.56 ±1.80	8.97 ±2.93
Boltzmann	9.1143 ±0.6361	0.13797 ±0.00998	81.1205 ±2.5449	137.40 ±2.35	10.39 ±4.47

All statistical tests in this subsection use the original EA-QD-CVaR as the reference. EA-SA, EA-Greedy, and EA-MPC-CVaR use new seeds different from the original EA trials. They are assessed using two-sided Wilcoxon rank-sum tests (Mann–Whitney U), with tie and continuity corrections. Uniform-MH and Boltzmann are paired with the original EA by search seed. Their two-sided Wilcoxon signed-rank tests follow the paired procedure in Section 4.3. In Table 15, p denotes the corresponding raw two-sided test value. The value p_{Holm} is adjusted across all 15 tests in this subsection, covering five methods and three risk metrics. Adjusted differences are evaluated using $p_{\text{Holm}} < 0.05$. Arrivals and mean ramp waiting vehicles are reported descriptively only.

Under the shared adaptive-weighting rule, EA-QD-CVaR reduces mean TTC risk by 11.50%, 23.45%, and 7.33% relative to EA-SA, EA-Greedy, and EA-MPC-CVaR, respec-

tively. The corresponding CVaR reductions are 11.51%, 24.54%, and 8.28%. All six risk comparisons have $p_{\text{Holm}} < 0.05$. The mean joint-risk reductions are 3.21%, 8.48%, and 2.08%, respectively. The EA-Greedy comparison remains significant after adjustment, with $p_{\text{Holm}} = 2.881 \times 10^{-5}$. The adjusted values for EA-SA and EA-MPC-CVaR are both 0.104690, above the 0.05 threshold. Thus, search strategies retain different TTC-risk and tail-loss outcomes under a shared feedback and weight-update rule.

Table 15. Risk tests comparing five additional methods with the original EA-QD-CVaR. Each comparison uses 25 versus 25 trials; adjustment covers all 15 tests in this table. The first three methods use independent-sample rank-sum tests, and the last two use paired signed-rank tests.

Comparison	Metric	p	p_{Holm}
EA-SA	TTC risk	1.671×10^{-4}	9.276×10^{-4}
EA-SA	CVaR	1.546×10^{-4}	9.276×10^{-4}
EA-SA	Joint risk	0.059826	0.104690
EA-Greedy	TTC risk	4.529×10^{-9}	6.341×10^{-8}
EA-Greedy	CVaR	2.828×10^{-9}	4.242×10^{-8}
EA-Greedy	Joint risk	2.619×10^{-6}	2.881×10^{-5}
EA-MPC-CVaR	TTC risk	2.646×10^{-4}	9.276×10^{-4}
EA-MPC-CVaR	CVaR	2.106×10^{-4}	9.276×10^{-4}
EA-MPC-CVaR	Joint risk	0.052345	0.104690
Uniform-MH	TTC risk	1.967×10^{-6}	2.360×10^{-5}
Uniform-MH	CVaR	3.278×10^{-6}	3.278×10^{-5}
Uniform-MH	Joint risk	3.195×10^{-5}	2.236×10^{-4}
Boltzmann	TTC risk	1.132×10^{-6}	1.472×10^{-5}
Boltzmann	CVaR	2.211×10^{-5}	1.769×10^{-4}
Boltzmann	Joint risk	1.234×10^{-5}	1.110×10^{-4}

Relative to Uniform-MH, EA-QD-CVaR reduces mean TTC risk, CVaR, and joint risk by 11.36%, 10.43%, and 4.99%, respectively. The corresponding reductions relative to Boltzmann are 11.47%, 9.68%, and 5.43%. All six paired comparisons remain significant at the 0.05 level after Holm adjustment. The two classical proposal configurations retain MH acceptance or energy-weighted sampling, respectively, while EA-QD-CVaR achieves lower means for all three risks. Given their weight settings, these results characterize the complete adaptive framework relative to fixed-objective classical proposal schemes.

EA-QD-CVaR averages 137.28 downstream arrivals, above the means for the three adaptive classical methods and Uniform-MH but slightly below the Boltzmann mean of 137.40 vehicles. Its mean ramp waiting count is 10.30 vehicles, higher than the means for the three adaptive classical methods and Uniform-MH and close to the Boltzmann mean of 10.39. These results reveal a trade-off among lower risk, higher downstream arrivals, and upstream ramp waiting. Overall, the additional comparisons support EA-QD-CVaR risk control in this dual-ramp scenario under a finite evaluation budget. Its performance reflects the overall framework of feedback weighting, candidate organization, and search-execution rules.

4.9. Joint Risk and Threshold Sensitivity

Table 16 presents the joint-risk results for both net-headway thresholds side by side. The main evaluation uses 3 s, whereas control feedback uses 1.5 s. Both measures count the union of TTC below 7 s and net headway below the respective threshold. A following pair satisfying both conditions is counted once, and both measures use the same 25 paired runs. Figure 8 shows the distribution of joint-risk-event proportions across the comparison methods, using a TTC threshold of 7 s and a net-headway threshold of 3 s.

At the 3 s threshold, mean joint risk decreases by 4.71%. The paired mean difference is -3.792 percentage points, with a 95% BCa interval of $[-5.627, -1.418]$ percentage points

and raw $p = 0.002255$. The Holm-adjusted value across the 54 main-comparison tests is 0.069913. At the 1.5 s threshold used for control feedback, mean joint risk decreases by only 0.15%. The paired mean difference is -0.076 percentage points, with a 95% BCa interval of $[-1.085, +1.117]$ percentage points and $p = 0.633827$. No significant difference is detected at this threshold. The magnitude of joint-risk improvement depends on the net-headway threshold and should be considered separately from TTC risk and CVaR. Relative to MPC-CVaR, mean joint risk at the 3 s threshold decreases by 6.74%, with raw $p = 1.830 \times 10^{-5}$.

Table 16. Joint risk of EA-QD-CVaR and fixed-weight QD-CVaR at two net-headway thresholds. Values are 25-run means, with a TTC threshold of 7 s in both measures. Relative changes use the QD-CVaR mean as the denominator; p denotes the raw two-sided Wilcoxon signed-rank value.

Net-Headway Threshold	QD-CVaR (%)	EA-QD-CVaR (%)	Relative Change (%)	p
3 s (main evaluation)	80.504	76.712	−4.71	0.002255
1.5 s (control feedback)	51.068	50.992	−0.15	0.633827

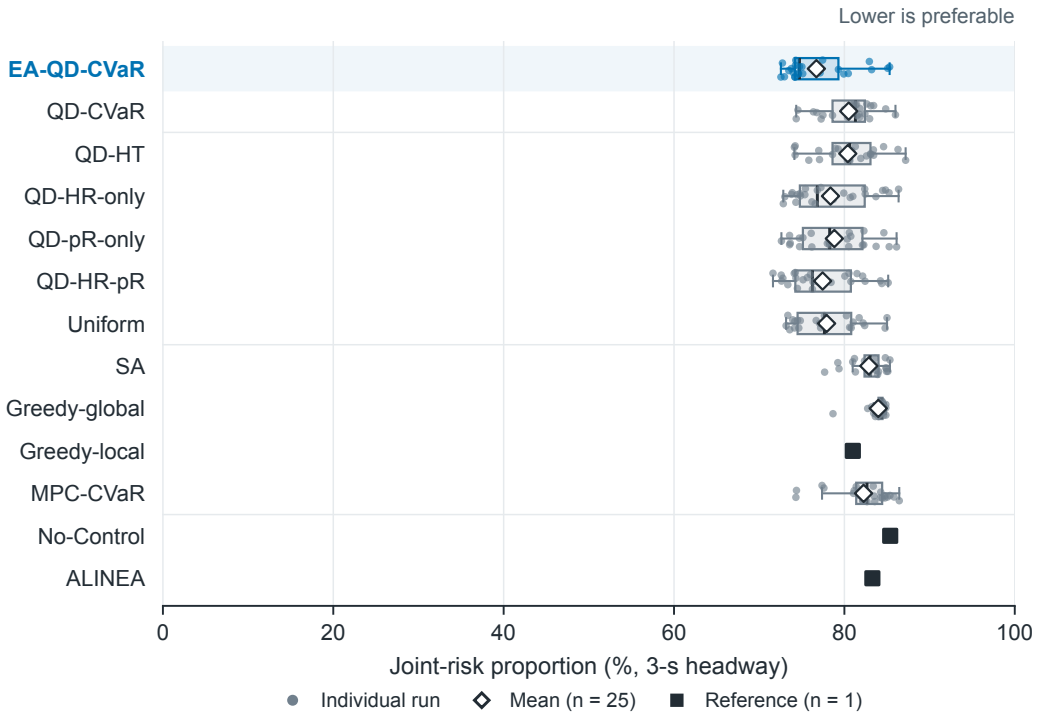


Figure 8. Joint-risk-event proportion in the dual-ramp scenario: TTC below 7 s or net headway below 3 s; lower is preferable. Each run averages 80 frames equally. Sample sizes and box-plot definitions follow Figure 6.

The union relationship $P(\mathcal{T} \cup \mathcal{H}) = P(\mathcal{T}) + P(\mathcal{H}) - P(\mathcal{T} \cap \mathcal{H})$ makes joint risk sensitive to both conflict proximity and short-headway exposure. Threshold sensitivity reflects the differing exposure ranges of the two-event union.

4.10. Traffic Operations and Command-Smoothness Costs

Table 17 reports arrivals, mean speed, and control variation. EA-QD-CVaR yields 137.28 ± 1.59 arrivals and QD-CVaR 136.64 ± 1.73. The paired difference is $+0.64$ vehicles, with a relative change of $+0.47$%, a 95% BCa interval of $[-0.20, +1.64]$ vehicles, and $p = 0.219830$. Among 25 pairs, EA-QD-CVaR has higher arrivals in 12, equal arrivals in 4, and lower arrivals in 9. The difference does not meet the 0.05 level. Figure 9 shows the distribution of downstream arrivals across the comparison methods.

EA-QD-CVaR's mean speed is 11.407 ± 0.275 m/s, a -0.75% change relative to QD-CVaR. The paired difference is -0.087 m/s, with a 95% BCa interval of $[-0.225, +0.071]$ m/s and $p = 0.241066$.

EA-QD-CVaR has $F_u = 0.28125 \pm 0.05980$, versus 0.19281 ± 0.04366 for QD-CVaR, an increase of 45.87%. The paired difference is $+0.08844$. Its 95% BCa interval is $[+0.06094, +0.11625]$, with $p = 6.199 \times 10^{-6}$. Lower TTC risk and CVaR accompany greater command variation. Table 18 summarizes the paired differences in arrivals, mean speed, and command variation between EA-QD-CVaR and QD-CVaR, together with the corresponding 95% confidence intervals and relative changes. Figure 10 shows the distribution of executed control variation across the comparison methods.

Table 17. Operating and command-variation metrics for 13 methods in the dual-ramp scenario. For stochastic methods ($n = 25$ each), upper entries are the mean±sample SD and lower entries are the mean 95% BCa intervals. The three deterministic references ($n = 1$ each) report single observations.

Method	*n*	Arrivals (Vehicles)	Speed (m/s)	F_u
EA-QD-CVaR	25	137.28 ± 1.59 [136.68, 137.92]	11.407 ± 0.275 [11.314, 11.528]	0.2812 ± 0.0598 [0.2584, 0.3044]
QD-CVaR	25	136.64 ± 1.73 [135.92, 137.24]	11.493 ± 0.208 [11.410, 11.570]	0.1928 ± 0.0437 [0.1762, 0.2097]
QD-HT	25	136.80 ± 1.85 [136.12, 137.56]	11.468 ± 0.283 [11.372, 11.591]	0.1988 ± 0.0402 [0.1847, 0.2159]
QD-HR-only	25	136.68 ± 2.21 [135.84, 137.52]	11.399 ± 0.313 [11.287, 11.529]	0.2447 ± 0.0425 [0.2291, 0.2616]
QD-pR-only	25	136.44 ± 2.08 [135.56, 137.16]	11.433 ± 0.275 [11.334, 11.545]	0.2906 ± 0.0572 [0.2697, 0.3138]
QD-HR-pR	25	136.64 ± 2.18 [135.76, 137.44]	11.379 ± 0.303 [11.267, 11.501]	0.3000 ± 0.0473 [0.2806, 0.3172]
Uniform	25	136.16 ± 1.65 [135.56, 136.80]	11.412 ± 0.280 [11.304, 11.520]	0.2859 ± 0.0620 [0.2591, 0.3072]
SA	25	134.24 ± 1.42 [133.76, 134.84]	11.308 ± 0.252 [11.219, 11.414]	0.0419 ± 0.0182 [0.0350, 0.0488]
Greedy-global	25	133.40 ± 0.82 [133.16, 133.80]	11.536 ± 0.319 [11.412, 11.656]	0.0209 ± 0.0087 [0.0175, 0.0241]
Greedy-local	1	136.00	11.255	0.1328
MPC-CVaR	25	135.92 ± 2.38 [135.16, 137.00]	11.394 ± 0.289 [11.291, 11.513]	0.0516 ± 0.0268 [0.0422, 0.0628]
No-Control	1	135.00	11.247	0.0000
ALINEA	1	135.00	11.140	0.0469

Table 18. Operating and smoothness costs of EA-QD-CVaR relative to QD-CVaR, with 25 pairs per metric. Difference units are vehicles, m/s, and dimensionless, respectively. Δ is EA-QD-CVaR minus the comparison; relative changes use the comparison mean as the denominator. Intervals are the paired mean-difference 95% BCa intervals.

Metric	Δ	Difference 95% CI	Change (%)
Arrivals	+0.64	[−0.20, +1.64]	+0.47
Mean speed	−0.087	[−0.225, +0.071]	−0.75
F_u	+0.0884	[+0.0609, +0.1163]	+45.87

The mean F_u is 0.04188 for SA, 0.02094 for Greedy-global, and 0.05156 for MPC-CVaR, all below EA-QD-CVaR, with all three $p = 5.960 \times 10^{-8}$. EA-QD-CVaR increases arrivals by 3.04 vehicles relative to SA and 3.88 relative to Greedy-global, with $p = 1.311 \times 10^{-6}$ and 5.960×10^{-8}, respectively. Relative to MPC-CVaR, arrivals increase by 1.36 vehicles, with $p = 0.040834$. Operating benefits and command-variation costs depend on the comparison method.

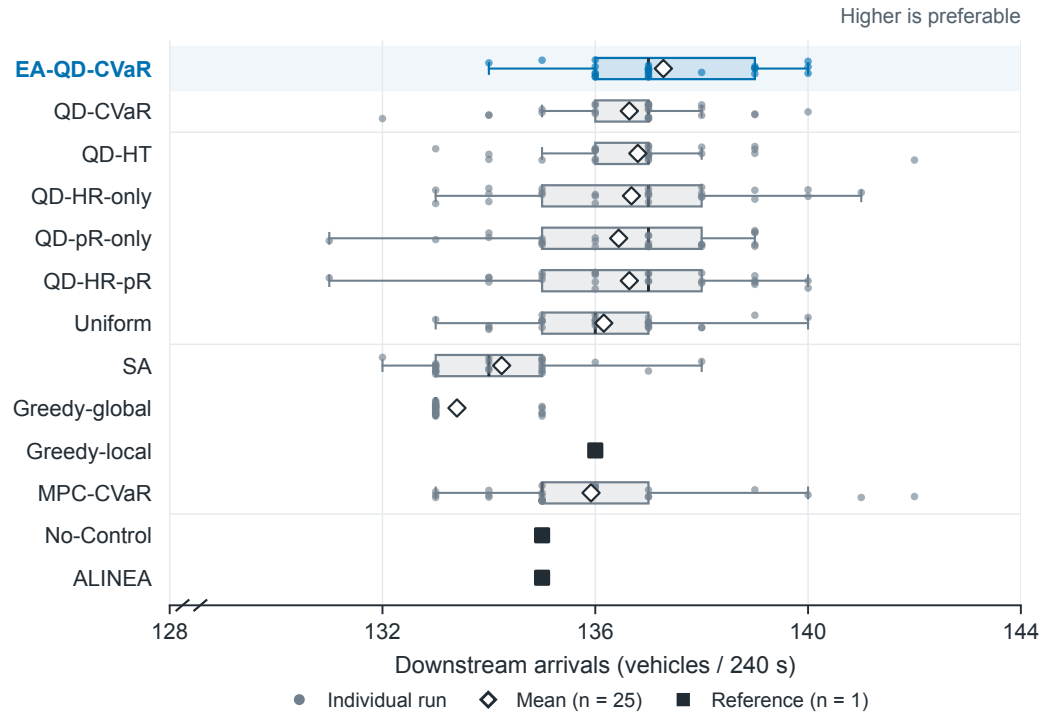


Figure 9. Downstream arrivals during the 240 s evaluation window; higher is preferable. The horizontal axis spans 128–144 vehicles, with a break marking the truncated origin. Sample sizes and box-plot definitions follow Figure 6.

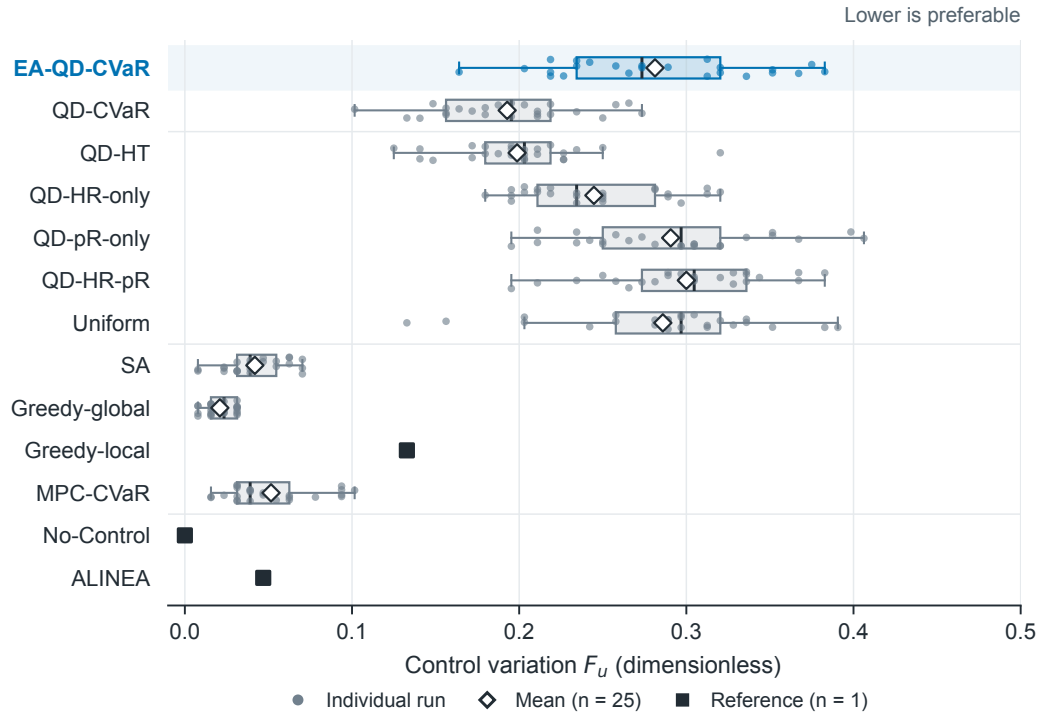


Figure 10. Executed control variation calculated using Equation (56). Lower values indicate smaller changes between adjacent commands. Sample sizes and box-plot definitions follow Figure 6.

Together, the risk and operational results show that EA-QD-CVaR reduces mean TTC risk by 10.73% relative to QD-CVaR. Mean joint risk at the 3 s net-headway threshold decreases by 4.71%. Mean downstream arrivals increase by 0.47%, while mean speed changes by −0.75%. Lower mean risks accompany slightly higher mean arrivals and a small change in mean speed. These outcomes support EA-QD-CVaR's effectiveness in

coordinating risk and operations in this dual-ramp scenario. The associated cost is higher F_u, reflecting more active command adjustment.

4.11. Interpretation of Entropy and Risk Feedback

H_T describes the dispersion across three speed bins, H_R describes the coexistence of four risk levels, p_R measures the event proportions, and CVaR measures upper-tail TTC loss. The three feedback signals alter candidate rankings and exploration paths through the objective weights, smoothness weight, and migration scale within the finite budget.

Complete feedback gives raw $p < 0.05$ for both TTC risk and CVaR relative to QD-HT, QD-pR-only, QD-HR-only, and Uniform. Relative to QD-HR-pR, only the raw TTC-risk p value is below 0.05. After Holm adjustment across 54 tests, both risk comparisons against QD-HT and QD-pR-only remain below 0.05. Both adjusted risk comparisons against each of the other three configurations exceed 0.05. Together with lower risk and higher F_u relative to fixed-parameter QD-CVaR, these results characterize the combined feedback modules' risk–command-variation trade-off in this scenario.

5. Discussion and Conclusions

5.1. Effects of Multiple Feedback Paths and Risk Trade-Offs

EA-QD-CVaR's main advantage is traffic-risk reduction under a finite candidate-evaluation budget. In 25 independent search trials in the dual-ramp scenario, it achieves the lowest mean TTC risk, CVaR, and joint risk (at the 3 s net-headway threshold) among ten stochastic methods and the highest mean arrivals. Relative to fixed-weight QD-CVaR, mean TTC-event proportion and CVaR decrease by 10.73% and 9.51%, with lower values in 21 and 20 paired trials, respectively. Both paired-difference 95% BCa intervals lie below zero, with $p = 1.007 \times 10^{-5}$ and 0.000103. Adaptive feedback therefore reduces both TTC-event exposure and upper-tail TTC loss in this scenario.

Relative to SA, Greedy-global, and MPC-CVaR, mean TTC risk decreases by 24.78%, 24.95%, and 20.76%, and mean CVaR by 25.77%, 25.69%, and 21.31%, respectively. All six differences have $p = 5.960 \times 10^{-8}$. With the same predictor, horizon, and evaluation budget, these results demonstrate the control utility of risk-sensitive search. Traffic-state entropy, risk entropy, and event proportion feed information on speed distributions, risk categories, and exposure into the objective weights, smoothness weight, and circuit migration scale. Control objectives and local exploration consequently update with observed conditions. Both risk comparisons against QD-HT and QD-pR-only also have $p < 0.05$, supporting the combined feedback paths. This effect includes the joint parameter updates in Table 6.

At the 3 s threshold, mean joint risk decreases by 4.71% relative to QD-CVaR ($p = 0.002255$) and 6.74% relative to MPC-CVaR ($p = 1.830 \times 10^{-5}$). For the comparison with QD-CVaR, the 3 s result has Holm $p = 0.069913$. At the 1.5 s net-headway threshold used for control feedback, mean joint risk decreases by only 0.15% ($p = 0.633827$). Joint risk supplements the individual measures by describing the union of conflict proximity and short-headway exposure. Section 4.9 reports both headway thresholds. Mean downstream arrivals are 137.28 vehicles, exceeding the QD-CVaR mean by 0.64 vehicles, or 0.47%, with a paired-difference 95% interval of $[-0.20, 1.64]$ vehicles and $p = 0.219830$. Mean speed changes by -0.75% ($p = 0.241066$). Relative to SA and Greedy-global, arrivals increase by 3.04 and 3.88 vehicles, with $p = 1.311 \times 10^{-6}$ and 5.960×10^{-8}, respectively. The method therefore improves vehicle output relative to these two methods alongside risk reduction.

Section 4.4 further examines fixed-safety-weight risk–efficiency trade-offs using 25 trials per setting. The empirical frontier of the tested fixed-weight points comprises

$\beta = 2, 4, 6, 8$. Relative to $\beta = 6$, which has the lowest fixed-weight mean risks, EA-QD-CVaR reduces TTC risk and CVaR by 5.37% and 3.85%, respectively. It also increases arrivals by 1.36 vehicles. Its mean arrivals exceed those at $\beta = 2$, the highest-arrival fixed setting, by 0.24 vehicles. Within the tested range, this combination of means indicates risk improvement alongside relatively high downstream throughput.

These risk improvements incur greater command variation: F_u increases from QD-CVaR's 0.19281 to 0.28125, a rise of 45.87% ($p = 6.199 \times 10^{-6}$). EA-QD-CVaR uses traffic and risk feedback to reduce TTC events and tail loss within a fixed evaluation budget while attaining higher mean arrivals. Engineering applications can choose this risk–control-variation trade-off according to actuator adjustment frequency and command-smoothness requirements.

5.2. Theoretical Convergence and Experimental Validation of Local Search

The analysis provides an explicit convergence target and guarantee for EA-QD-CVaR's fixed local probability search. With full-support proposals and finite candidate energies, a local MH kernel is irreducible and aperiodic, with a unique Gibbs stationary distribution. It converges asymptotically to that distribution from any initial distribution. Full support preserves access to every local candidate, while Hastings acceptance aligns probability migration with the energy-defined target, providing an analyzable basis for local exploration.

Section 3.5 evaluates distributional distances at different transition counts using 560 local kernels from five independent 16-step diagnostic trials. Mean TV distances at steps 8 and 16 are 0.08851 and 0.04517, respectively, a further decrease of 48.96%. The proportions of blocks satisfying TV ≤ 0.01 are 23.57% and 51.79%, respectively. Mean distributional distance continues to decrease in all five trials as transitions increase, indicating a closer approach to the target for these fixed kernels. These finite-step results complement the asymptotic properties in describing local-search probability evolution. Most diagnostic blocks nevertheless remain above the reference precision under the eight-step budget.

Actual rolling control uses eight MH transitions per complete local block and 64 complete-sequence evaluations per cycle. Transition states update subsequent search centers, while the controller executes the first action of the lowest-energy sequence evaluated in that cycle. The 25-trial closed-loop main comparison reports risk, traffic output, and command variation under this actual budget. Fixed-kernel analysis, finite-step diagnostics, and closed-loop trials characterize asymptotic properties, finite-budget distributional error, and control utility, respectively.

5.3. Applicability and Future Work

Statistics from the 25 independent trials characterize search variability in this scenario. The three deterministic controls each contribute one observation, and mean intervals are pointwise. Conclusions apply to the evaluated dual-ramp simulation, driving composition, and 240 s evaluation window. Other demand realizations, real networks, long-term queues, and actual crash risk require further validation.

This study evaluates risk and traffic operations under offline synchronous closed-loop simulation. Per-cycle timings in Section 4.2 quantify the computational requirements of the current implementation. Online deployment requires profiling computational modules and evaluating state-restoration reuse, parallel candidate evaluation, and simplified prediction models. End-to-end latency and its control effects should then be tested under specified hardware and load conditions. Further extensions include fixed-weight and feedback-gain comparisons, safety–efficiency frontier analysis, direct MH-step comparisons under identical traffic inputs, and command-variation constraints.

5.4. Conclusions

We propose EA-QD-CVaR, a cooperative freeway dual-ramp control method driven by entropy and risk feedback. Speed distributions, risk categories, event exposure, and tail TTC loss inform dynamic objective and search-parameter updates. The method coordinates two VSLs and two ramp release rates within 64 complete-candidate evaluations per cycle. Its advantage is feedback-driven adjustment of control priorities under a finite budget, reducing TTC-event exposure and high-loss following risk.

Across 25 independent search trials in the dual-ramp scenario, EA-QD-CVaR achieves the lowest mean TTC risk and CVaR among ten stochastic methods. It also achieves the lowest mean joint risk at the 3 s net-headway threshold and the highest mean arrivals. Relative to fixed-weight QD-CVaR, mean TTC risk and CVaR decrease by 10.73% and 9.51%, with raw $p = 1.007 \times 10^{-5}$ and 0.000103, respectively. Mean joint risk decreases by 4.71% at the 3 s net-headway threshold (raw $p = 0.002255$, Holm $p = 0.069913$). At the 1.5 s threshold used for control feedback, the decrease is only 0.15% ($p = 0.633827$). Mean arrivals increase by 0.47% ($p = 0.219830$). Relative to SA, Greedy-global, and MPC-CVaR, mean TTC risk decreases by 20.76–24.95%, and mean CVaR by 21.31–25.77%. These gains involve more active regulation: F_u increases by 45.87% relative to QD-CVaR, quantifying the command-variation cost of risk reduction.

The fixed-safety-weight scan, with 25 trials per group, further shows lower mean TTC risk and CVaR for EA-QD-CVaR than for all five fixed settings. Mean arrivals are 0.24–1.36 vehicles higher. This combination of means extends beyond the empirical safety–efficiency trade-off of the fixed-weight points tested here.

Fixed-local-kernel analysis establishes asymptotic convergence properties, while five independent diagnostic trials characterize finite-step distributional error at steps 8 and 16. The 25-trial closed-loop main comparison quantifies control performance under the actual eight-step budget. This study provides a feedback-adaptive method with an explicit evaluation budget and supporting simulation evidence. The findings inform dual-ramp coordination, safety-parameter selection, and subsequent deployment validation under a downstream bottleneck.

Author Contributions: Conceptualization, K.C. and Y.J.; methodology, K.C. and Y.J.; software, Y.J., K.Z., and Y.T.; validation, Y.J., K.Z., and Z.L.; formal analysis, K.C. and H.L.; investigation, Y.J. and H.L.; resources, Z.L. and X.Z.; data curation, Y.J., K.Z., and Z.L.; writing—original draft preparation, K.C.; writing—review and editing, Y.J., Y.T., and X.Z.; visualization, Y.J., H.L., and Y.T.; supervision, X.Z.; project administration, K.C. and X.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the Southeast University–Sichuan Digital Joint Research Center for Intelligent Transportation, Southeast University and Sichuan Digital Transportation Technology Co., Ltd.

Institutional Review Board Statement: Not applicable. This study used traffic simulation and secondary, publicly available vehicle-trajectory data and did not involve human participants or animals.

Informed Consent Statement: Not applicable.

Data Availability Statement: Traffic-demand scaling is informed by the public NGSIM I–80 data, and vehicle trajectories are generated by SUMO. The implementation code for this study will be released in a public code repository upon acceptance of the manuscript.

Conflicts of Interest: K.C., K.Z., Y.J., H.L., Z.L., and Y.T. are employed by Sichuan Digital Transportation Technology Co., Ltd. and Sichuan Shudao Intelligent Transportation Group Co., Ltd. Their employment is disclosed for transparency. The authors declare that this study received funding

from the Southeast University–Sichuan Digital Joint Research Center for Intelligent Transportation, Southeast University and Sichuan Digital Transportation Technology Co., Ltd. The funders were involved in the theoretical research, algorithm design, data analysis, and writing of the manuscript.

Abbreviations

The following abbreviations are used in this manuscript:

CVaR	Conditional value-at-risk
EA-QD-CVaR	Entropy-driven adaptive quantum-inspired optimization with CVaR
KDE	Kernel density estimation
MH	Metropolis–Hastings
MPC	Model predictive control
NGSIM	Next-Generation Simulation
QD-CVaR	Fixed-weight quantum-inspired optimization with CVaR
SUMO	Simulation of Urban Mobility
TTC	Time to collision
VSL	Variable speed limit

References

1. Lighthill, M.J.; Whitham, G.B. On kinematic waves II. A theory of traffic flow on long crowded roads. *Proc. R. Soc. London. Ser. A. Math. Phys. Sci.* **1955**, *229*, 317–345. https://doi.org/10.1098/rspa.1955.0089.
2. Richards, P.I. Shock waves on the highway. *Oper. Res.* **1956**, *4*, 42–51. https://doi.org/10.1287/opre.4.1.42.
3. Daganzo, C.F. The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory. *Transp. Res. Part B: Methodol.* **1994**, *28*, 269–287. https://doi.org/10.1016/0191-2615(94)90002-7.
4. Aw, A.; Rascle, M. Resurrection of "second order" models of traffic flow. *SIAM J. Appl. Math.* **2000**, *60*, 916–938. https://doi.org/10.1137/s0036139997332099.
5. Helbing, D. Traffic and related self-driven many-particle systems. *Rev. Mod. Phys.* **2001**, *73*, 1067–1141. https://doi.org/10.1103/revmodphys.73.1067.
6. Treiber, M.; Hennecke, A.; Helbing, D. Congested traffic states in empirical observations and microscopic simulations. *Phys. Rev. E* **2000**, *62*, 1805–1824. https://doi.org/10.1103/physreve.62.1805.
7. Treiber, M.; Helbing, D. Memory effects in microscopic traffic models and wide scattering in flow-density data. *Phys. Rev. E* **2003**, *68*, 046119. https://doi.org/10.1103/physreve.68.046119.
8. Nagel, K.; Schreckenberg, M. A cellular automaton model for freeway traffic. *J. Phys. I* **1992**, *2*, 2221–2229. https://doi.org/10.1051/jp1:1992277.
9. Kerner, B.S.; Rehborn, H. Experimental properties of complexity in traffic flow. *Phys. Rev. E* **1996**, *53*, R4275–R4278. https://doi.org/10.1103/physreve.53.r4275.
10. Kerner, B.S.; Rehborn, H. Experimental properties of phase transitions in traffic flow. *Phys. Rev. Lett.* **1997**, *79*, 4030–4033. https://doi.org/10.1103/physrevlett.79.4030.
11. Sugiyama, Y.; Fukui, M.; Kikuchi, M.; Hasebe, K.; Nakayama, A.; Nishinari, K.; Tadaki, S.i.; Yukawa, S. Traffic jams without bottlenecks-experimental evidence for the physical mechanism of the formation of a jam. *New J. Phys.* **2008**, *10*, 033001. https://doi.org/10.1088/1367-2630/10/3/033001.
12. Tadaki, S.i.; Kikuchi, M. Jam phases in a two-dimensional cellular-automaton model of traffic flow. *Phys. Rev. E* **1994**, *50*, 4564–4570. https://doi.org/10.1103/physreve.50.4564.
13. Bando, M.; Hasebe, K.; Nakayama, A.; Shibata, A.; Sugiyama, Y. Dynamical model of traffic congestion and numerical simulation. *Phys. Rev. E* **1995**, *51*, 1035–1042. https://doi.org/10.1103/physreve.51.1035.
14. Papageorgiou, M.; Hadj-Salem, H.; Blosseville, J.M. ALINEA: A local feedback control law for on-ramp metering. *Transp. Res. Rec.* **1991**, *1320*, 58–64.
15. Papageorgiou, M.; Kiakaki, C.; Dinopoulou, V.; Kotsialos, A.; Wang, Y. Review of road traffic control strategies. *Proc. IEEE* **2003**, *91*, 2043–2067. https://doi.org/10.1109/jproc.2003.819610.
16. Hegyi, A.; De Schutter, B.; Hellendoorn, H. Model predictive control for optimal coordination of ramp metering and variable speed limits. *Transp. Res. Part C Emerg. Technol.* **2005**, *13*, 185–209. https://doi.org/10.1016/j.trc.2004.08.001.
17. Hegyi, A.; De Schutter, B.; Hellendoorn, J. Optimal coordination of variable speed limits to suppress shock waves. *IEEE Trans. Intell. Transp. Syst.* **2005**, *6*, 102–112. https://doi.org/10.1109/tits.2004.842408.
18. Kotsialos, A.; Papageorgiou, M.; Mangeas, M.; Haj-Salem, H. Coordinated and integrated control of motorway networks via non-linear optimal control. *Transp. Res. Part C Emerg. Technol.* **2002**, *10*, 65–84. https://doi.org/10.1016/s0968-090x(01)00005-5.

19. Carlson, R.C.; Papamichail, I.; Papageorgiou, M.; Messmer, A. Optimal motorway traffic flow control involving variable speed limits and ramp metering. *Transp. Sci.* **2010**, *44*, 238–253. https://doi.org/10.1287/trsc.1090.0314.
20. Papamichail, I.; Kotsialos, A.; Margonis, I.; Papageorgiou, M. Coordinated ramp metering for freeway networks – a model-predictive hierarchical control approach. *Transp. Res. Part C Emerg. Technol.* **2010**, *18*, 311–331. https://doi.org/10.1016/j.trc.2008.11.002.
21. Frejo, J.R.D.; Camacho, E.F. Global versus local MPC algorithms in freeway traffic control with ramp metering and variable speed limits. *IEEE Trans. Intell. Transp. Syst.* **2012**, *13*, 1556–1565. https://doi.org/10.1109/tits.2012.2195493.
22. Lu, X.Y.; Varaiya, P.; Horowitz, R.; Su, D.; Shladover, S.E. Novel freeway traffic control with variable speed limit and coordinated ramp metering. *Transp. Res. Rec. J. Transp. Res. Board* **2011**, *2229*, 55–65. https://doi.org/10.3141/2229-07.
23. Ghods, A.; Kian, A.; Tabibi, M. Adaptive freeway ramp metering and variable speed limit control: A genetic-fuzzy approach. *IEEE Intell. Transp. Syst. Mag.* **2009**, *1*, 27–36. https://doi.org/10.1109/mits.2009.932718.
24. Hou, Z.; Xu, J.X.; Zhong, H. Freeway traffic control using iterative learning control-based ramp metering and speed signaling. *IEEE Trans. Veh. Technol.* **2007**, *56*, 466–477. https://doi.org/10.1109/tvt.2007.891431.
25. Vrbanić, F.; Ivanjko, E.; Kušić, K.; Čakija, D. Variable speed limit and ramp metering for mixed traffic flows: A review and open questions. *Appl. Sci.* **2021**, *11*, 2574. https://doi.org/10.3390/app11062574.
26. Zheng, S.; Li, Z.; Li, M.; Ke, Z. Enhancing reinforcement learning-based ramp metering performance at freeway uncertain bottlenecks using curriculum learning. *IET Intell. Transp. Syst.* **2024**, *18*, 1863–1878. https://doi.org/10.1049/itr2.12494.
27. Jin, J.; Li, Y.; Huang, H.; Dong, Y.; Liu, P. A variable speed limit control approach for freeway tunnels based on the model-based reinforcement learning framework with safety perception. *Accid. Anal. Prev.* **2024**, *201*, 107570. https://doi.org/10.1016/j.aap.2024.107570.
28. He, Z.; Wang, L.; Su, Z.; Ma, W. Integrating variable speed limit and ramp metering to enhance vehicle group safety and efficiency in a mixed traffic environment. *Phys. A Stat. Mech. Its Appl.* **2024**, *641*, 129754. https://doi.org/10.1016/j.physa.2024.129754.
29. Minderhoud, M.M.; Bovy, P.H.L. Extended time-to-collision measures for road traffic safety assessment. *Accid. Anal. Prev.* **2001**, *33*, 89–97. https://doi.org/10.1016/s0001-4575(00)00019-1.
30. Laureshyn, A.; Svensson, Å.; Hydén, C. Evaluation of traffic safety, based on micro-level behavioural data: Theoretical framework and first implementation. *Accid. Anal. Prev.* **2010**, *42*, 1637–1646. https://doi.org/10.1016/j.aap.2010.03.021.
31. Mahmud, S.M.S.; Ferreira, L.; Hoque, M.S.; Tavassoli, A. Application of proximal surrogate indicators for safety evaluation: A review of recent developments and research needs. *IATSS Res.* **2017**, *41*, 153–163. https://doi.org/10.1016/j.iatssr.2017.02.001.
32. Johnsson, C.; Laureshyn, A.; De Ceunynck, T. In search of surrogate safety indicators for vulnerable road users: A review of surrogate safety indicators. *Transp. Rev.* **2018**, *38*, 765–785. https://doi.org/10.1080/01441647.2018.1442888.
33. Ozbay, K.; Yang, H.; Bartin, B.; Mudigonda, S. Derivation and validation of new simulation-based surrogate safety measure. *Transp. Res. Rec. J. Transp. Res. Board* **2008**, *2083*, 105–113. https://doi.org/10.3141/2083-12.
34. Saunier, N.; Sayed, T. Probabilistic framework for automated analysis of exposure to road collisions. *Transp. Res. Rec. J. Transp. Res. Board* **2008**, *2083*, 96–104. https://doi.org/10.3141/2083-11.
35. Singh, D.; Das, P.; Ghosh, I. Surrogate safety assessment of traffic facilities under ordered and disordered traffic condition: Systematic literature review. *KSCE J. Civ. Eng.* **2023**, *27*, 5008–5029. https://doi.org/10.1007/s12205-023-0979-y.
36. Venthuruthiyil, S.P.; Chunchu, M. Anticipated collision time (ACT): A two-dimensional surrogate safety indicator for trajectory-based proactive safety assessment. *Transp. Res. Part C Emerg. Technol.* **2022**, *139*, 103655. https://doi.org/10.1016/j.trc.2022.103655.
37. Shannon, C.E. A mathematical theory of communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x.
38. Bandt, C.; Pompe, B. Permutation entropy: A natural complexity measure for time series. *Phys. Rev. Lett.* **2002**, *88*, 174102. https://doi.org/10.1103/physrevlett.88.174102.
39. Pincus, S.M. Approximate entropy as a measure of system complexity. *Proc. Natl. Acad. Sci. USA* **1991**, *88*, 2297–2301. https://doi.org/10.1073/pnas.88.6.2297.
40. Costa, M.; Goldberger, A.L.; Peng, C.K. Multiscale entropy analysis of complex physiologic time series. *Phys. Rev. Lett.* **2002**, *89*, 068102. https://doi.org/10.1103/physrevlett.89.068102.
41. Rosso, O.A.; Larrondo, H.A.; Martin, M.T.; Plastino, A.; Fuentes, M.A. Distinguishing noise from chaos. *Phys. Rev. Lett.* **2007**, *99*, 154102. https://doi.org/10.1103/physrevlett.99.154102.
42. Liao, G.; Shang, P. Scaling and complexity-entropy analysis in discriminating traffic dynamics. *Fractals* **2012**, *20*, 233–243. https://doi.org/10.1142/s0218348x12500211.
43. Zhang, Y.; Shang, P. The complexity–entropy causality plane based on multivariate multiscale distribution entropy of traffic time series. *Nonlinear Dyn.* **2019**, *95*, 617–629. https://doi.org/10.1007/s11071-018-4586-2.
44. Velasco, R.M. The informational entropy in traffic flow. In *Proceedings of the AIP Conference Proceedings*; AIP: Melville, NY, USA, 2005; Volume 757; pp. 200–206. https://doi.org/10.1063/1.1900501.

45. Dolbeault, J.; Illner, R. Entropy methods for kinetic models of traffic flow. *Commun. Math. Sci.* **2003**, *1*, 409–421. https://doi.org/10.4310/cms.2003.v1.n3.a2.
46. Knowles, J.K. On entropy conditions and traffic flow models. *ZAMM J. Appl. Math. Mech./Z. Für Angew. Math. Und Mech.* **2008**, *88*, 64–73. https://doi.org/10.1002/zamm.200700093.
47. Zubillaga, D.; Cruz, G.; Aguilar, L.D.; Zapotécatl, J.; Fernández, N.; Aguilar, J.; Rosenblueth, D.A.; Gershenson, C. Measuring the complexity of self-organizing traffic lights. *Entropy* **2014**, *16*, 2384–2407. https://doi.org/10.3390/e16052384.
48. Mehdi, M.Z.; Kammoun, H.M.; Benayed, N.G.; Sellami, D.; Masmoudi, A.D. Entropy-based traffic flow labeling for CNN-based traffic congestion prediction from meta-parameters. *IEEE Access* **2022**, *10*, 16123–16133. https://doi.org/10.1109/ACCESS.2022.3149059.
49. Huang, P.; Ding, H.; Chen, H. An entropy-based model for quantifying multi-dimensional traffic scenario complexity. *IET Intell. Transp. Syst.* **2024**, *18*, 1289–1305. https://doi.org/10.1049/itr2.12510.
50. Finnila, A.B.; Gomez, M.A.; Sebenik, C.; Stenson, C.; Doll, J.D. Quantum annealing: A new method for minimizing multidimensional functions. *Chem. Phys. Lett.* **1994**, *219*, 343–348. https://doi.org/10.1016/0009-2614(94)00117-0.
51. Kadowaki, T.; Nishimori, H. Quantum annealing in the transverse Ising model. *Phys. Rev. E* **1998**, *58*, 5355–5363. https://doi.org/10.1103/physreve.58.5355.
52. Farhi, E.; Goldstone, J.; Gutmann, S.; Lapan, J.; Lundgren, A.; Preda, D. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science* **2001**, *292*, 472–475. https://doi.org/10.1126/science.1057726.
53. Farhi, E.; Goldstone, J.; Gutmann, S. A quantum approximate optimization algorithm. *arXiv* 2014, arXiv.1411.4028. https://doi.org/10.48550/arXiv.1411.4028.
54. Hadfield, S.; Wang, Z.; O'Gorman, B.; Rieffel, E.G.; Venturelli, D.; Biswas, R. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms* **2019**, *12*, 34. https://doi.org/10.3390/a12020034.
55. McClean, J.R.; Romero, J.; Babbush, R.; Aspuru-Guzik, A. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* **2016**, *18*, 023023. https://doi.org/10.1088/1367-2630/18/2/023023.
56. Cerezo, M.; Arrasmith, A.; Babbush, R.; Benjamin, S.C.; Endo, S.; Fujii, K.; McClean, J.R.; Mitarai, K.; Yuan, X.; Cincio, L.; et al. Variational quantum algorithms. *Nat. Rev. Phys.* **2021**, *3*, 625–644. https://doi.org/10.1038/s42254-021-00348-9.
57. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2018**, *2*, 79. https://doi.org/10.22331/q-2018-08-06-79.
58. Wang, P.; Chen, Y.; Xin, G.; Jiao, Y.; Yin, X.; Yang, G.; Zhou, Y.; Mu, L.; Wang, F. A brief study on quantum dynamics of optimization algorithm. *J. Southwest Minzu Univ. (Nat. Sci. Ed.)* **2021**, *47*, 288–296. https://doi.org/10.11920/xnmdzk.2021.03.008.
59. Jiao, Y.W.; Wang, P. Multi-scale quantum harmonic oscillator algorithm based on subpopulation number adaptive. *Acta Autom. Sin.* **2023**, *49*, 1587–1600. https://doi.org/10.16383/j.aas.c200247.
60. Neukart, F.; Compostella, G.; Seidel, C.; von Dollen, D.; Yarkoni, S.; Parney, B. Traffic flow optimization using a quantum annealer. *Front. ICT* **2017**, *4*, 29. https://doi.org/10.3389/fict.2017.00029.
61. Hussain, H.; Javaid, M.B.; Khan, F.S.; Dalal, A.; Khalique, A. Optimal control of traffic signals using quantum annealing. *Quantum Inf. Process.* **2020**, *19*, 312. https://doi.org/10.1007/s11128-020-02815-1.
62. Inoue, D.; Okada, A.; Matsumori, T.; Aihara, K.; Yoshida, H. Traffic signal optimization on a square lattice with quantum annealing. *Sci. Rep.* **2021**, *11*, 3303. https://doi.org/10.1038/s41598-021-82740-0.
63. Stollenwerk, T.; O'Gorman, B.; Venturelli, D.; Mandrà, S.; Rodionova, O.; Ng, H.K.; Sridhar, B.; Rieffel, E.G.; Biswas, R. Quantum annealing applied to de-conflicting optimal trajectories for air traffic management. *IEEE Trans. Intell. Transp. Syst.* **2020**, *21*, 285–297. https://doi.org/10.1109/tits.2019.2891235.
64. Zhuang, Y.; Azfar, T.; Wang, Y.; Sun, W.; Wang, X.; Guo, Q.; Ke, R. Quantum computing in intelligent transportation systems: A survey. *Chain* **2024**, *1*, 138–149. https://doi.org/10.23919/CHAIN.2024.000007.
65. Liu, P.; Parkinson, S.; Best, K. Quantum and quantum-inspired optimisation in transport and logistics: A systematic review. *Smart Cities* **2025**, *8*, 206. https://doi.org/10.3390/smartcities8060206.
66. Yue, W.; Yang, H.; Wang, Y.; Zhou, Y.; Liu, G.; Zheng, P. A novel max-pressure-driven integrated ramp metering and variable speed limit control for port motorways. *Sustainability* **2025**, *17*, 5592. https://doi.org/10.3390/su17125592.
67. Zhang, Y.; Lei, Z.; Shi, Y.; Wang, Y.; Yin, X. A CVaR-based optimal perimeter control framework with consideration of boundary queue length. *J. Adv. Transp.* **2026**, *2026*, 7280111. https://doi.org/10.1155/atr/7280111.
68. Salloum, H.; Zhanalin, S.; Al Badr, A.; Kholodov, Y. Mini-scale traffic flow optimization: An iterative QUBOs approach converting from hybrid solver to pure quantum processing unit. *Sci. Rep.* **2025**, *15*, 22904. https://doi.org/10.1038/s41598-025-04568-2.
69. Hastings, W.K. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **1970**, *57*, 97–109. https://doi.org/10.1093/biomet/57.1.97.
70. Yang, G.; Xu, H.; Wang, Z.; Tian, Z. Truck acceleration behavior study and acceleration lane length recommendations for metered on-ramps. *Int. J. Transp. Sci. Technol.* **2016**, *5*, 93–102. https://doi.org/10.1016/j.ijtst.2016.09.006.
71. Lu, Y.; Cheng, K.; Zhang, Y.; Chen, X.; Zou, Y. Analysis of lane-changing conflict between cars and trucks at freeway merging sections using UAV video data. *J. Transp. Saf. Secur.* **2023**, *15*, 943–961. https://doi.org/10.1080/19439962.2022.2147611.

72. Chu, T.D.; Miwa, T.; Morikawa, T. Discrete choice models for gap acceptance at urban expressway merge sections considering safety, road geometry, and traffic conditions. *J. Transp. Eng. Part A Syst.* **2017**, *143*, 04017025. https://doi.org/10.1061/JTEPBS.0000053.
73. Zhang, S.; Xing, Y.; Lu, J.; Zhang, H.M. Exploring the influence of truck proportion on freeway traffic safety using adaptive network-based fuzzy inference system. *J. Adv. Transp.* **2019**, *2019*, 3879385. https://doi.org/10.1155/2019/3879385.
74. Chang, J. *Summary of NHTSA Heavy-Vehicle Vehicle-to-Vehicle Safety Communications Research*; Technical Report DOT HS 812 300; National Highway Traffic Safety Administration: Washington, DC, USA, 2016.
75. Rockafellar, R.T.; Uryasev, S. Optimization of conditional value-at-risk. *J. Risk* **2000**, *2*, 21–41. https://doi.org/10.21314/JOR.2000.038.
76. Barkoutsos, P.K.; Nannicini, G.; Robert, A.; Tavernelli, I.; Woerner, S. Improving variational quantum optimization using CVaR. *Quantum* **2020**, *4*, 256. https://doi.org/10.22331/q-2020-04-20-256.
77. Tierney, L. Markov chains for exploring posterior distributions. *Ann. Stat.* **1994**, *22*, 1701–1728. https://doi.org/10.1214/aos/1176325750.
78. Mengersen, K.L.; Tweedie, R.L. Rates of convergence of the Hastings and Metropolis algorithms. *Ann. Stat.* **1996**, *24*, 101–121. https://doi.org/10.1214/aos/1033066201.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.