REVS-T: Trust-Tier-Aware Provider Selection for Secure Vehicular Computation Offloading

Sharifah Fayi*†, Ferheen Ayaz‡, Zhengguo Sheng*
*School of Engineering and Informatics, University of Sussex, Brighton, United Kingdom
†College of Computer Science, King Khalid University, Abha, Saudi Arabia
‡Department of Computer Science, City St. George's, University of London, London, United Kingdom
Email: s.fayi@sussex.ac.uk, ferheen.ayaz@city.ac.uk, z.sheng@sussex.ac.uk

***Abstract*—Vehicular computation offloading (VCOff) enables resource-constrained vehicles to delegate delay-sensitive tasks to nearby providers. However, it remains vulnerable to strategic malicious nodes that withhold results, or behave intermittently to evade detection. Although reputation values evolve across repeated interactions, long-term security depends on how these signals are governed and enforced at the decision layer. This paper introduces REVS-T, a four-tier governance and tier-aware selection mechanism using reputation bands, warning-streak escalation, and pool partitioning. Under persistent attack at 50 % adversarial ratio, REVS-T achieves 91.9 % task success and 96.7 % malicious avoidance with zero false exclusions, outperforming Threshold by 5.5 % and Beta by 14.9 %. A four-step ablation under shared reputation-update logic shows composite scoring and four-tier governance as the dominant drivers, with ST-conditioned initialization providing phase-shift adaptation and a fairness guarantee no evaluated baseline achieves.**

***Index Terms*—vehicular edge computing, provider selection, trust governance, security, computation offloading.**

I. Introduction

Vehicular edge computing (VEC) enables latency-sensitive intelligent transportation applications by offloading computation-intensive tasks to nearby edge providers [1]. Vehicle-to-vehicle (V2V) offloading leverages idle computational resources without fixed infrastructure, but introduces security risks: providers may withhold results, introduce delays, or behave intermittently to evade detection [2]. Strategic on–off attacks are particularly challenging because malicious providers alternate between honest and adversarial behavior to remain eligible over time.

Trust and reputation mechanisms provide soft-security control by penalizing misbehavior and representing uncertainty [7]. We adopt the common view that trust is a current, situation-specific expectation about behavior, whereas reputation is a longer-term measure derived from historical interactions [11], [12]. Prior work has focused primarily on trust computation and aggregation, including Beta reputation systems (BRS) [5], Subjective Logic [6], Dirichlet-based extensions [9], reward–penalty schemes [4], dual-threshold models [3], and blockchain-assisted aggregation [10], with initialization strategies ranging from uniform priors [3], [4], [10] to social trust level (ST) schemes [8]. While these methods improve trust estimation, most enforce decisions through single- or dual-threshold exclusion, neither partitioning the eligible provider pool during selection nor maintaining behavioral memory across interactions, allowing strategic providers to recover eligibility between attack windows. BRS-style evidence updates additionally produce false exclusions of honest providers under adversarial conditions (5.7 % in our evaluation).

These limitations persist even in trust-aware selection frameworks that improve scoring without restructuring the underlying threshold-based exclusion logic. Our prior REVS framework [13] improved trust-aware selection through ST-levels initialization and composite scoring, but did not incorporate structured governance across continuously evolving reputation signals or partition the eligible provider pool by tier. This paper addresses that gap by introducing Reputation-Enhanced Vehicle Selection with Trust-Tier-aware governance (REVS-T), a four-tier governance and tier-aware selection mechanism combining threshold-defined reputation bands (Boost, Stable, Warning, Blacklist) with outcome-conditioned warning-streak escalation. The design captures both long-term trust signals and short-term behavioral deviations. All non-Beta ablation models share identical reputation-update logic, enabling clean isolation of governance depth, tier-aware selection, and initialization effects. The main contributions are:

- **Four-tier governance with tier-aware pool partitioning.** REVS-T combines reputation-band governance, warning-streak escalation, and tier-aware composite selection. Under 50 % persistent attack, it achieves 91.9 % SR and 96.7 % MAR with zero false exclusions, outperforming Threshold by 5.5 % and BRS by 14.9 % relative.
- **Resilience against strategic attacks.** By preserving behavioral memory across interactions, REVS-T sustains 90.9 % MAR under on–off attack (+24.7 % over Threshold-MaxRep) and improves MAR by 49.2 % over BRS under phase-shift attack.
- **Controlled four-step ablation** (shared reputation update): composite scoring and governance dominate; ST-conditioned initialization is most threat-robust (phase-shift adaptation, zero false exclusions); tier-aware partitioning is marginal.

The remainder covers the system model and the governance

framework Section II, the evaluation and results Section III, and the conclusion Section IV.

II. System Model and Governance-Aware Provider Selection

A. Network Model

We adopt the vehicular computation offloading (VCOff) model and vehicle-to-vehicle (V2V) interaction pipeline of our previously proposed REVS framework [13] as the operational baseline. The architecture comprises three logical layers:

1) *Cloud layer:* The Trusted Authority (TA) registers vehicles, issues cryptographic credentials and certified Social Trust (ST) levels $\mathrm{ST} \in \{\text{high}, \text{interm}, \text{low}\}$, and maps real identities to pseudonyms.
2) *Stationary edge layer:* Roadside units (RSUs) equipped with vehicular edge computing (VEC) servers host the consortium blockchain, execute smart contracts, and perform reputation updates. RSUs provide governance enforcement; computation tasks are executed by vehicles.
3) *Vehicular layer:* Smart vehicles within RSU coverage act as a requester V_{req}^{i} issuing an offloading request, or as one of $V_{\text{pro}}^{j} \in V_{\text{pro}}$, the set of candidate providers. Roles are determined by current resource availability.

Upon entering RSU coverage, vehicles broadcast authenticated beacons; providers are also expected to lock a digital-coin deposit, forfeited on failure and redeemed on success. These mechanisms are part of the security architecture but fall outside the simulated decision layer, which focuses on reputation governance and selection. Figure 1 illustrates the end-to-end REVS-T workflow across the TA, RSU, and vehicular layers.

B. Reputation Signal Model

Each provider j carries a global scalar reputation $R_j \in [0,1]$, stored immutably on-chain and updated after every interaction based on task outcome (success or deadline failure). Reputation evolves during a simulation run from accumulated interaction evidence for each requester–provider pair (i, j). These evidence parameters $(\alpha_{ij}, \beta_{ij}, \gamma_{ij})$ representing counts of positive, negative, and uncertain interactions, from which the requester forms a subjective opinion. This opinion is then fused with a recommended opinion aggregated from other requesters via the Jøsang consensus operator [6]. The RSU scalarizes the fused opinion to a single reputation value [3]:

$$R_j = b_{ij}^{\text{final}} + a_0 \cdot u_{ij}^{\text{final}}, \tag{1}$$

where b_{ij}^{final} and u_{ij}^{final} are the belief and uncertainty masses of the fused opinion and a_0=0.5 is the base rate [6]. Unlike systems where each requester independently maintains a private R_{ij}, the ij subscripts on b^{final}, u^{final} reflect that fusion consumed pairwise opinions from all requesters i and R_j is the resulting network-level consensus independent of any individual requester. This paper isolates the contribution of governance and selection policies acting on this trust signal; a detailed comparison of reputation-update mechanisms is left to future work. All non-Beta variants share the same evidence-based update; BRS [5] replaces it with a classical two-parameter scheme and serves as our external baseline.

C. ST-Conditioned Initial Reputation

Inspired by the ST-based evidence initialization in [8] and adopted in REVS [13], each provider is assigned a prior subjective opinion $\omega_{ij}^{0} = (b_{ij}^{0}, d_{ij}^{0}, u_{ij}^{0})$ from evidence parameters (α, β, γ) conditioned on its TA-issued ST level. $\mathrm{ST}^{\text{high}}$, $\mathrm{ST}^{\text{interm}}$, and ST^{low} map to $(3,2,1)$, $(2,2,2)$, and $(1,2,3)$ [8], where $b_{ij}^{0} = \alpha/(\alpha+\beta+\gamma)$, $u_{ij}^{0} = \gamma/(\alpha+\beta+\gamma)$. At t=0, the evidence parameters depend only on j's ST credential and are identical across all i; applying the scalarization from (1) yields $R_j^{0} \in \{0.583, 0.500, 0.417\}$. ST^{low} providers therefore begin with smaller initial reputation, reducing early selection probability. Under sustained operation, accumulated evidence dominates the prior at a rate proportional to selection frequency.

D. Tier-Aware Selection via REVS Smart Contract

When V_{req}^{i} submits an offloading request, the smart contract executes Algorithm 1, which combines two distinct mechanisms: *Feasibility filtering* forms the index set $\mathscr{B} \subseteq \{1, \ldots, |\mathscr{V}_{\text{pro}}|\}$ of providers traveling in the same direction as V_{req}^{i}, with $R_j \geq R_{\min}$, and not currently restricted by governance (see §II-E). *Tier-aware partitioning* then splits $\mathscr{B}$ into Good-tier $\mathscr{G} = \{j \in \mathscr{B} : R_j > R_{\text{warn}}\}$ and Warning-tier $\mathscr{W} = \mathscr{B} \setminus \mathscr{G}$, with active set $\mathscr{A} = \mathscr{G}$ if non-empty, else $\mathscr{W}$, else $\emptyset$ (no provider available). Selection maximizes the composite score $w_1 R_j + w_2 \hat{T}_{ij}^{\text{stay}}$ [13], where $\hat{T}_{ij}^{\text{stay}} \in [0,1]$ normalizes the per-pair stay time by $\max_{u \in \mathscr{A}} T_{iu}^{\text{stay}}$. Partitioning operates at the selection layer, distinct from the post-update governance that determines band membership.

After selection, the RSU returns a signed response; the encrypted task data is transmitted and result is returned. Execution success is determined by whether latency meets the task deadline. The outcome and updated tier state are recorded on-chain to seed the next interaction.

E. Four-Tier Governance with Warning-Streak Escalation

REVS-T governance (Figure 1) applies two interacting components after each interaction, based on R_j and outcome $\textit{success} \in \{0,1\}$: (i) reputation bands that determine provider eligibility and monitoring status, and (ii) a warning-streak counter that accumulates behavioral evidence and provides an escalation path independent of the current band (Table I). The streak update is $\textit{streak} \leftarrow \textit{streak} + 1$ on failure and $\textit{streak} \leftarrow \max(0, \textit{streak} - 0.5)$ on success, where gradual decay retains prior evidence across honest intervals. The escalation threshold adapts to the current reputation tier:

$$L_{\text{esc}}(R_j) = \begin{cases} \bar{L}+2, & R_j \geq R_{\text{boost}}, \\ \bar{L}+1, & R_{\text{warn}} \leq R_j < R_{\text{boost}}, \\ \bar{L}, & R_j < R_{\text{warn}}, \end{cases} \tag{2}$$

with $\bar{L} = 2$ is the base tolerance, yielding tolerances of 4, 3, and 2 failures for the Boost, Stable, and Warning tiers respectively.

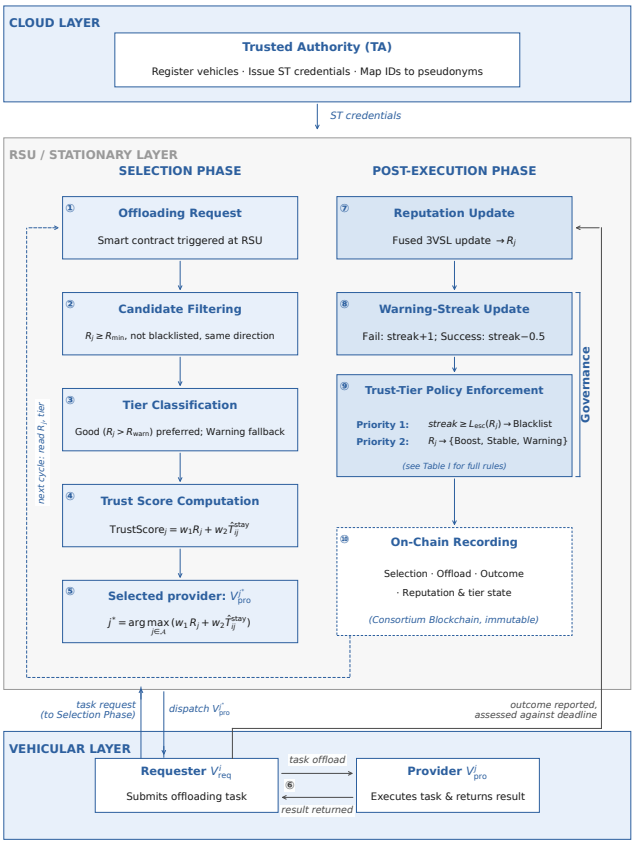

Fig. 1. REVS-T end-to-end offloading cycle. ST initializes reputation; the RSU enforces tier-aware selection and records (R_j, tier) on-chain.

Algorithm 1 REVS-T Tier-Aware Provider Selection

function SELECTPROVIDER($\mathcal{V}_{\text{pro}}, V_{\text{req}}^{i}$)
  (1) Feasibility filtering
  $\mathcal{B} \leftarrow \emptyset$
  for all $j \in \{1, \ldots, |\mathcal{V}_{\text{pro}}|\}$ **do**
    if sameDirection$(V_{\text{pro}}^{j}, V_{\text{req}}^{i})$ **and** $R_j \geq R_{\min}$ **and** $isActive_j$ **and** $blackoutCounter_j = 0$ **then**
      $\mathcal{B} \leftarrow \mathcal{B} \cup \{j\}$
    end if
  end for
  (2) Tier partition
  $\mathcal{G} \leftarrow \{j \in \mathcal{B} : R_j > R_{\text{warn}}\}$
  $\mathcal{W} \leftarrow \mathcal{B} \setminus \mathcal{G}$
  if $\mathcal{G} \neq \emptyset$ **then**
    $\mathcal{A} \leftarrow \mathcal{G}$
  else if $\mathcal{W} \neq \emptyset$ **then**
    $\mathcal{A} \leftarrow \mathcal{W}$
  else
    return $\perp$ ▷ no provider available
  end if
  (3) Stay-time normalization
  $\hat{T}_{ij}^{\text{stay}} \leftarrow \dfrac{T_{ij}^{\text{stay}}}{\max_{u \in \mathcal{A}} T_{iu}^{\text{stay}}}, \quad \forall j \in \mathcal{A}$
  (4) Composite scoring
  $j^* \leftarrow \arg\max_{j \in \mathcal{A}} \left(w_1 R_j + w_2 \hat{T}_{ij}^{\text{stay}}\right)$
  return $V_{\text{pro}}^{j^*}$
end function

TABLE I
REVS-T FOUR-TIER GOVERNANCE: REPUTATION BANDS AND STREAK ESCALATION

(a) Tier Definitions

Tier	Condition	Elig.	Effect
Boost	$R_j \geq R_{\text{boost}}$	Primary	If success: $R_j \leftarrow \min(R_j + \Delta R, 1)$
Stable	$R_{\text{warn}} < R_j < R_{\text{boost}}$	Primary	Standard monitoring
Warning	$R_{\min} \leq R_j \leq R_{\text{warn}}$	Fallback	Monitored (no freeze)
Blacklist	Policy-triggered (see (b))	Excluded	Freeze K_n (progressive)

(b) Escalation and Recovery Rules

Trigger	Action
Priority 1 — Streak-based	
$streak \geq L_{\text{esc}}(R)$	Blacklist; deactivate; reset streak; freeze K_n
Priority 2 — Reputation-based	
$R_j < R_{\min} \wedge fail$	Blacklist; deactivate; reset streak; freeze K_n
$R_j < R_{\min} \wedge success$	Warning (active; no freeze)
$R_j \geq R_{\text{boost}} \wedge success$	Boost: $R_j \leftarrow \min(R_j + \Delta R, 1)$
$R_j \geq R_{\text{boost}} \wedge fail$	No boost; remain active (non-warning)
Otherwise	Assign tier by R_j per (a)
Recovery and Progressive Penalty	
Freeze expires (K_n events)	Re-eligible: enters Warning with $streak$=1; tier re-evaluated on next update
Each blacklist event	$K_n = K_{\text{base}} \cdot 2^{n-1}$ for $n < maxStrikes$
$totalBlacklists \geq maxStrikes$	Permanent exclusion

Note: Boost is applied only on success; blacklisted providers are excluded from selection until the freeze expires; permanent exclusion triggers at the third blacklist event ($maxStrikes$=3).

Higher-reputation providers therefore absorb more failures before escalation. The Boost offset distinguishes demonstrably reliable providers from steady-state Stable ones; removing it collapses escalation to a binary Active/Warning split.

Each blacklist event doubles the freeze window, $K_n = K_{\text{base}} \cdot 2^{n-1}$ for $n < maxStrikes$, yielding 3 and 6 selection events for the first two strikes; the third triggers permanent exclusion. On freeze expiry, providers re-enter the Warning tier with $streak = 1$ (probation), so a single subsequent failure suffices to re-blacklist a Warning-tier provider. Unlike legacy REVS, where reputation is re-initialized per run, REVS-T reputation persists and accumulates across all interactions within a provider's operational lifetime at the RSU, so governance operates over the full provider life-cycle.

III. EVALUATION

A. Compared Models

All non-Beta models share the identical subjective-logic reputation update, so pairwise comparisons isolate governance, selection, and initialization; BRS is the external baseline. Table II summarizes the models.

TABLE II
COMPARED MODELS. GOV.=GOVERNANCE; SEL.=SELECTION; INIT.=INITIALIZATION; COMP.=$w_1 R_j + w_2 \hat{T}_{ij}$.

Model	Gov.	Sel.	Init.	Purpose
REVS-T (prop.)	4-tier	TierAware+Comp	ST-lev.	Proposed
REVS-T-Uniform	4-tier	TierAware+Comp	Uniform	Init. ablation
Threshold	Single	TierAware+Comp	ST-lev.	Gov. ablation
Threshold-NoTierSel	Single	NoTier+Comp	ST-lev.	Sel. ablation
Threshold-MaxRep	Single	Max-Rep	ST-lev.	Scoring ablation
BRS [5]	Single	Max-Rep	Uniform	Baseline

- **REVS-T** (proposed): four-tier governance with warning-streak escalation, tier-aware trust–mobility selection, and ST-conditioned initialization (Section II-C).
- **REVS-T-Uniform**: REVS-T with a uniform prior $(\alpha,\beta,\gamma)=(2,2,2)$, $R_j^0=0.500$ in (1); isolates ST-conditioned initialization.
- **Threshold**: REVS-T governance replaced by a single $R_{\min}$ blacklist rule; isolates governance depth.
- **Threshold-NoTierSel**: Threshold without the Good/Warning partition (all $R \geq R_{\min}$ compete under composite scoring); isolates tier-aware selection.
- **Threshold-MaxRep**: Threshold with max-reputation selection and no $R_{\min}$ admission filter; jointly isolates composite trust–mobility scoring and admission filtering.
- **BRS** [5]: classical Beta evidence, max-reputation selection, single $R_{\min}$ threshold, uniform prior $(\alpha=\beta=1)$.

All single-threshold models (BRS and the three Threshold variants) blacklist only when $R_j < R_{\min}$ *and* the interaction fails (fixed freeze *penaltyFreezeRuns*=3); a low-R success is never blacklisted. REVS-T additionally blacklists via streak escalation (*streak* $\geq L_{\text{esc}}$) and applies a progressive freeze ($K_n=K_{\text{base}}\,2^{n-1}$), as defined in Section II-E.

B. Adversary Model

Adversaries are insider vehicles with valid credentials that strategically degrade service, modeled as reduced computation capacity and transmission rate when attacking, inducing deadline failures, via three strategies: *always-attack* (persistent failure), *on–off attack* (alternating behavior to exploit reputation recovery), and *Behavioral Phase Shift attack* (honest for runs 1–500, then switch to persistent failure). Honest RSU/requester are assumed, focusing the analysis on provider behavior [13].

Two ST assignment strategies apply to malicious providers. **Xu-Aligned**: all malicious providers are ST^{low} and all honest providers are ST^{high}, following the distribution in [8]. **Moderate**: 5% ST^{high}, 15% $\text{ST}^{\text{interm}}$, 80% ST^{low} among malicious providers, reflecting a realistic adversarial correlation between social trust and cooperative behavior: most attackers possess low ST, while a small tail retains partially legitimate credentials at the intermediate or high level.

C. Evaluation Metrics

SR (%): task success rate. **MAR** (%): malicious-avoidance rate over the full run; for phase-shift, Phase-2 MAR (runs 501–1,000) is reported separately. $\textbf{FPR}_{\text{gov}}$ (%): honest providers wrongly blacklisted. **PMTI**: mean selections of a malicious provider before its first blacklist (lower = faster detection). **Staleness Score** (0–100, lower = faster adaptation; Phase Shift only): composite of the Phase-2/Phase-1 malicious-selection ratio, runs-to-detect (rolling MAR $\geq 70\%$), and Phase-2 MAR recovery. **Parameters**: $R_{\min}$=0.41, R_{warn}=0.50, R_{boost}=0.70, ΔR=0.05, K_{base}=3, $\bar{L}$=2, w_1=0.70, w_2=0.30. Thresholds anchor to the ST priors $R_j^0 \in \{0.583, 0.500, 0.417\}$: $R_{\min}$ sits below the ST^{low}, R_{warn} matches $\text{ST}^{\text{interm}}$, and R_{boost} exceeds the highest prior. $\bar{L}$=2 yields tolerances of $2/3/4$ failures across Warning/Stable/Boost. Network and task parameters are in Table III. Unless noted, results use 50 % malicious, moderate assignment, and 1,000 sequential interactions. The 50 % regime is where governance differences are most discriminative.

TABLE III
SIMULATION PARAMETERS

Platform		Network	
Platform	MATLAB R2025b	Req./Prov.	30/70
Runs/config.	1,000	Cand. pool	6–20
Mal. ratio	10,30,50%	RSU & V2V	250 m
		Veh. speed	60 km/h ±5,10
		ST (H/M/L)	30/30/40%
Tasks		**Provider**	
Input size	50–500 KB	Capacity	10 GHz
CPU cycles	$0.2–3.2\times10^9$	Tx rate	86 Mbps

D. Persistent Attack: Governance and Scoring Drivers

Table V summarizes persistent-attack results. Three drivers emerge, in decreasing order of within-selection impact:

Composite scoring with $R_{\min}$ admission filtering is the dominant within-selection gain. Threshold-NoTierSel achieves SR = 86.8 % vs. Threshold-MaxRep's 79.6 % (+9.0 %). Mobility weighting avoids deadline-prone providers and the $R_{\min}$ filter excludes degraded ones, neither of which max-reputation selection applies.

Four-tier governance accelerates containment. Compared with Threshold, REVS-T improves SR from 87.1 % to 91.9 % (+5.5 %) and MAR from 92.7 % to 96.7 %, cutting PMTI from 3.07 to 1.40 (−54 %, Table IV): streak escalation excludes a provider once *streak* reaches $L_{\text{esc}}(R_j)$=2 in the Warning band, whereas single-threshold governance must wait for R_j to fall below $R_{\min}$. Beta's PMTI is even lower (1.00) but at FPR_{gov}=5.7 %: its $(\alpha=\beta=1)$ prior places all providers at R_j=0.5, and one failure drops R_j to $1/3 < R_{\min}$, blacklisting honest providers as readily as malicious. REVS-T's selection overhead (0.4) over Beta buys zero FPR_{gov}. The gain is load-dependent: the REVS-T–vs.–Threshold SR gap stays ≤ 0.11 % at 10–30 % malicious, widening to 4.8 points at 50 % as sustained pressure drives providers into the Warning band and streak escalation compounds (Fig. 2).

Tier-aware selection contributes only +0.3 % SR. Composite scoring (w_1=0.70) already makes reputation the dominant selection factor, and malicious reputations collapse below R_{warn} within two failures (e.g. an $\text{ST}^{\text{interm}}$ provider from

TABLE IV
CONTAINMENT DIAGNOSTICS. PMTI AND STALENESS SCORE: LOWER IS BETTER. STALENESS FOR PHASE SHIFT ATTACK ONLY. BEST NON-BETA IN **BOLD**.

Model	PMTI Pers.	PMTI On–Off	Stale. Mod.	Stale. xu
REVS-T	**1.40**	2.30	**23.5**	**3.5**
REVS-T-Uniform	2.00	4.75	46.5	49.3
Threshold	3.07	2.56	48.9	**3.5**
Threshold-NoTierSel	3.29	1.71	51.2	**3.5**
Threshold-MaxRep	2.43	**1.59**	27.8	27.3
BRS	1.00	1.00	31.7	31.1

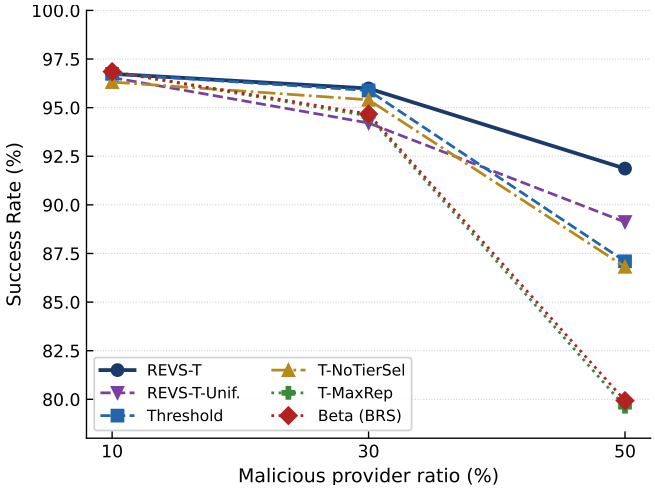

Fig. 2. SR vs. malicious ratio (moderate, persistent). Governance gap opens above 30 %.

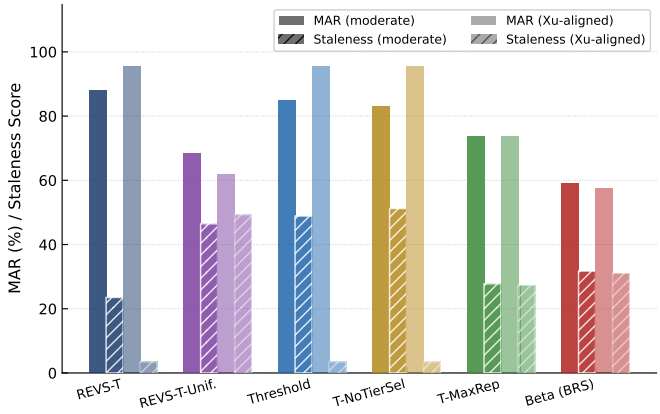

Fig. 3. MAR (full-run) and Staleness Score under Phase Shift attack at 50 % malicious, moderate and xu-aligned assignments. Lower staleness is better.

R_j^0=0.500 to 0.417), leaving little Warning-band residence for the mobility bonus to exploit. Threshold and Threshold-NoTierSel thus allocate near-identical pools (Good-tier 90.8 % vs. 90.6 %; Warning-tier 9.2 % vs. 9.4 %).

E. Strategic Behavior: Intermittent and Phase-Shift Attacks

1) On–Off Attack (Intermittent): On–off alternation degrades all models (malicious providers recover reputation during honest intervals; Table V): the smallest MAR points

TABLE V
RESULTS AT 50 % MALICIOUS (MODERATE, 1,000 EVENTS). SR AND MAR IN %; HIGHER IS BETTER. FPR_{gov}: PERSISTENT ATTACK.

Model	Persistent		On–Off		PS MAR	FPR gov
	SR	MAR	SR	MAR		
REVS-T	**91.9**	**96.7**	**90.3**	**90.9**	**87.9**	**0.0**
REVS-T-Unif.	89.1	94.2	85.0	79.8	68.2	2.9
Threshold	87.1	92.7	89.0	88.6	84.9	**0.0**
T-NoTierSel	86.8	92.5	88.8	88.1	83.1	**0.0**
T-MaxRep	79.6	84.6	82.8	72.9	73.6	**0.0***
BRS	79.9	84.9	84.3	75.6	58.9	5.7

*T-MaxRep: 2.9 % under on–off and Phase Shift (PS); REVS-T-Unif.: 2.9 % persistent only.

drops (4–6) are the ST-initialized composite models, the largest REVS-T-Uniform (−14.4) and Threshold-MaxRep (−11.7). The main findings are:

Four-tier governance provides the on–off-specific gain absent from single-threshold designs. REVS-T attains 90.9 % MAR versus 88.6 % for Threshold, a +2.6 % relative gain. On success the warning streak decays by 0.5 rather than resetting, so failure evidence persists across the honest half of each cycle and escalates on the next attack whereas single-threshold governance lacks this cross-cycle memory and repeatedly re-admits recovered providers.

Low PMTI does not imply containment. Threshold-MaxRep (1.59) and Beta (1.00) blacklist fastest yet score the worst MAR (72.9 %, 75.6 %), as the absent R_{min} admission filter lets recovered providers be re-selected during honest intervals (Table IV).

Tier-aware selection contributes non-monotonically. The partition (Threshold vs. Threshold-NoTierSel) gives +0.5 points MAR at 50 % malicious, peaking at +1.4 points at 30 % (96.6 % vs. 95.2 %) and ≈ 0 at 10 %, where the dense Good pool rarely needs the Warning fallback. As under persistent attack, composite scoring already concentrates selection in the Good pool.

2) Behavioral Phase Shift Attack: SR ranges from 83.1 % (Beta) to 94.7 % (REVS-T) under moderate assignment; MAR is reported for both the attack phase (runs 501–1,000) and the full run. REVS-T achieves 94.8 % Phase-2 MAR vs. 70.6 % for Beta (+34.3 %) and 87.9 % vs. 58.9 % full-run (Table V). Adaptation speed is captured by the Staleness Score. Under moderate assignment REVS-T scores 23.5 vs. 48.9 for Threshold (2.1 *times* faster), Table IV. The gain coming from streak escalation catching well established $\text{ST}^{\text{high/interm}}$ malicious before their reputation collapses. Under xu-aligned, where all malicious are ST^{low} and never established, governance becomes irrelevant (REVS-T = Threshold = Threshold-NoTierSel = 3.5) and initialization dominates: REVS-T-Uniform's score 49.3 is 14 *times* higher, confirming the ST prior lets governance act almost immediately at the switch (Fig. 3).

F. Initialization Effects and Governance Fairness

ST-conditioned initialization is the dominant driver across three aspects: early selection accuracy, credential exploitation, and governance fairness.

Early selection. At 100 events (moderate, persistent, 50 % malicious), REVS-T attains 80.0 % MAR vs. 62.0 % for REVS-T-Uniform (+29.0 % on EarlyMAR$_{100}$); the gap persists to 1,000 events (91.9 % vs. 89.1 % SR).

Credential leverage. Under xu-aligned REVS-T improves from 91.9 % to 94.7 % SR (+3.1 %), while REVS-T-Uniform does not benefit (89.1 % vs. 88.2 %), the ST prior is required to exploit credential informativeness. REVS-T and Threshold consequently converge to 94.7 % SR / 98.2 % MAR, since malicious providers begin below R_{warn} and cannot enter the preferred pool.

Governance fairness. REVS-T maintains zero FPR$_{\text{gov}}$ across all attack patterns. REVS-T-Uniform incurs 2.9 % under persistent attack because the uniform prior places honest providers at the Warning boundary ($R_j^0 = 0.500$), where a short run of failures escalates to blacklist; ST-conditioned initialization starts most honest providers in higher bands ($R_j^0 = 0.583$ for ST$^{\text{high}}$), absorbing such failures. Threshold-MaxRep instead incurs 2.9 % under on–off and phase-shift: without an R_{min} admission filter, borderline providers are repeatedly chosen until they cross the governance threshold. Beta records its worst-case 5.7 % under persistent attack. FPR$_{\text{gov}}$ is governed by initialization and admission filtering, not by governance depth or tier-aware selection: REVS-T, Threshold, and Threshold-NoTierSel maintain zero FPR$_{\text{gov}}$ across all attacks, while the uniform prior (REVS-T-Uniform, BRS) and the absent R_{min} filter (Threshold-MaxRep) each introduce false exclusions.

G. Ablation Summary

Fig. 4 ranks the four drivers (moderate, persistent, 50 % malicious): composite scoring with R_{min} admission filtering and four-tier governance are the two dominant within-model gains, ST-conditioned initialization is the most threat-robust (+3.1 to +12.5 SR across regimes); paired with the R_{min} admission filter, it gives REVS-T zero FPR$_{\text{gov}}$ vs. 2.9 % for REVS-T-Uniform under persistent attack; tier-aware selection is marginal under persistent attack but provides targeted on–off defence (+1.4 points MAR peak at 30 % malicious). BRS, lacking all four, reaches 79.9 % SR at FPR$_{\text{gov}} = 5.7$ %.

IV. Conclusion

This paper introduced REVS-T, a four-tier governance mechanism for vehicular computation offloading with reputation bands, warning-streak escalation, and tier-aware partitioning. A controlled ablation identifies composite scoring with R_{min} filtering (+9.0 %) and four-tier governance (+5.5 % SR) as the dominant drivers; streak escalation sustains on–off resilience, ST-conditioned initialization is the most threat-robust (phase-shift adaptation; necessary for zero false exclusions), and tier-aware partitioning is marginal. BRS's distinct update consistently underperforms across attack conditions,

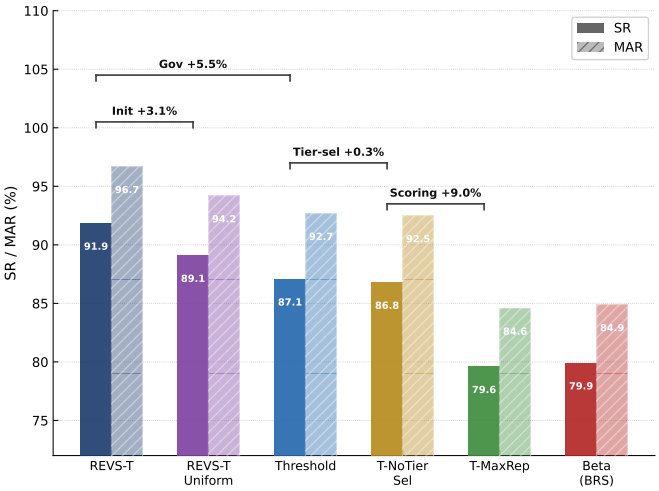

Fig. 4. Four-step ablation (50 % malicious, moderate, persistent): scoring +9.0 %, tier-sel +0.3 %, governance +5.5 %, init +3.1 % SR.

suggesting the reputation-update mechanism has an independent effect. Future work will examine uncertainty-aware updates under multi-tier governance.

References

[1] R. A. Dziyauddin *et al.*, "Computation offloading and content caching in vehicular edge networks: A survey," *Comput. Netw.*, vol. 197, p. 108228, 2021.

[2] S. Y. Fayi and Z. Sheng, "A survey of security, privacy, and trust issues in vehicular computation offloading and blockchain-based solutions," *Open Res. Eur.*, vol. 3, 2023.

[3] Z. Yang, K. Yang, L. Lei, K. Zheng, and V. C. M. Leung, "Blockchain-based decentralized trust management in vehicular networks," *IEEE Internet Things J.*, vol. 6, no. 2, pp. 1495–1505, Apr. 2019.

[4] S. Iqbal, R. M. Noor, and A. W. Malik, "Blockchain-based reputation management for task offloading in micro-level vehicular fog networks," *IEEE Access*, vol. 8, pp. 52968–52981, 2020.

[5] A. Jøsang and R. Ismail, "The beta reputation system," in *Proc. 15th Bled Electron. Commer. Conf.*, Bled, Slovenia, Jun. 2002, pp. 1–14.

[6] A. Jøsang, *Subjective Logic: A Formalism for Reasoning Under Uncertainty*. Cham, Switzerland: Springer, 2016.

[7] A. Jøsang, "The right type of trust for distributed systems," in *Proc. New Security Paradigms Workshop*, 1996, pp. 119–131.

[8] S. Xu, C. Guo, R. Q. Hu, and Y. Qian, "Blockchain-inspired secure computation offloading in vehicular cloud networks," *IEEE Internet Things J.*, vol. 8, no. 6, pp. 4315–4328, Mar. 2021.

[9] G. Liu *et al.*, "Assessment of multi-hop interpersonal trust in social networks by three-valued subjective logic," in *Proc. IEEE INFOCOM*, Apr. 2014, pp. 1698–1706.

[10] X. Liu *et al.*, "Blockchain-assisted reputation management scheme for Internet of Vehicles," *Sensors*, vol. 23, no. 10, p. 4624, 2023.

[11] A. Hbaieb, S. Ayed, and L. Chaari, "A survey of trust management in the Internet of Vehicles," *Comput. Netw.*, vol. 203, p. 108558, 2022.

[12] R. Hussain, J. Lee, and S. Zeadally, "Trust in VANET: A survey of current solutions and future research opportunities," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 5, pp. 2553–2571, 2020.

[13] S. Y. Fayi, F. Ayaz, and Z. Sheng, "A blockchain-based reputation-enhanced vehicle selection (REVS) for computation offloading," in *Proc. IEEE Veh. Technol. Conf. (VTC-Fall)*, 2025.