An Efficient Partial State Acquisition Scheme Under Communication Constraints in V2X

Mingyu Ma
School of Fundamental Science and Engineering
Waseda University
Tokyo, Japan
william.ma@asagi.waseda.jp

Hidenori Nakazato
School of Fundamental Science and Engineering
Waseda University
Tokyo, Japan
nakazato@waseda.jp

***Abstract*—Efficient partial state acquisition is a critical bottleneck for distributed state estimation in V2X edge networks, where nodes must make decisions under strict latency constraints and incomplete global knowledge. This paper proposes to introduce Information Centric Networking (ICN) to V2X edge networks and to apply Method-M, an ICN-aware roadside unit (RSU) selection framework that jointly optimizes information diversity and retrieval efficiency in data collection. We model ICN's hop-reduction effect through an analytical hop-reduction ratio, and evaluate its robustness across different topology types, request intensities, aggregation windows, and latency budgets. RSU selection is formulated as a latency-constrained knapsack problem and solved via Lagrangian relaxation with dynamic Jensen-Shannon Divergence based utility refresh. Experiments on SUMO-generated topologies and over 5,800 parameter configurations show that Method-M achieves 15–16% higher Total Marginal Information Gain (TMIG) than the other algorithms, reduces retrieval latency by up to 82% compared to the ICN-ablated variant, and trails the algorithm using global knowledge by only 6–7%.**

***Index Terms*—V2X; Information-Centric Networking; Distributed Particle Filter; Jensen-Shannon Divergence; Lagrangian Relaxation; Edge Computing**

I. Introduction

In traditional Vehicle-to-Everything (V2X) networks, roadside units (RSUs) serve as key elements for both computation and communication. Recent studies have explored offloading autonomous driving tasks from vehicles to the edge network [1], driven by the data- and computation-intensive nature of autonomous driving and the prohibitive cost of equipping every vehicle with sufficient on-board resources.

However, edge networks are highly heterogeneous, delay-sensitive, and resource-constrained, with system states that are often only partially observable [2]. This makes data collection in V2X networks a decision-making problem under incomplete information. Many existing reinforcement learning based methods overlook this issue by assuming full state knowledge [3]. In practice, however, vehicle mobility, changing channel conditions, and fluctuating server loads make accurate global state acquisition difficult within a short decision window [4]. More recent studies therefore model mobile edge computing (MEC) networks under partial observability, reflecting the decentralized nature of practical edge systems [5].

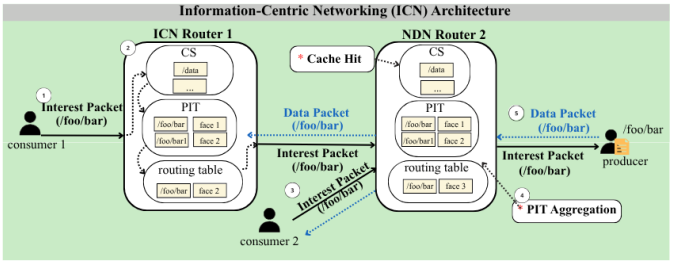


Fig. 1. ICN Architecture.

Particle filters provide a principled framework for estimating hidden states from partial observations [6], while distributed variants allow nodes to exchange local beliefs to approximate the global state [7]. In V2X edge networks, however, querying all RSUs may incur excessive communication delay, whereas querying only nearby or low-cost RSUs may miss informative state updates. We therefore focus on the preceding state-acquisition problem: deciding which remote RSU states should be queried under a limited communication budget before downstream offloading or other decisions are made.

We further exploit ICN caching and Interest aggregation as communication-side mechanisms that reduce the retrieval cost considered during state acquisition. Our contributions are: (1) an ICN-aware latency model; (2) a hop-reduction formulation based on temporal reuse and network topology; and (3) Method-M, named after its marginal information-driven selection principle, a Jensen-Shannon Divergence (JSD) based and Lagrangian guided RSU selection scheme that reduces redundancy under a latency budget.

II. Background

A. Information Centric Networking

Information-Centric Networking (ICN) replaces address-based routing with name-based content retrieval [8]. A consumer sends an *Interest packet* for named content, which is returned in a *Data packet*. ICN routers maintain a Pending Interest Table (PIT) to aggregate concurrent Interests for the same content and a Content Store (CS) to cache returned Data [9], [10]. Both mechanisms can reduce redundant transmissions and retrieval latency.

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

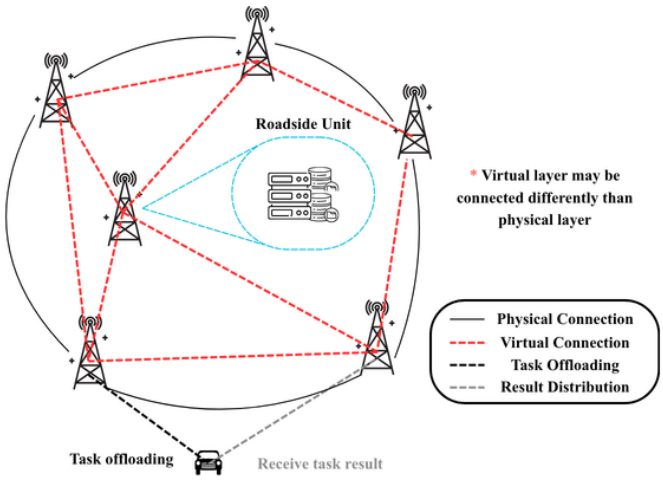


Fig. 2. V2X Network and Road Topology.

In the partial state acquisition problem considered in this paper, the content requested by an RSU is not an arbitrary file but a short-lived state summary produced by another RSU. Such state summaries may be requested repeatedly by multiple nearby decision makers within the same time window. This makes PIT aggregation particularly relevant: concurrent or near-concurrent Interests for the same state summary can be merged before the Data packet returns. In-network caching may also help when the cached state summary remains fresh within the decision window. Therefore, ICN provides a natural communication substrate for reducing redundant state retrieval in V2X edge networks.

B. Vehicle to Everything

V2X refers to a vehicular network in which vehicles and RSUs are interconnected, supporting applications such as road safety, traffic efficiency, and passenger infotainment [11]. With the growth of AI, offloading computation to edge RSUs has proven beneficial for autonomous driving [12]. As shown in Fig. 2, each RSU can maintain local knowledge of its covered road segment via particle filters. Since RSUs are resource-constrained and V2X networks generate large data volumes, efficient distributed data management is essential.

C. Distributed Particle Filter

Particle filtering is a sequential Monte Carlo method that represents the posterior state distribution using a set of weighted particles, updated iteratively from observations [13]. At each time step, particles are propagated through a system dynamics model and weighted by the likelihood of new observations. Periodic resampling favors high-weight particles and suppresses low-weight ones, thereby maintaining representational efficiency and avoiding particle degeneracy [14].

In distributed settings, each node maintains a local particle set and exchanges information with other nodes to refine the global state estimate without centralization [15]. This framework is well-suited for V2X environments, where observations are inherently distributed and complete global state is rarely available in real time. Each RSU acquires useful non-local state information under limited communication time and estimate the current global state.

D. Information Value in Partial State Acquisition

In distributed V2X edge networks, the usefulness of a remote state update depends not only on whether it is available, but also on how much new information it contributes to the current belief. State observations from nearby or highly correlated road segments may be redundant, while observations from distant segments may be less relevant to update the current belief at the RSU requesting the information. Therefore, partial state acquisition requires a mechanism for evaluating both the novelty and the relevance of candidate state information.

From an information-theoretic perspective, the value of acquiring an additional observation can be interpreted as the reduction in uncertainty after incorporating that observation. However, directly computing posterior information entropy reduction requires access to the joint distribution of the global system state, which is generally unavailable in distributed V2X settings. This motivates the use of distributional divergence as a practical proxy for marginal information value. In this work, we use Jensen–Shannon divergence to compare particle-based state distributions, as formalized in Section III-B.

E. Related Work and Positioning

Existing studies have investigated communication-efficient information acquisition from several related perspectives. In vehicular content-centric networks, cooperative caching has been used to improve content availability and retrieval efficiency by exploiting mobility and social relationships among vehicles [16]. Communication-constrained source selection has also been studied in distributed sensing systems. Godrich *et al.* formulated sensor selection for distributed localization as a knapsack problem under resource constraints [17], while Liu *et al.* investigated distributed sensor selection for field estimation [18]. These studies demonstrate the benefit of selectively acquiring information rather than collecting observations from all available sources. In contrast, our work considers RSU-level state acquisition in V2X edge networks and explicitly models the acquisition cost according to ICN-induced retrieval opportunities, including Interest aggregation and cache reuse.

Recent V2X studies increasingly employ learning-based approaches for decision making under incomplete or constrained information. Chai *et al.* studied learning-based resource allocation for ultra-reliable V2X communication under partial channel state information [19], while Dong *et al.* jointly considered information value, task offloading, and resource allocation for vehicular cooperative perception [20]. These approaches primarily optimize downstream communication, computation, or task-allocation decisions based on the available observations. Method-M instead focuses on the preceding state-acquisition stage and determines which remote RSU states should be retrieved under a communication budget before a downstream decision is made.

III. Problem Formulation

A. System Model

We model the V2X edge network as a graph $G = (V, E)$, where each node $i \in V$ is an RSU equipped with computation, storage, and sensing resources. As illustrated in Fig. 2, each RSU aggregates vehicle observations within its coverage area into particles characterized by four traffic features: mean speed, road occupancy, mean acceleration, and traffic flow rate. Each feature is discretized into three levels (low, medium, and high), and the joint frequencies of these discretized states are normalized into the probability distribution used for JSD-based utility computation.

When a vehicle offloads a task, the task should be executed at the RSU that has good information of surrounding area of the vehicle. We refer to this RSU as the requester RSU i. The requester RSU observes the vehicle's task request, collects additional state information from other RSUs, and then executes the task.

Since each RSU only observes its local traffic and workload state, it must retrieve particle-based state distributions from other RSUs to improve its belief of the global system state. However, retrieving states from all RSUs within a decision window is infeasible under strict latency constraints. Before executing the task, the requester RSU must first decide which subset of RSUs should be queried for state acquisition. In this work, we formulate this as an RSU subset selection problem that maximizes information gain under a hard retrieval budget.

Although offloading requests are generated by individual vehicles, the state representation in our formulation is maintained at the RSU level. We therefore treat RSUs as aggregation points, rather than treating individual vehicles as particles. The detailed particle generation process is outside the scope of this work; existing studies have shown the suitability of particle-filter-based state estimation in edge and vehicular settings [21]. In this work, each RSU's particle set is represented as a probability distribution over bucketized system states.

B. Information Utility

The value of querying an RSU depends on how much new information its distribution provides relative to the current belief. Since nearby or correlated RSUs may provide redundant state information, the objective is not simply to retrieve as many RSU states as possible, but to identify state summaries that provide useful incremental information under the latency budget.

Directly computing posterior entropy reduction would require access to the joint distribution of the global system state, which is unavailable in distributed V2X settings. We therefore use JSD as a bounded and symmetric proxy for marginal information novelty.

Let P and Q denote two particle distributions. Their JSD is defined as

$$\mathrm{JSD}(P\|Q) = \tfrac{1}{2} D_{\mathrm{KL}}(P\|M) + \tfrac{1}{2} D_{\mathrm{KL}}(Q\|M), \qquad (1)$$

where $M = \frac{1}{2}(P + Q)$ is the average of two distributions P and Q, and $D_{\mathrm{KL}}(P\|M)$ is the Kullback-Leibler divergence between distributions P and M.

Let $\mathcal{P}_j$ be the particle distribution of candidate RSU j, and let $\mathcal{P}_{i\cup\mathcal{S}}$ be the fused belief distribution at requester RSU i after incorporating already selected RSUs $\mathcal{S}$. The marginal utility of selecting j is

$$U_{j|i,\mathcal{S}} = \mathrm{JSD}(\mathcal{P}_j \parallel \mathcal{P}_{i\cup\mathcal{S}}) \cdot \exp(-\alpha h_{ij}) \cdot \min_{s\in\mathcal{S}} \mathrm{JSD}(\mathcal{P}_j \parallel \mathcal{P}_s)\,, \qquad (2)$$

where the first term measures divergence from the current belief, the second term applies distance decay, and the third term penalizes redundancy with already selected RSUs. When $\mathcal{S} = \emptyset$, the novelty term, which is the third term, is set to 1.

For a selected sequence $\mathcal{S} = \{j_1, \ldots, j_k\}$, we define Total Marginal Information Gain (TMIG) as

$$\mathrm{TMIG}_i(\mathcal{S}) = \sum_{m=1}^{k} U_{j_m|i,\,\mathcal{S}_{m-1}}, \qquad (3)$$

where $\mathcal{S}_{m-1} = \{j_1, \ldots, j_{m-1}\}$. TMIG is computed using raw shortest-path hops, i.e., the shortest-path distance before any ICN-induced hop reduction. This keeps information utility independent of ICN's retrieval benefit, so ICN affects how many RSUs can be reached within the budget but does not artificially increase the information value of any candidate.

The utility is evaluated independently in each decision window. Within a window, particle distributions are treated as fixed snapshots; across windows, they are updated as new observations arrive. The novelty term in Eq. (2) also provides a practical diminishing-novelty effect, since adding more RSUs to $\mathcal{S}$ can only reduce or maintain the minimum JSD to already selected distributions.

C. ICN-Aware Latency Model

Without ICN, the latency of retrieving data from RSU j is

$$T_j^{(0)} = 2h_{ij}t_{\mathrm{hop}} + \frac{s_j}{r}, \qquad (4)$$

where h_{ij} is the shortest-path hop count, t_{hop} is the average per-hop delay, s_j is the retrieved data size, and r is the effective transmission rate.

ICN reduces effective retrieval latency through PIT aggregation and in-network caching. We model the hop-reduction ratio as

$$\rho_{ij} = \big(1 - e^{-\lambda\Delta}\big)\beta_{ij}, \qquad (5)$$

where λ is the Interest arrival rate, Δ is the aggregation window, and β_{ij} is the path-average normalized betweenness centrality. The term $1 - e^{-\lambda\Delta}$ captures temporal request overlap under a Poisson arrival model, while β_{ij} approximates topology-dependent aggregation opportunity through path centrality. This is a high-level topology- and request-aware latency abstraction; lower-layer PHY/MAC effects are represented through the effective per-hop delay t_{hop} and transmission rate r.

The ICN-aware retrieval latency is therefore

$$T_{ij} = 2h_{ij}(1 - \rho_{ij})t_{\mathrm{hop}} + \frac{s_j}{r}. \qquad (6)$$

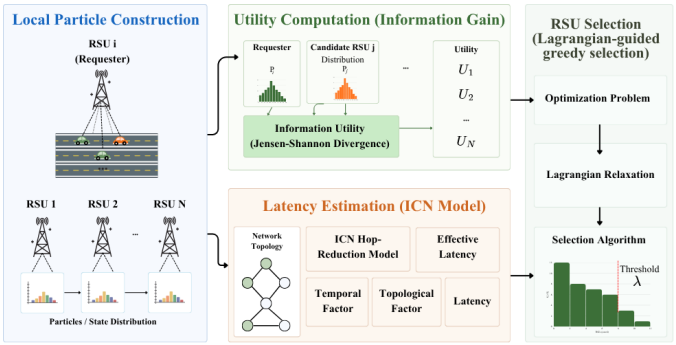


Fig. 3. System Overview.

In practice, λ can be estimated from recent Interest arrivals for the corresponding state prefix, while Δ is determined by the state acquisition or decision window. The topology-dependent factor β_{ij} can be computed from the RSU-level communication graph and updated when the network topology changes. Thus, the model does not require current-window global state collection; it only requires local request statistics and a known RSU topology.

D. Latency-Constrained Selection

Given a requester RSU i, the objective is to select candidate RSUs that maximize TMIG under a latency budget:

$$\begin{aligned} \max_{x_j} \quad & \sum_{j \in V} x_j U_{j|i,\mathcal{S}} \\ \text{s.t.} \quad & \sum_{j \in V} x_j T_{ij} \le T_{\text{budget}}, \\ & x_j \in \{0,1\}, \quad \forall j \in V. \end{aligned} \tag{7}$$

This is a latency-constrained 0–1 knapsack problem. To obtain a practical selection rule, we apply a Lagrangian relaxation and select candidates according to their information gain per unit latency:

$$\frac{U_{j|i,\mathcal{S}}}{T_{ij}} > \mu, \tag{8}$$

where μ controls the marginal cost of consuming latency budget.

IV. Method

A. Overview

Figure 3 illustrates the overall framework of Method-M, following the problem setting introduced in Section III. When a requester RSU needs to collect state information, Method-M addresses the state-acquisition step: it selects a subset of RSUs to query so that the requester RSU can improve its belief of the global system state within a latency budget. The framework proceeds in three stages.

Method-M is a budget-aware state-acquisition framework that can operate with either baseline or ICN-aware retrieval latency. ICN affects only the retrieval-cost term through caching and Interest aggregation, while leaving the information-utility definition unchanged.

Stage 1 – Local particle construction. Each RSU aggregates observations from its road segment into a particle-based distributional representation, mapping raw measurements into a unified bucket format suitable for utility computation.

Stage 2 – Utility and cost estimation. The incremental utility of each candidate RSU is computed via JSD with distance decay (Eq. (2)), using the globally synchronized network state from the *previous* time window as the current belief. This temporally shifted snapshot avoids real-time global communication while remaining compatible with standard V2X time-synchronization [22]. Retrieval cost T_j is estimated via Eq. (6): ICN lowers T_j without affecting $U_{j|i,\mathcal{S}}$, so its benefit is captured entirely through the ratio $U_{j|i,\mathcal{S}}/T_j$.

Stage 3 – Lagrangian-guided greedy selection. Given utility scores and latency costs, RSU selection is solved via Lagrangian relaxation and greedy selection, as described in Section IV-B.

B. RSU Selection Algorithm

Algorithm 1 summarizes the proposed Method-M. In each decision window, a requester RSU i selects a subset of candidate RSUs from the candidate RSU set $\mathcal{R}$ and retrieves their state information before making the downstream decision. The selection framework itself is not restricted to ICN: it can also operate with the baseline retrieval latency in a conventional request–response network. ICN is introduced as a communication-layer enhancement that reduces the effective retrieval cost through in-network caching and Interest aggregation. Therefore, Method-M remains applicable without ICN, while the ICN-aware version can achieve a better information–latency trade-off under the same retrieval budget. The function $\textsc{Utility}(i, j, \mathcal{S})$ implements the topology-aware marginal information gain $U_{j|i,S}$ defined in Eq. (2) with requester/source RSU i, candidate RSU j, and the set of already selected RSUs $\mathcal{S}$. When $\mathcal{S} = \emptyset$, the utility $u_0[j] = U_{j|i,\emptyset}$ represents the initial information value of candidate j before considering redundancy with previously selected RSUs. The function $\textsc{LatencyICN}(i, j)$ implements the ICN-aware retrieval latency T_{ij} defined in Eq. (6) and evaluates the latency for the retrieval path from requester RSU i to candidate RSU j yielding $t[j] = T_{ij}$. Thus, Eq. (2) is used to estimate the information benefit of selecting j, while Eq. (6) is used to estimate the corresponding retrieval cost.

The configuration parameter δ is the tolerance for budget matching in $\textsc{LagrangianPrune}$, K is the maximum number of bisection iterations, θ is the target TMIG threshold, and `stop_mode` determines whether the selection stops by budget exhaustion or by reaching the TMIG threshold.

Remark. The pruning step filters the candidate space but does not determine the final ranking. Final ordering is decided by the greedy loop, where utilities are recomputed after each selection to reflect updated novelty.

C. Design Rationale and Complexity

Method-M separates candidate pruning from final greedy ordering. $\textsc{LagrangianPrune}$ is applied once before the

Algorithm 1 Method-M for ICN-Aware State Acquisition

Input: requester RSU i, candidate RSU set $\mathcal{R}$, retrieval budget T_{budget}, configuration $\Omega = (\delta, K, \theta, \texttt{stop_mode})$
Output: Selected query set $\mathcal{S}$

$\mathcal{S} \leftarrow \emptyset$, $\tau \leftarrow 0$, $\sigma_{\text{TMIG}} \leftarrow 0$
for each $j \in \mathcal{R}$ **do**
  $u_0[j] \leftarrow \textsc{Utility}(i, j, \emptyset)$ ▷ Initial utility from Eq. (2)
  $t[j] \leftarrow \textsc{LatencyICN}(i, j)$ ▷ ICN-aware latency from Eq. (6)
end for
$A \leftarrow \textsc{LagrangianPrune}(\mathcal{R}, u_0, t, T_{\text{budget}}, \Omega.\delta, \Omega.K)$
if $A = \emptyset$ **then**
  $A \leftarrow \mathcal{R}$
end if
while $A \neq \emptyset$ **and** $\tau < T_{\text{budget}}$ **do**
  for each $j \in A$ **do**
    $u[j] \leftarrow \textsc{Utility}(i, j, \mathcal{S})$ ▷ Recompute novelty given $\mathcal{S}$
  end for
  $b \leftarrow \arg\max_{j \in A}\ u[j]/t[j]$
  if $\tau + t[b] > T_{\text{budget}}$ **then**
    $A \leftarrow A \setminus \{b\}$
    continue
  end if
  $\mathcal{S} \leftarrow \mathcal{S} \cup \{b\}$
  $\tau \leftarrow \tau + t[b]$
  $\sigma_{\text{TMIG}} \leftarrow \sigma_{\text{TMIG}} + u[b]$
  $A \leftarrow A \setminus \{b\}$
  if $\Omega.\texttt{stop_mode} = \texttt{tmig_threshold}$ **and** $\sigma_{\text{TMIG}} \geq \Omega.\theta$ **then**
    break
  end if
end while
return $\mathcal{S}$

Algorithm 2 LagrangianPrune

Input: Candidate set $\mathcal{R}$, initial utilities u, latencies t, budget B, tolerance δ, maximum iterations K
Output: Filtered candidate set A^*

$\ell \leftarrow 0$
$h \leftarrow \max_{c \in \mathcal{R}} u[c]/t[c]$
$A^* \leftarrow \mathcal{R}$
for $k = 1, \ldots, K$ **do**
  $\mu \leftarrow (\ell + h)/2$
  $A \leftarrow \{c \in \mathcal{R} \mid u[c]/t[c] > \mu\}$
  $\tau_A \leftarrow \sum_{c \in A} t[c]$
  if $A = \emptyset$ **then**
    $h \leftarrow \mu$
    continue
  end if
  if $\tau_A \leq B$ **then**
    $A^* \leftarrow A$
    $h \leftarrow \mu$
  else
    $\ell \leftarrow \mu$
  end if
  if $|\tau_A - B| \leq \delta$ **then**
    break
  end if
end for
return A^*

greedy loop using the initial utility u_0 and ICN-aware latency. Since u_0 is computed with $\mathcal{S} = \emptyset$, the pruning step serves only as budget-aware candidate screening and does not imply a global approximation guarantee. Therefore, the output of the pruning step is interpreted as a budget-aware candidate pool, not as the final selected sequence.

The final ordering is determined by the greedy loop. After each RSU is selected, Method-M recomputes the marginal utility of the remaining candidates with respect to the updated selected set $\mathcal{S}$. This refresh updates the JSD-based novelty term and prevents the algorithm from repeatedly selecting RSUs that provide similar particle distributions. Thus, the pruning step reduces the search space, while the greedy loop determines the final selection order based on updated information novelty.

The complexity follows directly from this separation. Let $|C|$ be the number of candidate RSUs and $|A|$ be the number of candidates retained after pruning. Computing initial utilities and latencies requires $O(|C|)$ evaluations, and LagrangianPrune requires $O(|C|K)$, where K is the maximum number of bisection iterations. The greedy loop recomputes utilities over the remaining candidate set after each selection, resulting in $O(|A|^2)$ utility evaluations in the worst case. The overall complexity is therefore $O(|C|K + |A|^2)$, instead of $O(|C|^2 K)$ if pruning were repeated inside every greedy iteration.

V. Evaluation

A. Experimental Setup

a) Simulation environment: We generate road network topologies using SUMO [23]. Each topology is simulated over a 3,600-second horizon with a 300-second warm-up period. Vehicle flows are parameterized with arterial rates of 800 vph and side-road rates of 80 vph. Each RSU aggregates vehicle observations within a 5-second time window into a particle-based distributional representation, which is then bucketized into a unified format for utility computation.

b) Topologies: We evaluate three representative road structures. Topology A is a linear arterial network with 38 RSUs, Topology B is a grid network with 42 RSUs, and Topology C is an irregular mixed network with 56 RSUs. These topologies are selected to cover different ICN-relevant structural properties. Topology A stresses distance decay under limited path redundancy, where longer shortest paths can create larger absolute latency savings after hop reduction. Topology B provides a relatively uniform grid structure, where betweenness centrality is more evenly distributed and ICN hop reduction is less concentrated. Topology C contains more

irregular and hub-like paths, leading to higher betweenness-centrality variance and potentially stronger opportunities for cache reuse and PIT aggregation.

c) Baselines: We compare Method-M with six baselines. **B0 Random-k** randomly selects candidate RSUs until the latency budget is exhausted. **B1 Nearest-first** selects candidates in ascending hop distance from the requester RSU. **B2 Static Utility** ranks candidates by a precomputed utility score without updating redundancy after each selection. **B3 Greedy U/C without ICN** greedily selects candidates by the utility-to-cost ratio using the non-ICN latency model. **B4 ICN-only Hop** selects candidates only by ICN-reduced retrieval latency, isolating the effect of ICN without information-utility optimization. **B5 Oracle-Myopic** uses current-window marginal utility at each greedy step. Since this ground-truth information is unavailable in practice, B5 serves as an upper-bound reference rather than a deployable method.

For ablation, we additionally evaluate **Method-M (no ICN)** and **Method-M (no JSD)**.

d) Evaluation suites and metrics: We use two complementary evaluation settings. The *parameter-sweep* setting evaluates information gain under a latency budget $T_{\text{budget}} \in \{20, 40, 60, 80, 100, 120\}$ ms and reports average TMIG. This setting sweeps the available retrieval budget to examine the information–latency trade-off under different communication constraints. The same budgeted setting is also used for component ablation.

The *target_info* suite uses a non-binding latency budget and stops once a fixed target TMIG threshold is reached, measuring the latency required to reach the same information target. Specifically, T_{budget} is set to 999999 ms, so the latency budget does not truncate the selection process. We set the target threshold to $\theta = 0.45$. During selection, the accumulated TMIG is computed as $\sigma_{\text{TMIG}} = \sum_{m=1}^{k} U_{j_m|i,\mathcal{S}_{m-1}}$, and the process stops once $\sigma_{\text{TMIG}} \geq \theta$.

Across the experiments, we sweep $\lambda \in \{0.05, 0.1, 0.3, 0.5, 1.0, 2.0\}$, $\Delta \in \{0.5, 1, 2, 3, 5, 8\}$, and $\alpha \in \{0.08, 0.15, 0.25\}$. Each configuration is repeated three times under realized latency accounting.

The latency advantage is most visible in Topologies A and C. In Topology A, long baseline path lengths amplify the absolute benefit of hop reduction, while in Topology C, hub-like paths increase betweenness-centrality variance and create stronger PIT-aggregation and cache-reuse opportunities. In the more uniform grid topology, latency differences are smaller because ICN hop reduction is less concentrated.

B. Baseline Comparison

a) Information gain: Figure 4 reports paired comparison of $\Delta\text{TMIG} = \text{TMIG}_{\text{Method-M}} - \text{TMIG}_{\text{baseline}}$ in the parameter-sweep suite. Each paired comparison is computed under the same topology, requester RSU, parameter configuration, random repeat, and latency budget T_{budget}. The results are first evaluated separately for each swept budget value $T_{\text{budget}} \in \{20, 40, 60, 80, 100, 120\}$ ms. We then aggregate all paired samples across the swept budget values and report the

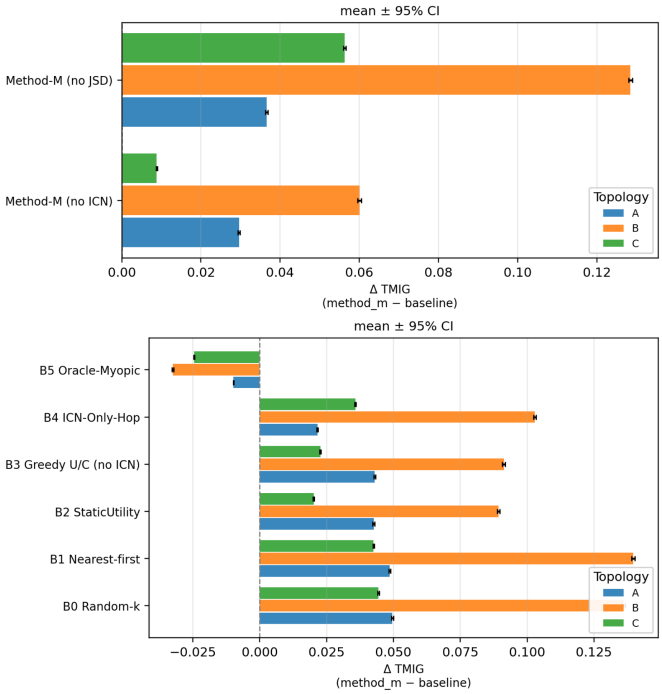


Fig. 4. Paired ΔTMIG (Method-M minus baseline), mean ± 95% CI. Top: ablation variants; bottom: all baselines.

TABLE I
BASELINE COMPARISON

Method	parameter-sweep ΔTMIG	target_info Lat.(ms)	target_info ΔLat.
Method-M	–	731.5	–
B0 Random-k	−23.1%	2536.1	+246.7%
B1 Nearest	−23.2%	1681.4	+129.9%
B2 Static	−14.9%	1114.5	+52.4%
B3 Greedy U/C	−15.4%	1226.0	+67.6%
B4 ICN-hop	−16.4%	927.2	+26.8%
B5 Oracle	+6.7%	718.7	−1.8%

overall mean and confidence interval. Therefore, the figure summarizes Method-M's average information-gain advantage over the full parameter-sweep budget range, rather than under a single fixed budget. Method-M consistently outperforms all practical baselines, and the confidence intervals for B0–B4 exclude null hypothesis. As summarized in Table I, Method-M improves over B3 (Greedy U/C without ICN) by 15.4%, B4 (ICN-only Hop) by 16.4%, and B0 (Random-k) by 23.1% in mean TMIG. This shows that neither ICN-aware latency alone nor utility-to-cost selection alone is sufficient; Method-M benefits from jointly considering information novelty and ICN-aware retrieval cost.

Figure 5 further shows that Method-M's advantage is stable across the swept values of α, λ, Δ, and the latency budget. Compared with B5 Oracle-Myopic, Method-M trails by 6.7% in large-scale TMIG. This is expected because B5 uses current-window ground-truth marginal utility, whereas Method-M relies only on the previous synchronized belief state. Therefore, B5 should be interpreted as an upper-bound reference rather

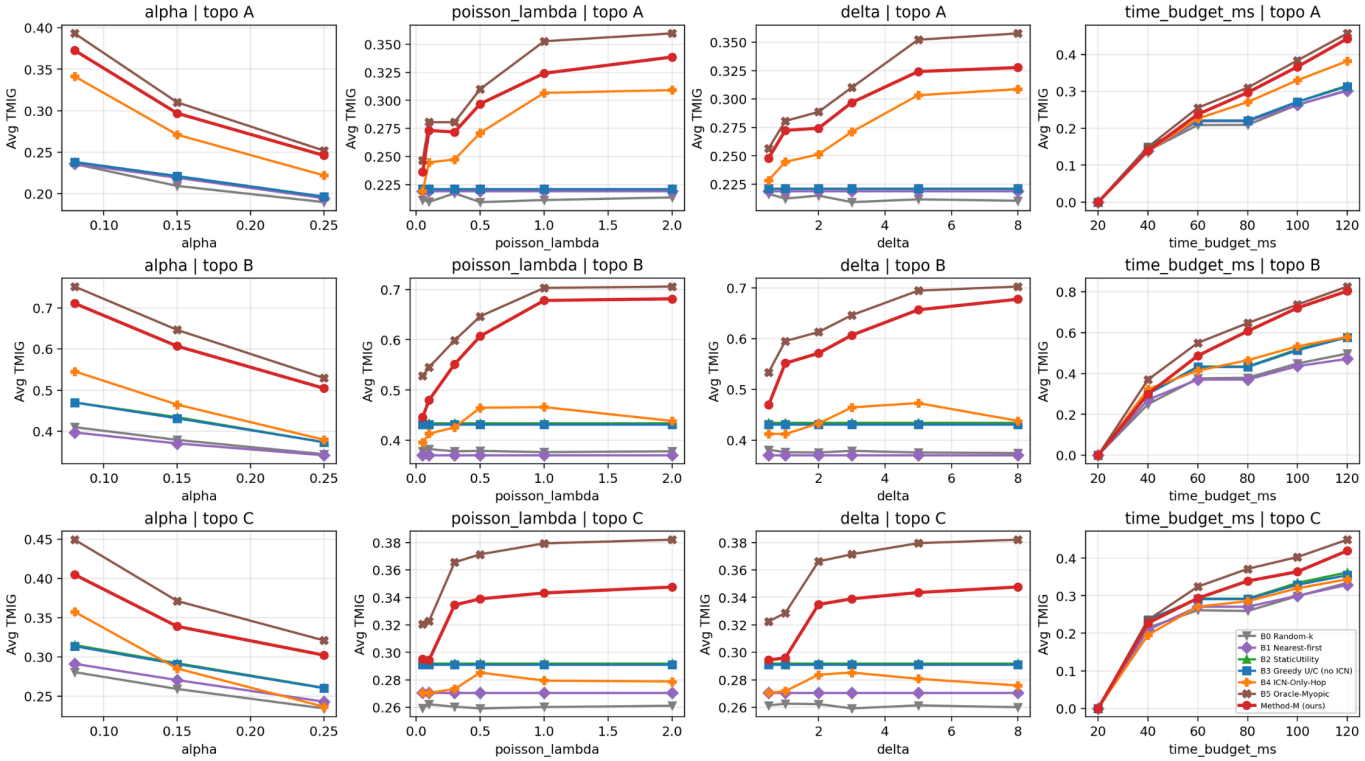


Fig. 5. Control sensitivity on large-scale topologies in the parameter-sweep setting. Method-M dominates the practical baselines across control parameters and topology types.

TABLE II
TOPOLOGY-WISE PARAMETER-SWEEP RESULTS.

Topology	Method-M	Best practical	Gain
A	0.2377	0.2161	+10.0%
B	0.4601	0.3708	+24.1%
C	0.2704	0.2502	+8.1%

TABLE III
ABLATION SUMMARY.

Variant	ΔTMIG$_p$	ΔLat.$_p$	ΔTMIG$_t$	ΔLat.$_t$
w/o ICN	-9.2%	-6.8%	-1.2%	$+81.8\%$
w/o JSD	-22.2%	$+6.4\%$	-10.6%	$+8.6\%$

than a deployable method.

b) Latency to target information: Table I also reports the latency required to reach the target TMIG $\theta = 0.45$ in the target_info suite. Method-M achieves the lowest latency among practical methods. It reduces latency by 26.8% compared with B4 (ICN-only Hop), 52.4% compared with B2 (Static Utility), 67.6% compared with B3 (Greedy U/C without ICN), and 246.7% compared with B0 (Random-k). Method-M is also close to B5 Oracle-Myopic, trailing by only 1.8% in latency while avoiding the need for current-window ground-truth state information.

Table II further reports topology-wise parameter-sweep results. Method-M improves over the best practical baseline in all three topologies, but the magnitude of the gain varies with network structure. The largest gain appears in Topology B, where the uniform grid makes simple locality-based or ICN-only selection less discriminative and increases the value of utility-aware novelty selection. Topology A and Topology C show smaller but still consistent gains, indicating that Method-M remains effective under both long-path arterial structures and irregular hub-like structures.

C. Ablation Study and Robustness

a) Ablation study: Table III separates the effects of ICN-aware cost reduction and JSD-based novelty. In the table, subscript p denotes the parameter-sweep setting and subscript t denotes the target_info setting. Under the parameter-sweep setting, removing JSD causes a 22.2% TMIG drop, much larger than the 9.2% drop caused by removing ICN.

In contrast, under the target_info suite, removing ICN increases the latency to reach the same information target by 81.8%, while removing JSD increases it by only 8.6%. This confirms the intended separation of roles: JSD improves what information is selected by reducing redundancy, whereas ICN improves how efficiently the selected information can be retrieved.

b) Parameter robustness: Finally, Figure 5 shows that Method-M remains robust across the full parameter ranges: $\lambda \in [0.05, 2.0]$, $\Delta \in [0.5, 8]$, $\alpha \in [0.08, 0.25]$, and $T_{\text{budget}} \in [20, 120]$ ms. The line-plot results show that Method-M consistently outperforms practical baselines across the tested ranges,

indicating that its advantage is not tied to a single favorable setting. As predicted by Eq. (5), increasing λ or Δ generally increases the expected ICN hop-reduction ratio by creating stronger temporal overlap among Interests, while non-ICN baselines remain mostly insensitive to these two parameters. The sweeps over α and T_{budget} further show that Method-M remains effective under different distance-decay strengths and latency-budget regimes.

VI. Conclusion

This paper proposes to introduce ICN to V2X edge network and to apply Method-M, an ICN-aware partial state acquisition framework for V2X edge networks under communication constraints. We formulated RSU selection as a latency-constrained information acquisition problem, modeled ICN-induced hop reduction through temporal request reuse and topological path overlap, and combined JSD-based novelty with Lagrangian-guided greedy selection.

Evaluation on SUMO-generated topologies shows that Method-M consistently improves information gain over practical baselines and approaches the current-state Oracle without requiring current-window global knowledge. Ablation results further confirm the complementary roles of JSD and ICN: JSD improves information quality by reducing redundancy, while ICN improves retrieval efficiency by lowering the latency required to reach a target information level. Overall, the results show that partial state acquisition should jointly consider information utility and communication cost.

The current evaluation uses a topology- and request-aware latency abstraction rather than packet-level ICN simulation and assumes that recent Interest statistics and RSU-level topology information are available for retrieval-cost estimation. Future work will therefore investigate estimation errors, cache freshness, non-stationary request patterns, and packet-level validation with an ICN/NDN simulator.

Acknowledgments

This work was partly supported by the National Institute of Information and Communications Technology (NICT), Japan, under Grant 24101.

References

[1] J. Liu, M. Ahmed, M. A. Mirza, W. U. Khan, D. Xu, J. Li, A. Aziz, and Z. Han, "Rl/drl meets vehicular task offloading using edge and vehicular cloudlet: A survey," *IEEE Internet of Things Journal*, vol. 9, no. 11, pp. 8315–8338, 2022.

[2] K. Xiong, S. Leng, C. Huang, C. Yuen, and Y. L. Guan, "Intelligent task offloading for heterogeneous v2x communications," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 4, pp. 2226–2238, 2020.

[3] A. A. Ismail, N. E. Khalifa, and R. A. El-Khoribi, "A survey on resource scheduling approaches in multi-access edge computing environment: a deep reinforcement learning study," *Cluster Computing*, vol. 28, no. 3, p. 184, 2025.

[4] P. Li, X. Wang, C. Li, M. Iqbal, A. Al-Dulaimi, P. Casaseca-de-la Higuera *et al.*, "Deep reinforcement learning-based task scheduling and resource allocation for vehicular edge computing: A survey," *IEEE Transactions on Intelligent Transportation Systems*, 2025.

[5] S. A. Ullah, M. Bibi, S. A. Hassan, H. Abou-Zeid, H. K. Qureshi, H. Jung, A. Mahmood, M. Gidlund, and E. Hossain, "From nodes to roads: Surveying drl applications in mec-enhanced terrestrial wireless networks," *IEEE Communications Surveys & Tutorials*, 2025.

[6] M. Rosencrantz, G. Gordon, and S. Thrun, "Decentralized sensor fusion with distributed particle filters," *arXiv preprint arXiv:1212.2493*, 2012.

[7] N. E. Marinică, A. Sarlette, and R. K. Boel, "Distributed particle filter for urban traffic networks using a platoon-based model," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 4, pp. 1918–1929, 2013.

[8] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A survey of information-centric networking," *IEEE Communications Magazine*, vol. 50, no. 7, pp. 26–36, 2012.

[9] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *Proceedings of the 5th international conference on Emerging networking experiments and technologies*, 2009, pp. 1–12.

[10] M. Amadeo, C. Campolo, J. Quevedo, D. Corujo, A. Molinaro, A. Iera, R. L. Aguiar, and A. V. Vasilakos, "Information-centric networking for the internet of things: challenges and opportunities," *IEEE Network*, vol. 30, no. 2, pp. 92–100, 2016.

[11] S. Gyawali, S. Xu, Y. Qian, and R. Q. Hu, "Challenges and solutions for cellular based v2x communications," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 1, pp. 222–255, 2020.

[12] S. Raza, S. Wang, M. Ahmed, M. R. Anwar, M. A. Mirza, and W. U. Khan, "Task offloading and resource allocation for iov using 5g nr-v2x communication," *IEEE Internet of Things Journal*, vol. 9, no. 13, pp. 10 397–10 410, 2021.

[13] P. M. Djuric, J. H. Kotecha, J. Zhang, Y. Huang, T. Ghirmai, M. F. Bugallo, and J. Miguez, "Particle filtering," *IEEE signal processing magazine*, vol. 20, no. 5, pp. 19–38, 2003.

[14] M. Bolic, P. M. Djuric, and S. Hong, "Resampling algorithms and architectures for distributed particle filters," *IEEE Transactions on Signal Processing*, vol. 53, no. 7, pp. 2442–2450, 2005.

[15] M. Coates, "Distributed particle filters for sensor networks," in *Proceedings of the 3rd international symposium on Information processing in sensor networks*, 2004, pp. 99–107.

[16] L. Yao, Y. Wang, X. Wang, and G. Wu, "Cooperative caching in vehicular content centric network based on social attributes and mobility," *IEEE Transactions on Mobile Computing*, vol. 20, no. 2, pp. 391–402, 2019.

[17] H. Godrich, A. P. Petropulu, and H. V. Poor, "Sensor selection in distributed multiple-radar architectures for localization: A knapsack problem formulation," *IEEE Transactions on Signal Processing*, vol. 60, no. 1, pp. 247–260, 2011.

[18] S. Liu, S. P. Chepuri, G. Leus, and A. O. Hero, "Distributed sensor selection for field estimation," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 4257–4261.

[19] G. Chai, W. Wu, Q. Yang, R. Liu, and F. R. Yu, "Learning-based resource allocation for ultra-reliable v2x networks with partial csi," *IEEE Transactions on Communications*, vol. 70, no. 10, pp. 6532–6546, 2022.

[20] M. Dong, Y. Fu, C. Li, M. Tian, F. R. Yu, and N. Cheng, "Task offloading and resource allocation in vehicular cooperative perception with integrated sensing, communication, and computation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 26, no. 6, pp. 8481–8496, 2025.

[21] S. Liu, L. Liu, J. Tang, B. Yu, Y. Wang, and W. Shi, "Edge computing for autonomous driving: Opportunities and challenges," *Proceedings of the IEEE*, vol. 107, no. 8, pp. 1697–1716, 2019.

[22] K. Manolakis and W. Xu, "Time synchronization for multi-link d2d/v2x communication," in *2016 IEEE 84th Vehicular Technology Conference (VTC-Fall)*. IEEE, 2016, pp. 1–6.

[23] P. Alvarez Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using SUMO - Simulation of Urban Mobility," in *The 21st IEEE Intelligent Transportation Systems Conference (ITSC)*, Nov 2018, pp. 217–224.