Maximising Data Rate in HAP Assisted Cooperative Vehicular Networks with Optimised Relay Selection

Kayode Popoola[1*], David Grace[2†], Tim Clarke[2†], Muheeb Ahmed[1†], Stuart Nicholson[1†]

[1*]Department of Computer Science, Dyson Institute of Engineering and , Malmesbury, SN16 0RP, Wiltshire, United Kingdom.
[2]School of Physics, Engineering and Technology, University of York, Heslington, York, YO10 5EZ, North Yorkshire, United Kingdom.

*Corresponding author(s). E-mail(s): kayode.popoola@dysoninstitute.ac.uk;
Contributing authors: david.grace@york.ac.uk; tim.clarke@york.ac.uk; muheeb.ahmed@dysoninstitute.ac.uk; stuart.nicholson@dysoninstitute.ac.uk;
†These authors contributed equally to this work.

Abstract

In this paper, we address the challenge of improving network performance in High-Altitude Platform-assisted cooperative vehicular networking (HAP-CVN) through optimised relay selection. We propose the Kuhn-Munkres Relay Assignment (KMRA) algorithm, an approach based on a weighted bipartite matching framework. Our primary objective is to maximise sum transmission rates for vehicle user equipment (VUE) experiencing out-of-coverage or cell-edge channel conditions, which in turn significantly minimises their outage probability. By formulating the relay selection as a maximum weighted bipartite matching problem, our approach optimally pairs relays to destinations, outperforming conventional benchmark schemes. The proposed KMRA algorithm delivers significant performance improvements, particularly in terms of reduced outage probability, enhanced transmission success probability, and increased user throughput.

Keywords: Cooperative vehicular networks, relay selection, relay assignment, bipartite matching, KMRA, VUE, outage probability, throughput.

1 Introduction

Recent advances in vehicular networking have driven intense research and development efforts toward seeking reliable, high-capacity, and ultra-low-latency communications for connected and autonomous vehicles. Such networks, generally referred to as Cooperative Vehicular Networks (CVNs), leverage cooperation between vehicle user equipment (VUE) to enhance performance metrics such as transmission rate, coverage probability, and network connectivity [1, 2].

In High-altitude Platform (HAP)-assisted CVNs (HAP-CVN), vehicular relays are essential for mitigating connectivity gaps. HAPs offer a compelling solution for vehicular networking due to their ability to provide ubiquitous coverage over large geographical areas, overcome terrain-induced blockages common in terrestrial networks, and support rapid deployment for disaster relief or temporary coverage extension [3, 4]. These characteristics make them particularly suitable for serving VUEs in rural or underserved regions where consistent terrestrial connectivity is often lacking, especially for those at cell edges or outside direct coverage areas. While VUEs within the HAP's direct coverage can establish reliable communication, those at the cell edge or outside the coverage area often experience severe outages or significantly reduced throughput. Vehicular relay User Equipment (UEs) address this challenge by assisting in data forwarding, thereby improving connectivity and network performance. Cooperative relaying, where selected relay VUEs forward data from infrastructure nodes (e.g., HAPs, RSUs, or base stations) or other VUEs to users with poor connectivity [5], effectively extends coverage, enhances throughput, and increases network robustness.

However, the effectiveness of cooperative relaying is highly dependent on optimal relay selection. Poorly chosen relays can introduce additional interference, suffer from unstable links, or inefficiently utilise network resources, ultimately degrading network performance. This is particularly important in HAP-CVNs where the large coverage area can lead to a diverse range of channel conditions and significant numbers of VUEs at the cell edge or in coverage holes. Ensuring reliable, high-rate connectivity for these out-of-coverage VUEs through efficient cooperative relaying is therefore important, but presents unique challenges in optimal relay selection. Consequently, the critical challenge lies in dynamically selecting the most suitable relay UEs to maximise transmission efficiency while minimising potential network disruptions. To this end, this paper proposes an optimised relay selection framework tailored to the unique characteristics and demands of HAP-CVNs. However, while various relay selection schemes have been proposed for general wireless and terrestrial vehicular networks [6, 7], they do not adequately address the specific constraints and opportunities inherent in HAP-assisted architectures. For instance, the quasi-stationary position of the HAP provides a unique vantage point for centralised coordination, yet existing distributed heuristics may not fully leverage this. Furthermore, specific challenges such as the wide HAP footprint, distinct 3D channel characteristics between HAP and ground VUEs, and the need for computationally efficient assignment algorithms scalable to potentially large numbers of VUEs under HAP coverage remain open areas. Furthermore, existing approaches often lack a holistic optimisation tailored for sum-rate maximisation while considering these HAP-CVN specific factors. Our work aims to address this gap.

In this work, we focus on evaluating an optimised relay selection scheme in a HAP-CVN, specifically motivated by a scenario where a HAP serves ground VUEs. VUEs with strong channel conditions to the HAP are identified as candidate relay UEs for out-of-coverage or cell-edge VUEs. We formulate the relay selection process as a maximum weighted bipartite matching problem, in which the weights correspond to achievable data rates between potential relay-destination VUE pairs. We then apply the Kuhn-Munkres algorithm to find an assignment that maximises the total data rate while respecting interference and coverage constraints. Our proposed Kuhn-Munkres Relay Assignment (KMRA) approach is tested against several baseline approaches, demonstrating superior performance in terms of cell-edge user throughput and outage probability.

1.1 Key Contributions

- We present the application of the Kuhn-Munkres algorithm for relay selection (KMRA) in a vehicular network served by a quasi-stationary HAP. By modelling the network as a weighted bipartite graph, we pose the relay selection as a maximum weighted matching problem, specifically addressing the unique link characteristics of HAP-served vehicular networks.
- We provide a comprehensive performance evaluation of the proposed KMRA approach under varying signal-to-interference-plus-noise ratio (SINR) and interference conditions. Our algorithm is compared with three benchmark methods: Threshold-Based Relay Assignment (TBRA), Link Lifetime Relay Selection (LLRS), and Random Relay Assignment (RRA). Simulation results show the superior performance in terms of outage probability, transmission success probability and achievable user throughput.

The remainder of this paper is organised as follows. Section 2 reviews related work on relay selection in vehicular networks. In Section 3, we describe our system model, including channel propagation assumptions and the cooperative transmission protocols. Section 4 provides the formal optimisation problem for relay selection. Section 5 presents our bipartite matching-based relay selection methodology. In Section 6, we present the simulation environment and setup. Section 7 details the performance metric and benchmark schemes used for evaluation. Section 8 presents and discusses the simulation results. Finally, Section 9 concludes the paper and suggests future research directions.

2 Related Work

Relay networks play a crucial role in enhancing wireless communication by extending coverage, improving signal reliability, and mitigating the effects of fading. Traditional relay-based approaches have been widely studied in terrestrial networks, including fixed infrastructure relays that improve connectivity in dynamic environments. In the context of vehicular networks, cooperative relay communications have attracted significant attention as an effective strategy to improve network coverage, throughput, and reliability, especially in out-of-coverage or cell-edge areas [8, 9]. Similarly, high-altitude platforms (HAPs) have also been proposed as a promising solution for enhancing communication, particularly in areas where terrestrial infrastructure is limited or during emergencies when ground networks are unavailable. As aerial base stations, HAPs enable reliable Vehicle-to-Infrastructure (V2I) communications, supporting applications such as real-time traffic management, safety alerts, and infotainment services. Numerous studies have explored the integration of HAPs into vehicular networks [10–12].

Furthermore, research in related aerial and sensor networking domains highlights other important performance aspects. For instance, energy efficiency is critical for battery-powered nodes, with studies addressing energy-aware multihop, multipath cooperative routing [13] and the impact of transceiver hardware [14]. Similarly, for mobile platforms like UAVs, which share some characteristics with HAPs, trajectory optimisation is key to ensuring both safety and energy efficiency with research in [15] proposing an energy-efficient trajectory planning framework which enhances UAV flight endurance by optimising energy consumption during operation. Secure data transmission and the design of resilient relaying techniques also remain critical considerations in cooperative networks [15]. While our current work focuses on optimising for rate and connectivity, we acknowledge these areas as important dimensions for future, more comprehensive frameworks.

Existing relay selection methods for vehicular networks often focus on specific terrestrial or UAV-based scenarios. For instance, [16] considered a multi-hop bi-directional relay selection based on a link estimation algorithm, demonstrating improved packet delivery and reduced delay in dense vehicular networks. Other approaches have utilised metrics like direction of travel, traffic density, or received signal power [17] for selecting the best relay. Research has also explored maximising CVN capacity through relay assignment, with [18] suggesting that a one-to-one assignment between relay and destination VUEs is often ideal for total capacity. Furthermore, multi-metric methods incorporating mobility and received signal strength [19] and "nearest-first" systems based on distance [20] have been assessed for routing performance.

While these prior investigations offer valuable insights into terrestrial or UAV-based relay networks [21, 22], they often do not fully address the unique characteristics and constraints of HAP-CVNs. HAP systems introduce specific challenges such as wide coverage patterns, altitude-dependent path loss, and unique power consumption trade-offs or antenna payload considerations [22, 23]. These factors significantly influence channel conditions, interference management, and optimal relay placement. Motivated by these distinct challenges, we propose a bipartite matching perspective that simplifies the combinatorial assignment problem while effectively maintaining interference and link constraints specifically for a HAP-CVN scenario.

3 System Model

We consider a High-altitude Platform (HAP) that provides wireless access over a large coverage region. Let the total set of Vehicle User Equipment (VUEs) be $\mathcal{N} = \{1, 2, \ldots, N\}$. A subset of these VUEs, denoted by $\mathcal{H}$, can be served directly by the

HAP, satisfying a minimum Signal-to-Interference-plus-Noise Ratio (SINR) threshold $\chi_{\text{HAP-direct}}$. The VUEs that cannot be served directly due to poor channel conditions (e.g., out-of-coverage or at the cell edge) form the set of destination VUEs $\mathcal{T} = \{t_1, t_2, \ldots, t_{|\mathcal{T}|}\}$. Our goal is to relay data to these destination VUEs $\mathcal{T}$ via cooperative transmissions from a subset of the "good" VUEs in $\mathcal{H}$. We define $\mathcal{M}$ as the set of candidate relay VUEs (a subset of $\mathcal{H}$), such that:

$$\mathcal{M} = \{k_j \in \mathcal{H} : \tilde{\gamma}_{\text{HAP},k_j} \geq \chi_{\text{HAP-direct}}\}.$$

The overall system architecture is as depicted in Fig. 1.

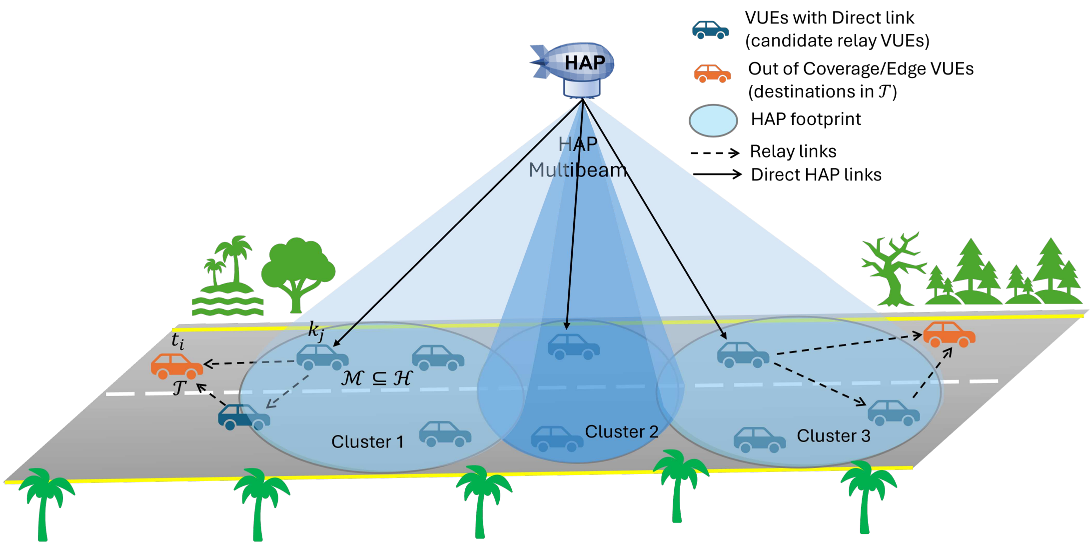

Fig. 1 System model for HAP-assisted Cooperative Vehicular Network. The HAP provides beam coverage to vehicular clusters on the ground. VUEs within the beam coverage footprint have good direct links and form the set of candidate relay UEs ($\mathcal{M} \subset \mathcal{H}$). VUEs at the edge or outside coverage form the set of destinations $\mathcal{H}$. The full two-hop path (HAP-to-relay k_j and relay k_j to destination t_i) is shown.

Here, $\tilde{\gamma}_{\text{HAP},k_j}$ denotes the SINR between the HAP and candidate relay VUE k_j. The final set of chosen relay VUEs is $\mathcal{R} \subseteq \mathcal{M}$. Each relay VUE $k_j \in \mathcal{R}$ is assigned to forward data to exactly one destination VUE $t_i \in \mathcal{T}$ (i.e., a one-to-one assignment).

3.1 Cooperative Relaying Protocols

Various relaying protocols have been proposed in the literature, including Amplify-and-Forward (AF) and Decode-and-Forward (DF) [5]. In this work, we adopt the DF protocol due to its generally superior link reliability, which is crucial for minimising the outage probability, albeit at the cost of potentially higher relay complexity. We assume a half-duplex relaying operation, where the HAP-to-relay (first hop) and relay-to-destination (second hop) transmissions occur in orthogonal time slots using the same channel bandwidth B. Thus, B represents the bandwidth available for transmission in each hop during its active phase. For a two-hop DF link where the HAP transmits to a destination VUE t_i via a relay VUE k_j, the equivalent SINR at the destination is limited by the weaker of the two hops. This can be modelled with an effective end-to-end SINR:

$$\tilde{\gamma}_{\mathrm{eff},k_j,t_i} = \min\{\, \tilde{\gamma}_{\mathrm{HAP},k_j},\ \tilde{\gamma}_{k_j,t_i}\}. \tag{1}$$

Hence, the achievable data rate (in bps/Hz) for such a DF link is:

$$r_{k_j,t_i} \; = \; B \log_2\bigl(1 + \tilde{\gamma}_{\mathrm{eff},k_j,t_i}\bigr), \tag{2}$$

where B is the system bandwidth utilised during each transmission slot. It is important to note that due to the two orthogonal time slots required for this half-duplex relaying scheme, the effective end-to-end throughput for a complete message relayed over both hops is $0.5 \times r_{k_j,t_i}$ compared to a direct link that could transmit the same message content using bandwidth B in a single equivalent time slot. The resource utilisation and rate implications of such cooperative schemes, including the impact of orthogonalisation, are extensively discussed in [24]. Our formulation of r_{k_j,t_i} therefore represents the instantaneous achievable rate during the active relaying phase over bandwidth B.

For direct transmissions by the HAP to a VUE t_i, the transmission rate is obtained as:

$$r_{\mathrm{direct},t_i} = B \log_2\bigl(1 + \tilde{\gamma}_{\mathrm{HAP},t_i}\bigr).$$

If $\tilde{\gamma}_{\mathrm{HAP},t_i}$ is below a pre-defined link quality threshold or if the direct link is in an outage, we rely on the relay path via the chosen set $\mathcal{R}$.

3.2 Channel Model, Interference, and Link Duration

The overall system performance is heavily influenced by the channel conditions, which encompass path loss, shadowing, small-scale fading, and interference. Furthermore, vehicular mobility dynamically affects link connectivity duration. We detail our channel model and interference assumptions for both HAP-to-VUE and V2V communication links below.

3.2.1 Channel Model

The instantaneous channel power gain between any two communicating entities i and j (e.g., HAP to VUE, VUE to VUE) at time τ is expressed as:

$$g_{i,j}(\tau) = L(d_{i,j}(\tau))\, S_{i,j}(\tau)\, H_{i,j}(\tau). \tag{3}$$

Here, $L(d_{i,j}(\tau))$ denotes the path-loss component as a function of the distance $d_{i,j}(\tau)$ between i and j, $S_{i,j}(\tau)$ represents the log-normal shadowing component, and $H_{i,j}(\tau)$ captures the small-scale fading effects.

The path loss in dB, $L_{\mathrm{dB}}(d)$, as a function of distance d is modelled as:

$$L_{\mathrm{dB}}(d) = L_{0,\mathrm{dB}} + 10\alpha \log_{10}\left(\frac{d}{d_0}\right), \tag{4}$$

where $L_{0,\mathrm{dB}}$ is the path loss dB at a reference distance d_0, and α is the path loss exponent. The linear path-loss component is then $L(d) = 10^{-L_{\mathrm{dB}}(d)/10}$. The shadowing component $S_{i,j}(\tau)$ is modelled such that $10\log_{10}(S_{i,j}(\tau))$ is a zero-mean Gaussian random variable with standard deviation $\sigma_{\mathrm{sh,dB}}$ (in dB). The small-scale fading power gain $H_{i,j}(\tau)$ is modelled according to an appropriate statistical distribution (Rayleigh or Shadowed Rician, depending on the presence of a line-of-sight path). While the Nakagami-m fading model offers greater generality and has been used for vehicular channel propagation models as it can represent a wider range of fading conditions [25],its *m parameter* introduces an additional layer of complexity. Hence in this study and similar to works in [26, 27], we adopt the Shadowed Rician (for Line-of-Sight, LOS) and Rayleigh (for Non-Line-of-Sight, NLOS) distributions for describing the HAP downlink propagation. This allows for a clear distinction in modelling the different propagation characteristics expected for HAP-to-VUE (often LOS-dominant) and V2V (variable LOS/NLOS) links, while maintaining model tractability.

3.2.2 HAP-to-VUE Links

For the HAP-to-VUE link (e.g., HAP to candidate relay k_j), the Signal-to-Interference-plus-Noise Ratio (SINR) at time τ is given by:

$$\tilde{\gamma}_{\mathrm{HAP},k_j}(\tau) = \frac{P_{\mathrm{HAP},k_j}(\tau)\, g_{\mathrm{HAP},k_j}(\tau)}{N_p + I^{\mathrm{ext}}_{\mathrm{HAP},k_j}(\tau)}, \tag{5}$$

$$\text{where } I^{\mathrm{ext}}_{\mathrm{HAP},k_j}(\tau) = \sum_{m \in \mathcal{I}_{\text{HAP-ext}}} P_{m,k_j}(\tau)\, g_{m,k_j}(\tau). \tag{6}$$

Here, $P_{\mathrm{HAP},k_j}(\tau)$ is the time-varying transmit power from the HAP to VUE k_j, and $g_{\mathrm{HAP},k_j}(\tau)$ is the channel power gain as defined in (3). N_p is the thermal noise power. The term $I^{\mathrm{ext}}_{\mathrm{HAP},k_j}(\tau)$ in (6) represents the aggregated co-channel interference power at k_j. The set $\mathcal{I}_{\text{HAP-ext}}$ comprises all significant external co-channel interferers, which may include other beams from the HAP platform. If this external interference is negligible compared to N_p, the SINR approaches the Signal-to-Noise Ratio (SNR).

3.2.3 V2V Links

For the V2V link where a relay VUE k_j serves a destination VUE t_i, the SINR at t_i from k_j at time τ is expressed as:

$$\tilde{\gamma}_{k_j,t_i}(\tau) = \frac{P_{k_j,t_i}(\tau)\, g_{k_j,t_i}(\tau)}{N_p + I^{\mathrm{relay}}_{k_j,t_i}(\tau)}, \tag{7}$$

$$\text{where } I^{\mathrm{relay}}_{k_j,t_i}(\tau) = \sum_{x \in \mathcal{R}, x \neq k_j} P_{x,t_i}(\tau)\, g_{x,t_i}(\tau). \tag{8}$$

Here, $P_{k_j,t_i}(\tau)$ denotes the transmit power from relay k_j to destination VUE t_i, and $g_{k_j,t_i}(\tau)$ is the channel power gain. The term $I^{\text{relay}}_{k_j,t_i}(\tau)$ represents the aggregated interference power at t_i from other concurrently active relays that are transmitting on the same resource, potentially to other destination VUEs. This models inter-cluster interference if resources are reused across clusters.

3.2.4 Interference Management and Link Viability

Within each logical cluster of VUEs, orthogonal radio resources are typically allocated to minimise intra-cluster interference. However, inter-cluster interference, primarily from relays in adjacent clusters reusing the same frequency resources, remains a key factor and is accounted for in the SINR calculation for V2V links (specifically, in the term $I^{\text{relay}}_{k_j,t_i}(\tau)$ of (8)).

To further manage the impact of this inter-cluster interference on the relay selection process, we introduce an interference threshold, I_{th}. A potential V2V relay link from k_j to t_i is considered viable for the current assignment snapshot only if the estimated aggregated interference power at the receiver t_i, $I^{\text{relay}}_{k_j,t_i}(\tau)$, does not exceed this threshold I_{th}. That is, $I^{\text{relay}}_{k_j,t_i}(\tau) \leq I_{\text{th}}$ must hold. This condition is an additional prerequisite for an edge to exist in the bipartite graph used for KMRA, as further detailed in Sections 4 and 5. This allows us to study the performance of relay selection strategies under varying levels of permissible interference. Beyond this snapshot-based viability check, persistent high interference exceeding certain operational limits might trigger broader resource reallocation by the HAP, but that is outside the scope of our relay assignment problem.

3.3 V2V Connectivity and Link Duration

Due to the high mobility of vehicular UEs, channel gains $g_{x,y}(\tau)$ and consequently, SINR values vary rapidly with time τ. A critical aspect for effective V2V communication in the relay-assisted network is the link-contact duration: the relay VUE k_j must remain within effective communication range of the destination VUE t_i for a period sufficient to complete the intended data exchange.

We define a link-contact duration, C_{k_j,t_i}, based on current relative distance and velocity:

$$C_{k_j,t_i} = \begin{cases} \dfrac{d_{k_j,t_i}(\tau)}{|v_{\text{rel}}(\tau)|}, & \text{if } d_{k_j,t_i}(\tau) < R_{\text{eff}} \\ \quad \text{and } v_{\text{rel}}(\tau) > 0 \text{ indicates separation}, & \\ T_{\text{stable}}, & \text{if } d_{k_j,t_i}(\tau) < R_{\text{eff}} \\ \quad \text{and } v_{\text{rel}}(\tau) \leq 0 \text{ indicates stable link}, & \\ 0, & \text{if } d_{k_j,t_i}(\tau) \geq R_{\text{eff}}. \end{cases} \tag{9}$$

Here $d_{k_j,t_i}(\tau)$ is the instantaneous distance between relay k_j and destination t_i, $v_{\text{rel}}(\tau)$ is their relative velocity, and R_{eff} the maximum effective communication range for a V2V link to maintain a minimum quality. T_{stable} represents a sufficiently long duration assumed if vehicles are getting closer or maintaining distance within R_{eff}.

For the purpose of our relay selection problem (detailed in Section 5), a potential relay-destination pair (k_j, t_i) is considered a viable candidate for matching only if its estimated link-contact duration C_{k_j,t_i} exceeds a minimum required service time, T_{service}. That is, an edge (k_j, t_i) exists in the bipartite graph if $C_{k_j,t_i} > T_{\text{service}}$. The choice of T_{service} and further implications of this link duration model are discussed in Section 7.

4 Problem Formulation

The primary goal is to assign each destination VUE $t_i \in \mathcal{T}$ (those with poor direct HAP connectivity) to at most one candidate relay VUE $k_j \in \mathcal{M}$ (those with strong HAP connectivity) to maximise the total achievable data rate for the relayed transmissions. While other objectives like fairness (e.g., maximising the minimum rate) are possible, this work focuses on maximising the aggregate throughput.

Let s_{k_j,t_i} be a binary decision variable:

$$s_{k_j,t_i} = \begin{cases} 1, & \text{if relay VUE } k_j \text{ is assigned to destination VUE } t_i, \\ 0, & \text{otherwise.} \end{cases}$$

The achievable data rate r_{k_j,t_i} for a specific relay-destination pair (k_j, t_i), using the Decode-and-Forward (DF) protocol and considering the bandwidth B available during each active transmission slot, is determined by the minimum Signal-to-Interference-plus-Noise Ratio (SINR) of the two hops involved:

$$r_{k_j,t_i} = B \log_2 \left(1 + \min\{\tilde{\gamma}_{\text{HAP},k_j}, \tilde{\gamma}_{k_j,t_i}\}\right). \tag{10}$$

Here, $\tilde{\gamma}_{\text{HAP},k_j}$ and $\tilde{\gamma}_{k_j,t_i}$ denote the SINRs of the HAP-to-relay k_j and relay k_j-to-destination t_i links, respectively (as detailed in Section 3.1). The variable B is the system bandwidth used per transmission slot. Due to half-duplex operation, the effective end-to-end data rate is $0.5 \times r_{k_j,t_i}$, since relaying requires two orthogonal time slots. The rate r_{k_j,t_i} in Eq. (10) represents the instantaneous achievable rate during the active relaying phase over bandwidth B, and is used as the weight in our maximisation problem to prioritise links with higher instantaneous quality.

The optimisation problem can therefore be formulated as maximising the sum of rates from all active relay links:

$$\max_{s_{k_j,t_i}} \quad \sum_{k_j \in \mathcal{M}} \sum_{t_i \in \mathcal{T}} s_{k_j,t_i} \, r_{k_j,t_i} \tag{11a}$$

$$\text{s.t.} \quad \sum_{t_i \in \mathcal{T}} s_{k_j,t_i} \leq 1 \quad \forall \, k_j \in \mathcal{M}, \tag{11b}$$

$$\sum_{k_j \in \mathcal{M}} s_{k_j,t_i} \leq 1 \quad \forall \, t_i \in \mathcal{T}, \tag{11c}$$

$$s_{k_j,t_i} \in \{0, 1\} \quad \forall \, k_j \in \mathcal{M}, \, t_i \in \mathcal{T}. \tag{11d}$$

Constraints (11b) and (11c) ensure that each relay VUE serves at most one destination, and each destination VUE is served by at most one relay, respectively, thereby enforcing a one-to-one assignment. Constraint (11d) defines the decision variables as binary.

Several conditions must be met for a pair (k_j, t_i) to form a viable relay link, and these are used to construct the graph for the subsequent matching algorithm rather than being explicit constraints in the optimisation problem solved by KMRA itself:

1. HAP-to-Relay Link Quality: The candidate relay k_j must have a sufficiently strong link from the HAP. This is already ensured by the definition of the set $\mathcal{M}$, where for any $k_j \in \mathcal{M}$, $\tilde{\gamma}_{\text{HAP},k_j} \geq \chi_{\text{HAP-direct}}$. Here, $\chi_{\text{HAP-direct}}$ is the minimum SINR required for reliable decoding at the relay VUE, determined by the HAP's modulation and coding scheme.
2. Relay-to-Destination Link Quality: The SINR on the link from relay k_j to destination t_i must meet a minimum threshold: $\tilde{\gamma}_{k_j,t_i} \geq \gamma_{\text{V2V-th}}$. $\gamma_{\text{V2V-th}}$ is the minimum SINR for reliable decoding at the destination VUE, reflecting V2V link characteristics.
3. Relay Transmit Power: The transmit power P_{k_j,t_i} used by relay k_j (assumed to be a fixed operational value for a given link or relay, e.g., P_{k_j}) must not exceed its maximum capability $P_{k_j}^{\max}$.
4. Link Duration: The estimated link-contact duration C_{k_j,t_i} between relay k_j and destination t_i must exceed a minimum required service time T_{service}, as described in Section 3.3.
5. Interference Constraint: The estimated aggregated interference power at the destination VUE t_i from other relays, $I_{k_j,t_i}^{\text{relay}}(\tau)$, must not exceed a predefined interference threshold I_{th}, i.e.,
$$I_{k_j,t_i}^{\text{relay}}(\tau) \leq I_{\text{th}}.$$

This formulation, when considering all VUEs, leads to a combinatorial assignment problem. An exhaustive search for the optimal s_{k_j,t_i} values is computationally intractable for networks of practical size. To overcome this, we reformulate the relay selection as a maximum-weighted bipartite matching problem. An edge (k_j, t_i) exists in the bipartite graph only if all the aforementioned viability conditions (2, 3, 4 and 5, with condition 1 pre-screening $\mathcal{M}$) are met. The weight of such an edge is the achievable rate r_{k_j,t_i} from Eq. 10. The Kuhn-Munkres Relay Assignment (KMRA) algorithm can then efficiently find the assignment that maximises the sum of weights (i.e., total data rate).

5 Relay Selection via Bipartite Matching

To efficiently solve the relay assignment problem formulated in Section 4, we model it as a maximum weighted bipartite matching problem.

5.1 Bipartite Graph Construction

We construct a weighted bipartite graph $G = (\mathcal{M} \cup \mathcal{T}, E)$, where $\mathcal{M}$ is the set of candidate relay VUEs and $\mathcal{T}$ is the set of destination VUEs. An edge $(k_j, t_i) \in E$ is created between a candidate relay $k_j \in \mathcal{M}$ and a destination VUE $t_i \in \mathcal{T}$ if and only if the pair (k_j, t_i) constitutes a viable relay link. As detailed in Section 4, this means the following conditions must all be met:

- The HAP-to-relay link SINR $\tilde{\gamma}_{\text{HAP},k_j}$ meets the threshold $\chi_{\text{HAP-direct}}$ (inherent in the definition of $k_j \in \mathcal{M}$).
- The relay-to-destination link SINR $\tilde{\gamma}_{k_j,t_i}$ meets the threshold $\gamma_{\text{V2V-th}}$.
- The relay's transmit power P_{k_j,t_i} is within its maximum limit $P_{k_j}^{\max}$.
- The estimated link-contact duration C_{k_j,t_i} exceeds the required service time T_{service}.
- Interference Constraint: The estimated interference at the destination VUE t_i from other relays, denoted $I_{k_j,t_i}^{\text{relay}}(\tau)$, must not exceed the predefined interference threshold I_{th}.

If a pair (k_j, t_i) satisfies all these conditions, an edge is included in E. The *weight* $w(k_j, t_i)$ of this edge is defined as the achievable data rate r_{k_j,t_i} for that link, calculated using (10). If a pair does not meet all viability conditions, no edge is formed, or equivalently, its weight can be considered $-\infty$ for a maximisation problem.

Our objective is to find a matching $M \subseteq E$ such that each destination VUE $t_i \in \mathcal{T}$ is matched to at most one relay $k_j \in \mathcal{M}$, and each $k_j \in \mathcal{M}$ is matched to at most one $t_i \in \mathcal{T}$ (enforcing one-to-one assignment as per constraints (11b) and (11c)), in a way that maximises the sum of weights of the selected edges:

$$\max_{M} \sum_{(k_j,t_i) \in M} w(k_j, t_i). \tag{12}$$

5.2 Kuhn-Munkres Relay Algorithm for Assignment (KMRA)

The Maximum Weighted Bipartite Matching problem, as defined above, can be solved optimally in polynomial time. We employ the Kuhn-Munkres algorithm, also known as the Hungarian algorithm [28], for this purpose. KMRA typically finds a perfect matching with the maximum weight in a complete bipartite graph where $|\mathcal{M}| = |\mathcal{T}|$. Adjustments are needed if the graph is not complete or if the sets have unequal sizes (e.g., by adding dummy nodes and zero-weight edges, or by seeking a maximum weight matching that may not cover all nodes if $|\mathcal{M}| < |\mathcal{T}|$).

The core steps of the algorithm, adapted for our maximisation problem, involve:

1. Graph Preparation: If $|\mathcal{M}| \neq |\mathcal{T}|$, augment the smaller set with dummy nodes to make the sets of equal size, $n = \max(|\mathcal{M}|, |\mathcal{T}|)$. Add edges between real nodes and dummy nodes with zero weight. For pairs (k_j, t_i) where no viable link exists (i.e., an edge is not in E as per viability conditions), assign a weight of 0 (or a negative number if all actual rates are positive and non-zero to distinguish from dummy matches). This makes the graph effectively complete for the algorithm.
2. Initial Feasible Labelling: Assign labels (potentials) $\ell(v)$ to each vertex $v \in (\mathcal{M} \cup \mathcal{T})$. For maximisation, a common initialisation is $\ell(k_j) = \max_{t_i \in \mathcal{T}}\{w(k_j, t_i)\}$ for all $k_j \in \mathcal{M}$, and $\ell(t_i) = 0$ for all $t_i \in \mathcal{T}$. This ensures $\ell(k_j) + \ell(t_i) \geq w(k_j, t_i)$ for all edges.

3. Equality Subgraph Construction: Form an equality subgraph $G_\ell = (\mathcal{M} \cup \mathcal{T}, E_\ell)$, where $E_\ell = \{(k_j, t_i) \in E \mid \ell(k_j) + \ell(t_i) = w(k_j, t_i)\}$.

Algorithm 1 Kuhn-Munkres Relay Assignment (KMRA)

Require: Sets $\mathcal{M}$ (candidate relays), $\mathcal{T}$ (destination VUEs), and system parameters for rate and viability checks.

Ensure: Optimal matching $M^* \subseteq \mathcal{M} \times \mathcal{T}$ that maximises the sum of achievable rates $\sum r_{k_j,t_i}$.

Graph Construction:
Initialise edge set $E \leftarrow \emptyset$
for each $k_j \in \mathcal{M}$ **do**
    for each $t_i \in \mathcal{T}$ **do**
        if $\mathrm{SINR}_{k_j,t_i} \geq \gamma_{\text{V2V-th}}$ and $P_{k_j} \leq P_{k_j}^{\max}$ and $C_{k_j,t_i} > T_{\text{service}}$ **then**
            Compute weight $w(k_j, t_i) \leftarrow r_{k_j,t_i}$ using Eq. (10)
            Add edge (k_j, t_i) with weight $w(k_j, t_i)$ to E
        end if
    end for
end for
Construct bipartite graph $G = (\mathcal{M} \cup \mathcal{T}, E)$
if $|\mathcal{M}| \neq |\mathcal{T}|$ **then**
    Add dummy nodes/edges to square the bipartite graph
end if
Initialise vertex labels $\ell(k_j)$ and $\ell(t_i)$
while no full matching found in equality subgraph G_ℓ **do**
    Define $E_\ell \leftarrow \{(k_j, t_i) \mid \ell(k_j) + \ell(t_i) = w(k_j, t_i)\}$
    Find maximum matching M' in G_ℓ
    if M' covers $\min(|\mathcal{M}|, |\mathcal{T}|)$ nodes **then**
        break ▷ Optimal matching found
    else
        Update labels $\ell(v)$ via augmenting path to expand E_ℓ
    end if
end while
$M^* \leftarrow M'$
return M^*

4. Iterative Matching Improvement: Find a maximum cardinality matching M' in G_ℓ.
    - If M' covers all nodes in the smaller of the original $\mathcal{M}$ and $\mathcal{T}$ sets (or all n nodes if dummy nodes were added to achieve a perfect matching of size n), then M' is a maximum weight matching for G.
    - Otherwise, if M' is not yet of the desired size, update the labels $\ell(v)$ to include more edges in G_ℓ while maintaining feasibility. Repeat by finding a new maximum cardinality matching in the updated G_ℓ until a matching of the desired size is found.

The resulting matching provides the optimal set of relay-destination pairings that maximises the sum rate according to our problem formulation. We refer to this application of the algorithm as Kuhn-Munkres Relay Assignment (KMRA).

The KMRA algorithm is as outlined in Algorithm 1 and runs in $O(n^3)$ time, where $n = \max(|\mathcal{M}|, |\mathcal{T}|)$. In practical HAP-CVN scenarios, n is the number of candidate relays and out-of-coverage VUEs within a specific service area or cluster being managed. While a HAP covers a vast area, the number of active UEs involved in this specific relay selection problem at any given snapshot might be constrained, making this polynomial-time complexity significantly more tractable than an exponential-time exhaustive search.

5.3 Computational Complexity and Feasibility Analysis

A critical aspect of any relay selection algorithm is its computational efficiency and feasibility for real-time deployment. The KMRA algorithm has a well-established polynomial-time complexity of $O(n^3)$, where $n = \max(|\mathcal{M}|, |\mathcal{T}|)$ [24]. We can compare this analytically with our benchmark schemes:

- **RRA:** The Random Relay Assignment has a low complexity of approximately $O(|\mathcal{T}|)$, as it performs a single random selection for each destination VUE.
- **TBRA:** The Threshold-Based Relay Assignment, being a greedy algorithm, must iterate through each of the $|\mathcal{T}|$ destinations and, in the worst case, evaluate each of the $|\mathcal{M}|$ candidate relays for each one. This results in a complexity of approximately $O(|\mathcal{M}| \cdot |\mathcal{T}|)$, or $O(n^2)$.
- **LLRS:** The Link Lifetime Relay Selection heuristic, which also typically relies on a greedy assignment based on its stability metric, operates at a similar complexity of approximately $O(n^2)$.

This analysis reveals a clear trade-off: KMRA's $O(n^3)$ complexity is higher than the $O(n^2)$ complexity of the greedy heuristics. However, this increased computational investment yields a globally optimal assignment for the sum-rate objective, as demonstrated by our performance results in Section 8.

Regarding real-time feasibility, for a moderately sized problem within a managed HAP cluster, e.g., $n = 50$ (representing 25 relays and 25 destinations), the number of operations for an $O(n^3)$ algorithm is on the order of $50^3 = 125,000$. Modern processors can execute such a task in a timeframe on the order of milliseconds. Given that channel coherence times and vehicular topology update intervals are typically in the range of 100 ms to several seconds, this computational latency is well within acceptable limits for practical deployment. Therefore, while KMRA is more computationally demanding than the benchmark heuristics, its polynomial-time nature makes it a feasible and highly effective solution for the targeted HAP-CVN scenarios.

6 Simulation Environment and Setup

In this section, we detail the key aspects of our simulation setup including the HAP and traffic configuration. Our aim is to create a realistic traffic model for the HAP-assisted cooperative vehicular networking (HAP-CVN) scenario.

6.1 HAP and Traffic Configuration

A single High-Altitude Platform (HAP) is assumed to be quasi-stationary at an altitude of 20 km, providing multi-beam coverage over the entire simulation region. The HAP's multi-beam capability allows for spatial separation and potential frequency reuse across different geographical sectors; inter-beam interference (from other beams of the same HAP serving other VUEs) is considered a primary component of the external interference $I^{\text{ext}}_{\text{HAP},k_j}(\tau)$ in the SINR model for HAP-to-VUE links.

For the traffic configuration, we define Traffic Assignment Zones (TAZs) across the rural road network, representing key areas where trips originate and terminate. Using O/D matrices, we specify the volume of VUEs traveling between each TAZ pair within defined time periods. These matrices are then processed using SUMO's od2trips tool [29]. This tool converts the aggregate O/D demand into individual VUE trips, mapping TAZs to specific start and end edges on the network. A crucial feature for realism is that od2trips assigns a random departure time to each VUE within its specified O/D interval, inherently creating a non-uniform, Poisson-like arrival pattern across the network. Vehicles follow their computed routes towards their destination TAZs, departing the simulation upon reaching their destination or at the network boundaries. The overall traffic volumes in the O/D matrices are calibrated to ensure that this dynamic flow sustains an average of approximately N=100 VUEs throughout the simulation.

Additionally, to govern VUE movement, vehicle speeds are configured in SUMO to reflect typical rural driving. We set appropriate speed limits on the road network. For individual VUEs, speeds are varied using SUMO's speedFactor attribute within their vehicle type definitions. We employ a truncated normal distribution via speedFactor=normc(1.0, 0.2, 0.5, 1.5). This means VUEs generally target the legal speed limit (their SpeedFactor mean is 1.0) with a 20% standard deviation, but their speeds are realistically bounded (between 50% and 150% of the limit). This setup achieves an overall average speed near 70 km/h with a standard deviation of around 15 km/h.

At each simulation snapshot (which represents a specific point in time for evaluating channel conditions and VUE positions):

- VUEs are classified based on their instantaneous connectivity to the HAP. Using the HAP configuration, channel models (detailed in Table 1), and the SINR threshold $\chi_{\text{HAP-direct}}$, we observe that, on average, approximately 25% of the total VUEs have $\tilde{\gamma}_{\text{HAP},t_i} < \chi_{\text{HAP-direct}}$. These VUEs form the set of destination VUEs $\mathcal{T}$ requiring relay assistance.
- The remaining VUEs (approximately 75% of the total) satisfy $\tilde{\gamma}_{\text{HAP},k_j} \geq \chi_{\text{HAP-direct}}$ and thus constitute the set $\mathcal{H}$. As per our system model (Section 3), where $\mathcal{M} = \{k_j \in \mathcal{H} : \tilde{\gamma}_{\text{HAP},k_j} \geq \chi_{\text{HAP-direct}}\}$, all VUEs in $\mathcal{H}$ are considered candidate relays, meaning $\mathcal{M} = \mathcal{H}$ in this setup.

This configuration, resulting in approximately 25 destination VUEs and 75 candidate relays on average per snapshot, establishes a scenario with a clear need for relaying and a substantial pool of potential relays, making the selection problem non-trivial. The exact counts in $\mathcal{T}$ and $\mathcal{M}$ fluctuate in each snapshot due to VUE mobility and dynamic channel conditions. The wireless channel models, including path loss, shadowing, and small-scale fading for both HAP-to-VUE and V2V links, are based on the descriptions in Section 3, with specific parameters detailed in Table 1. Each data point presented in the results section is obtained by averaging over 500 independent simulation snapshots to ensure statistical reliability and to capture the effects of varying VUE positions and channel realisations.

Table 1 summarises the key system and channel model parameters used in our simulations.

Table 1 Key Simulation Parameters

Parameter	**Value / Assumption**
System and Scenario Parameters	
Simulation Area Size	30 km × 30 km (rural)
Mobility Generator	SUMO [29]
Total VUEs	100
Average VUE Speed	70 km/h
HAP Altitude	20 km
HAP Coverage Type	Multibeam
Communication Parameters	
Carrier Frequency	5.9 GHz
System Bandwidth	10 MHz
Noise Power	-101 dBm
HAP Tx Power	23 dBm
Max Relay VUE Tx Power	20 dBm
Channel Model Parameters	
Path-Loss at d_0 (V2V)	47.86 dB
Path-Loss Exponent (V2V)	2.75
Shadowing Std. Dev. (V2V)	4 dB
Path-Loss Exponent (HAP-VUE)	2.1
Shadowing Std. Dev. (HAP-VUE)	3 dB
Performance Thresholds	
HAP-Relay SINR Threshold	5 dB
Relay-Destination SINR Threshold	5 dB
Min. Service Time	100 ms
Max V2V Comm. Range	500 m
Min. Rate Threshold	1 bps/Hz
References	
V2V path loss based on 3GPP TR 37.885 [30]	
HAP-VUE path loss based on 3GPP TR 38.811 [31] ITU-R P.618 [32]	

In setting the SINR thresholds for link viability, we draw upon the principles of link adaptation in NR-V2X, where transmissions adjust their Modulation and Coding Scheme (MCS) based on channel quality to ensure reliability [33]. Channel Quality Indicator (CQI) feedback from vehicle UEs enables the selection of an appropriate MCS. For challenging link conditions, such as those experienced at the cell edge or by out-of-coverage VUEs, robust schemes like QPSK with a low code rate (e.g., effective rate around 0.3) are employed to enhance link resilience [33]. Following the methodology in [34] and considering the SINR requirements for various MCS levels, such as those detailed in 3GPP TR 38.901 (Tables 5.2.2.1-3 for baseline performance) [35], we establish the minimum SINR thresholds for a link to be considered viable for relaying.

For this study, both the HAP-to-relay SINR threshold ($\chi_{\text{HAP-direct}}$) and the relay-to-destination V2V SINR threshold ($\gamma_{\text{V2V-th}}$) are set to 5 dB. This value is selected as a conservative lower bound representative of the SINR required to reliably decode transmissions using a robust, low-rate MCS suitable for challenging vehicular links. It is important to note that these viability thresholds are distinct from, and more stringent than, the SINR threshold used to define an outage event. As specified in Table 1, an outage occurs if the effective end-to-end SINR for a relayed connection results in an achievable rate below $R_{\text{th}} = 1$ bps/Hz (which corresponds to an effective SINR of 0 dB). Thus, our selection criteria ensure that individual hops of a relay link exhibit a higher quality than the minimum required for basic service continuity.

7 Performance Metrics and Benchmark Schemes

7.1 Benchmark Schemes

In this subsection, we introduce three benchmark relay selection schemes against which the performance of the proposed KMRA approach is evaluated. We compare KMRA against three benchmark schemes representing greedy (TBRA), mobility-aware (LLRS), and random (RRA) strategies. These fundamental approaches were chosen to clearly establish the performance gains achievable through our optimal matching framework relative to common heuristics.

1. Threshold-Based Relay Assignment (TBRA): This is a greedy approach. For each unassigned destination VUE $t_i \in \mathcal{T}$, TBRA identifies all candidate relays $k_j \in \mathcal{M}$ that satisfy the HAP-to-relay SINR threshold (i.e., $\tilde{\gamma}_{\text{HAP},k_j} \geq \chi_{\text{HAP-direct}}$) and the V2V SINR threshold to t_i (i.e., $\tilde{\gamma}_{k_j,t_i} \geq \gamma_{\text{V2V-th}}$). Among these valid relays, the one providing the highest $\tilde{\gamma}_{k_j,t_i}$ to t_i is selected. If multiple destination VUEs can be served by the same "best" relay, a random tie-break occurs for the VUE assignment, and the relay is then considered assigned. The process repeats until all t_i are assigned or no more valid relays can be found.
2. Link Lifetime Relay Selection (LLRS): This mobility-aware heuristic, inspired by approaches like [36, 37], prioritises link stability. It attempts to form one-to-one assignments by selecting relay-destination pairs (k_j, t_i) that maximise a weighted metric combining current SINR ($\tilde{\gamma}_{k_j,t_i}$) and predicted link duration (C_{k_j,t_i}).
3. Random Relay Assignment (RRA): For each destination VUE $t_i \in \mathcal{T}$, a relay $k_j \in \mathcal{M}$ is chosen uniformly at random from the set of available (unassigned) candidate relays that are within the effective communication range R_{eff} of t_i. This serves as a basic performance baseline.

7.2 Performance Metrics and Benchmark Schemes

To evaluate the efficiency of our proposed relay selection scheme (KMRA), we utilise the following key performance metrics:

7.2.1 Outage Probability

The outage probability (P_{out}) for a destination VUE t_i served by a relay k_j is defined as the probability that its achieved data rate r_{k_j,t_i} falls below a predefined minimum acceptable rate threshold R_{th}. This is equivalent to the probability that its effective end-to-end SINR, $\tilde{\gamma}_{\text{eff},k_j,t_i}$, falls below the corresponding SINR threshold $\gamma_{\text{min\_th}} = 2^{R_{\text{th}}/B} - 1$.

Thus:

$$P_{\text{out}}(t_i) = \Pr\left(\tilde{\gamma}_{\text{eff},k_j,t_i} < \gamma_{\text{min\_th}}\right), \tag{13}$$

where $\tilde{\gamma}_{\text{eff},k_j,t_i} = \min\{\tilde{\gamma}_{\text{HAP},k_j}, \tilde{\gamma}_{k_j,t_i}\}$ as defined in (1). The overall system outage probability can be considered as the average outage experienced by all actively relayed destination VUEs.

7.2.2 Achievable User Throughput

The achievable user throughput for a destination VUE t_i being served by relay k_j is its data rate r_{k_j,t_i} as calculated by (2). For VUEs with direct HAP links, it is r_{direct,t_i}. To assess the distribution of service quality, we will examine the Cumulative Distribution Function (CDF) of the per-user throughput for all destination VUEs in $\mathcal{T}$. This allows us to understand the variability of throughput across users and the benefits provided by the relaying scheme, particularly for cell-edge or out-of-coverage VUEs.

8 Simulation Results and Discussion

Figure 2 presents the average Transmission Success Probability (P_{succ}) as a function of the V2V link distance between relay and destination pairs, comparing our proposed KMRA against the TBRA, LLRS, and RRA benchmarks.

As expected, all schemes show degradation in P_{succ} with increasing distance, primarily due to the higher path loss reducing the V2V SINR ($\tilde{\gamma}_{k_j,t_i}$) and thus increasing the chance of falling below the $\gamma_{\text{V2V-th}}$ threshold. KMRA consistently achieves the highest P_{succ}, significantly outperforming RRA and demonstrating a clear advantage over both TBRA and LLRS. This supports our central thesis that KMRA finds a globally optimal assignment to maximise the sum-rate, using r_{k_j,t_i} as edge weights. Since r_{k_j,t_i} is a direct function of the end-to-end SINR, and P_{succ} is heavily dependent on meeting SINR thresholds, maximising sum-rate strongly correlates with maximising overall success probability. Unlike TBRA, KMRA considers all pairings simultaneously, avoiding sub-optimal local assignments. Similarly, while LLRS prioritises link duration, it can underperform KMRA if it selects a long-duration link that has only marginal SINR quality. KMRA, by focusing on rate (and thus SINR), inherently selects high-quality links, which, provided they meet the minimum T_{service} constraint, have a higher chance of successful transmission. TBRA makes greedy decisions based on local thresholds, which can lead to situations where some VUEs get poor pairings or no pairings, even if a better global assignment exists, thus lowering the average P_{succ}. This confirms that an assignment strategy based on maximising a quality-of-service metric (rate) within a globally optimal framework (bipartite matching) leads to superior transmission success compared to baseline and heuristic approaches.

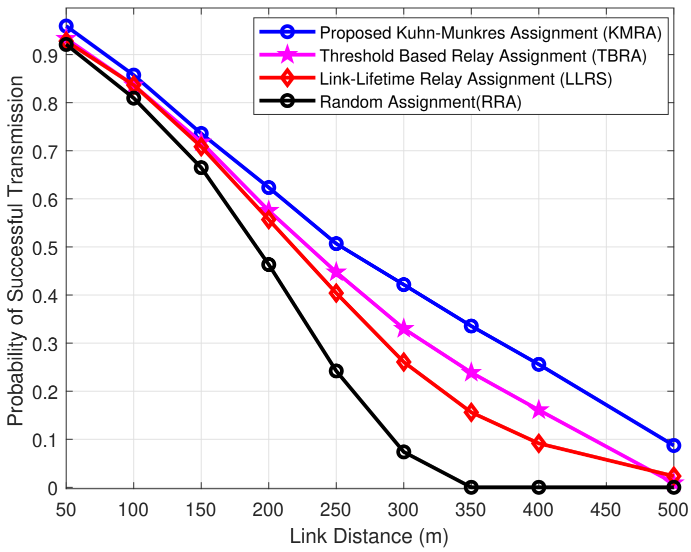

Fig. 2 Average Transmission Success Probability (P_{succ}) vs. V2V link distance (meters) for different relay selection algorithms.

8.1 Outage Probability Performance

We now assess the Outage Probability (P_{out}), defined according to (13) as the likelihood of falling below the minimum SINR threshold $\gamma_{\text{min\_th}}$.

Fig. 3 presents P_{out} as a function of the average V2V SINR. The y-axis is plotted on a logarithmic scale to better visualise performance improvements at low outage levels. As expected, all schemes exhibit a decreasing outage probability as the link quality improves. Specifically, KMRA demonstrates significantly superior performance, consistently achieving the lowest outage probability across the entire SINR range. Notably, at an average V2V SINR of 15 dB, KMRA reaches an outage probability close to 10^{-3}, which is orders of magnitude better than the baseline RRA and markedly better than both TBRA and LLRS. This advantage stems directly from KMRA's objective: by maximising the sum rate based on the minimum SINR of the two hops, it inherently prioritises assignments that offer high, balanced link quality, thus directly minimising the probability of either hop falling below the outage threshold. In contrast, RRA ignores link quality while TBRA and LLRS employ heuristic or greedy criteria that do not guarantee globally strong link assignments, leading to higher average outage rates.

8.2 Impact of Interference Threshold on Success Probability

Fig. 4 explores the system's robustness by illustrating the average Transmission Success Probability (P_{succ}) as a function of the Interference Threshold (I_{th}). As defined in Section 3.2.4 and considered in Section 4, I_{th} acts as a hard constraint: a V2V link (k_j, t_i) is deemed non-viable if the estimated aggregated interference $I^{\text{relay}}_{k_j,t_i}(\tau)$ at t_i

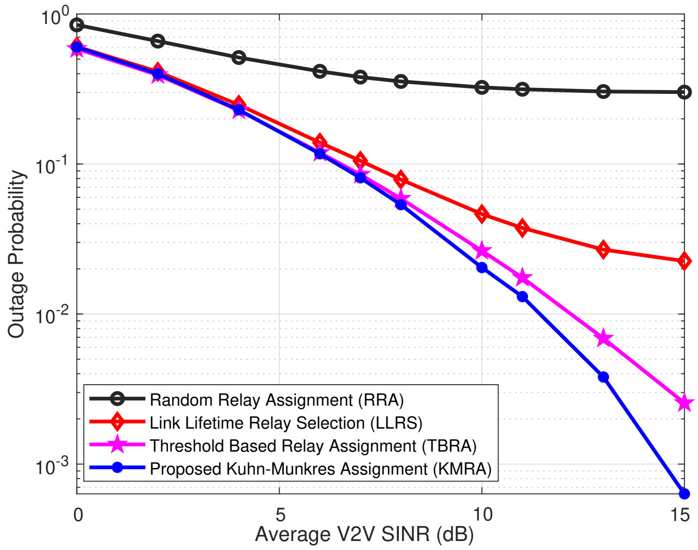

Fig. 3 Outage probability (P_{out}) vs. Average V2V Link SINR (dB) for different relay selection algorithms.

exceeds I_{th}. This means I_{th} directly limits the pool of available links for the assignment algorithms.

The figure demonstrates that for KMRA, TBRA, and LLRS, when the interference threshold (I_{th}) is very low (e.g., -105 dBm), the success probability is extremely low, approaching zero for KMRA. This is a direct consequence of the stringent interference constraint: very few, if any, potential relay links satisfy $I^{\text{relay}}_{k_j,t_i}(\tau) \leq I_{\text{th}}$, making it nearly impossible to find valid assignments. As I_{th} is increased, allowing V2V links to tolerate higher levels of interference, more links become viable. Consequently, the success probability rises sharply for KMRA, which quickly approaches optimal performance. This showcases KMRA's efficiency: once viable links – those satisfying the I_{th} constraint – become available, it optimally selects the best ones to maximise overall system performance. TBRA and LLRS also show an improvement in P_{succ} as I_{th} increases, but they generally lag behind KMRA, particularly in the transition region. This is attributed to their heuristic or greedy nature, which may not fully exploit the expanded set of viable links as effectively as KMRA's global optimisation.

The RRA scheme exhibits a high and relatively flat P_{succ} across the range of I_{th} values. This indicates that, for this baseline comparison, RRA is implemented without strictly enforcing the I_{th} constraint during its random selection process. It thus serves as a reference showing potential success rates if this specific interference management constraint were ignored. Overall, Fig. 4 highlights KMRA's ability to strictly adhere to defined interference constraints and to effectively leverage the available link opportunities as these constraints are relaxed, rapidly achieving high success rates.

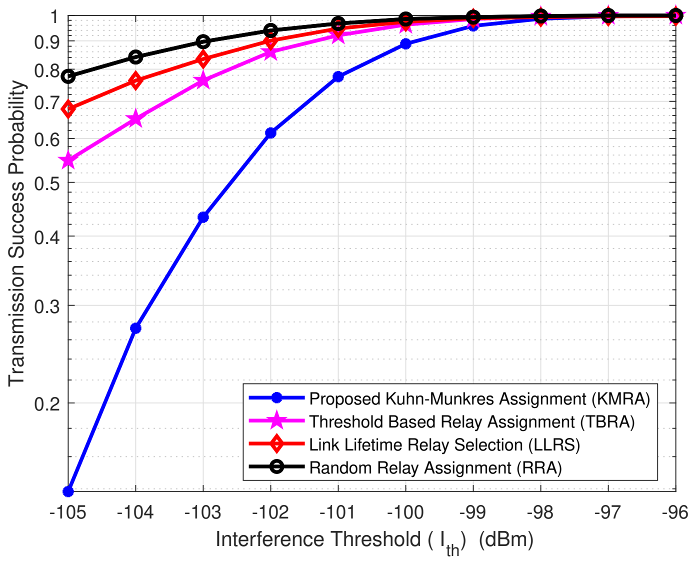

Fig. 4 Average Transmission Success Probability (P_{succ}) vs. Interference Threshold (I_{th} in dBm) for different relay selection algorithms.

8.3 Achievable User Throughput

Finally, we evaluate the achievable user throughput distribution for the destination VUEs in $\mathcal{T}$ when served by the different relay selection schemes. Fig. 5 presents the Cumulative Distribution Function (CDF) of the per-user throughput (in bps/Hz), generated from data aggregated over 500 independent simulation snapshots.

The figure compares the throughput achieved by VUEs in $\mathcal{T}$ when assisted by relays selected using KMRA, TBRA, LLRS, and RRA. The curves illustrate the significant performance differences among the relay selection strategies. Our proposed KMRA algorithm yields the most favourable throughput distribution, consistently providing higher data rates to the destination VUEs across the entire user population compared to the benchmark schemes. For instance, at the median (50th percentile), KMRA achieves a user throughput of approximately 1.3 bps/Hz, whereas TBRA, LLRS, and RRA provide notably lower median throughputs around 0.8-0.9 bps/Hz. This superior performance of KMRA is evident across all percentiles, indicating that it not only improves the situation for the worst-off users but also pushes a larger fraction of users to higher throughput levels. This demonstrates KMRA's effectiveness in selecting high-quality relay links that translate into substantial data rates for the assisted VUEs. While all presented schemes employ relaying, KMRA's optimal, rate-driven assignment ensures that the allocated relay resources are utilised most efficiently to enhance user capacity.

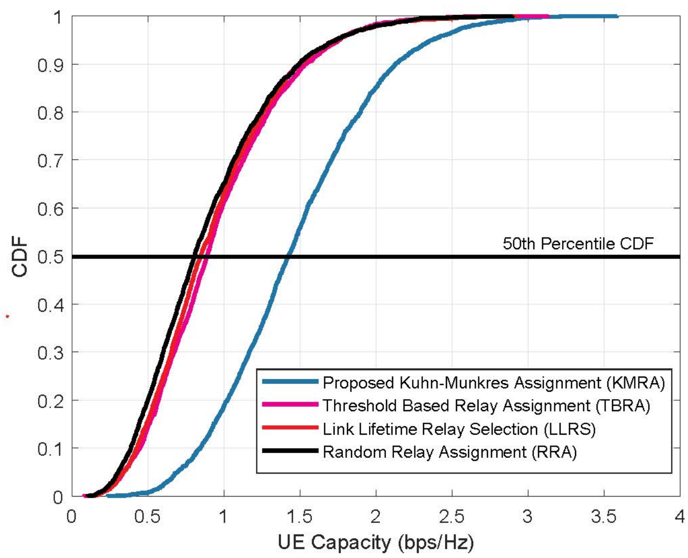

Fig. 5 CDF of achievable user throughput (bps/Hz) for destination VUEs in $\mathcal{T}$ under different relay selection schemes.

8.4 Summary of Simulation Findings

In summary, the presented simulation results consistently validate the advantage of our proposed KMRA algorithm over the benchmark relay selection schemes (TBRA, LLRS, and RRA) across several key performance indicators. Specifically, KMRA demonstrates:

- Higher transmission success probabilities under varying link durations and V2V distances.
- Significantly lower outage probabilities, especially as link quality (e.g., average V2V SINR) improves.
- Enhanced robustness against interference, effectively utilising links even under stringent interference thresholds.
- Markedly improved throughput distributions for VUEs at the network edge or in out-of-coverage regions, transforming unusable direct links into effective relayed connections.

These performance gains are attributed to KMRA's optimal, rate-driven assignment based on a global view of potential relay-destination pairings, ensuring efficient resource utilisation and robust link selection.

9 Conclusion and Future Work

This paper has presented a relay VUE selection scheme for High-altitude Platform (HAP)-assisted cooperative vehicular networks, where VUEs with strong direct HAP links are optimally assigned as relays for out-of-coverage or cell-edge VUEs. We formulated this critical relay-to-destination VUE assignment as a maximum-weighted bipartite matching problem, with edge weights intelligently defined by the achievable data rates of the two-hop links.

The efficacy of the proposed KMRA algorithm was validated through comprehensive simulations. Results demonstrate KMRA's advantage over common benchmark approaches—namely threshold-based, link-lifetime-based, and random assignments—particularly in its ability to significantly reduce outage probability and enhance data rates for VUEs experiencing poor direct channel conditions. This highlights KMRA's potential for improving overall network service quality and user experience in challenging HAP-CVN environments.

9.1 Future Work

Building upon this work, our future research directions include:

- **Dynamic Resource Allocation:** To investigate the joint optimisation of relay selection with dynamic frequency/time-slot allocation strategies to further mitigate inter-relay interference and enhance overall spectral efficiency.
- **Joint Scheduling and Multi-Hop Routing:** To extend the framework by combining relay selection with sophisticated scheduling and multi-hop routing algorithms, particularly for large-scale vehicular scenarios, and to support VUEs beyond the immediate single-relay coverage range.
- **Intelligent Data-Driven Approaches:** To explore the application of machine learning and multi-agent reinforcement learning techniques for developing low-overhead, adaptive, and distributed relay selection algorithms that can learn from real-time channel dynamics and complex vehicular mobility patterns.
- **Energy-Efficient and Secure Relay Selection:** To extend our model to a multi-objective optimisation framework that jointly considers sum-rate, the energy consumption of relay VUEs, and link security. This will involve selecting relays based on trust metrics or using physical layer security techniques to protect the two-hop links.
- **Hybrid Rate-Stability Optimisation:** To develop a more sophisticated weighting function for the bipartite graph that jointly considers both the achievable rate and a link stability metric derived from predicted link duration (C_{k_j,t_i}). This would create a multi-objective assignment problem to better balance throughput with the robustness required in high-mobility scenarios.
- **QoS-Aware Relay Selection:** To introduce Quality of Service (QoS) differentiation by modifying the edge weights $w(k_j, t_i)$ in the bipartite graph. The weights could be adjusted based on the traffic priority of the destination VUE (e.g., higher weights for emergency messages than for infotainment data), thereby prioritising critical communications in the assignment process.

In conclusion, the proposed KMRA algorithm offers a computationally efficient and demonstrably effective solution to the relay VUE assignment problem. We acknowledge certain simplifying assumptions in our current framework such as the fixed bandwidth allocation and predefined SINR thresholds for link viability. In dynamic 6G NR-V2X environments, these parameters could be adaptively controlled. Furthermore, our current work treats all destination VUEs with uniform importance. Extending the KMRA framework to incorporate these real-world dynamics represents important future work.

Author Contributions

K.P., D.G., and T.C. contributed to conceptualisation, investigation, and methodology. K.P. contributed to writing the original draft. K.P., D.G., T.C., M.A., and S.N. contributed to writing—review, visualisation, and editing. All authors reviewed the results and approved the final version of the manuscript.

Funding

This research received no external funding.

Declarations

Competing Interests

The authors declare no competing interests.

References

[1] Ahmed, E., Gharavi, H.: Cooperative vehicular networking: A survey. IEEE Transactions on Intelligent Transportation Systems **19**(3), 996–1014 (2018)

[2] Hasan, M.K., Jahan, N., Nazri, M.Z.A., Islam, S., Khan, M.A., Alzahrani, A.I., Alalwan, N., Nam, Y.: Federated learning for computational offloading and resource management of vehicular edge computing in 6g-v2x network. IEEE Transactions on Consumer Electronics **70**(1), 3827–3847 (2024)

[3] Arum, S.C., Grace, D., Mitchell, P.D.: Extending coverage and capacity from high altitude platforms with a two-tier cellular architecture. IEEE Transactions on Mobile Computing **23**(2), 1942–1953 (2023)

[4] Kurt, G.K., Khoshkholgh, M.G., Alfattani, S., Ibrahim, A., Darwish, T.S., Alam, M.S., Yanikomeroglu, H., Yongacoglu, A.: A vision and framework for the high altitude platform station (haps) networks of the future. IEEE Communications Surveys & Tutorials **23**(2), 729–779 (2021)

[5] Li, Y., Wu, R., Gan, L., He, P.: Development of an effective relay communication technique for multi-uav wireless network. IEEE Access (2024)

[6] Feteiha, M.F., Ahmed, M.H.: Multihop best-relay selection for vehicular communication over highways traffic. IEEE Transactions on Vehicular Technology **67**(10), 9845–9855 (2018)

[7] Li, J., Niu, Y., Wu, H., Ai, B., He, R., Wang, N., Chen, S.: Joint optimization of relay selection and transmission scheduling for uav-aided mmwave vehicular networks. IEEE Transactions on Vehicular Technology **72**(5), 6322–6334 (2023)

[8] Akdemir, B., Karabulut, M.A., İlhan, H.: A comparative study of deep leaning-aided relay selection protocols in cooperative communication. Wireless Personal Communications **136**(1), 365–384 (2024)

[9] Pan, B., Wu, H.: Success probability analysis of cooperative c-v2x communications. IEEE Transactions on Intelligent Transportation Systems **23**(7), 7170–7183 (2021)

[10] Popoola, K., Grace, D., Clarke, T.: Capacity and coverage analysis of high altitude platform (hap) antenna arrays for rural vehicular broadband services. In: 2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring), pp. 1–5 (2020). IEEE

[11] Nauman, A., Maashi, M., Alkahtani, H.K., Al-Wesabi, F.N., Aljehane, N.O., Assiri, M., Ibrahim, S.S., Khan, W.U.: Efficient resource allocation and user association in noma-enabled vehicular-aided hetnets with high altitude platforms. Computer Communications **216**, 374–386 (2024)

[12] Ren, Q., Abbasi, O., Kurt, G.K., Yanikomeroglu, H., Chen, J.: Handoff-aware distributed computing in high altitude platform station (haps)–assisted vehicular networks. IEEE Transactions on Wireless Communications **22**(12), 8814–8827 (2023)

[13] Muhammad Umer, M., Hong, J., Muhammad, O., Awwad, F.A., Ismail, E.A.: Energy-efficient and resilient secure routing in energy harvesting wireless sensor networks with transceiver noises: Ecosecnet design and analysis. Journal of Sensors **2024**(1), 3570302 (2024)

[14] Umer, M.M., Jiang, H., Zhang, Q., Manlu, L., Muhammad, O.: Time-slot based architecture for power beam-assisted relay techniques in cr-wsns with transceiver hardware inadequacies. Bulletin of the Polish Academy of Sciences. Technical Sciences **71**(5) (2023)

[15] Umer, M.M., Venter, H., Muhammad, O., Shafique, T., Awwad, F.A., Ismail, E.A.: Cognitive strategies for uav trajectory optimization: Ensuring safety and energy efficiency in real-world scenarios. Ain Shams Engineering Journal **16**(3), 103301 (2025)

[16] Dong, Y., Zhang, F., Joe, I., Lin, H., Jiao, W., Zhang, Y.: Learning for multiple-relay selection in a vehicular delay tolerant network. IEEE Access **8**, 175602–175611 (2020)

[17] Singh, A.K., Pamula, R., Jain, P.K., Srivastava, G.: An efficient vehicular-relay

selection scheme for vehicular communication. Soft Computing **27**(6), 3443–3459 (2023)

[18] Alam, M.S., Mark, J.W., Shen, X.S.: Relay selection and resource allocation for multi-user cooperative ofdma networks. IEEE Transactions on Wireless Communications **12**(5), 2193–2205 (2013)

[19] Xu, C., Xiong, Z., Han, Z., Zhao, G., Yu, S.: Link reliability-based adaptive routing for multilevel vehicular networks. IEEE Transactions on Vehicular Technology **69**(10), 11771–11785 (2020)

[20] Pan, B., Wu, H.: Modeling and analysis of multi-relay cooperative communications in c-v2x networks. IEEE Transactions on Intelligent Transportation Systems **23**(9), 16371–16385 (2022)

[21] Ma, M., Wang, C., Li, Z., Min, G.: A multi-objective optimization-based transmission design for v2v communication networks. IEEE Transactions on Vehicular Technology **72**(10), 13081–13093 (2023)

[22] Fan, X., Liu, D., Fu, B., Wen, S.: Optimal relay selection for uav-assisted v2v communications. Wireless Networks **27**(5), 3233–3249 (2021)

[23] Liu, R., Guo, K., An, K., Huang, Y., Zhou, F., Zhu, S.: Resource allocation for cognitive satellite-hap-terrestrial networks with non-orthogonal multiple access. IEEE Transactions on Vehicular Technology **72**(7), 9659–9663 (2023)

[24] Sharma, S., Shi, Y., Liu, J., Hou, Y.T., Kompella, S., Midkiff, S.F.: Network coding in cooperative communications: Friend or foe? IEEE Transactions on Mobile Computing **11**(7), 1073–1085 (2012)

[25] Wei, Z., Wang, L., Gao, Z., Wu, H., Zhang, N., Han, K., Feng, Z.: Spectrum sharing between high altitude platform network and terrestrial network: Modeling and performance analysis. IEEE Transactions on Communications **71**(6), 3736–3751 (2023)

[26] Jang, G., Choi, J.P.: Haps altitude optimization for downlink communications: Alleviating the effect of channel elevation angles. IEEE Transactions on Aerospace and Electronic Systems (2025)

[27] Yahia, O.B., Erdogan, E., Kurt, G.K., Altunbas, I., Yanikomeroglu, H.: Haps selection for hybrid rf/fso satellite networks. IEEE Transactions on Aerospace and Electronic Systems **58**(4), 2855–2867 (2022)

[28] ZHU, L., GU, D., TANG, S., CHEN, X.: Algorithm for layered bipartite graph maximum matching problem. Computer and Modernization (06), 59 (2024)

[29] Lopez, P.A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P.,

Hilbrich, R., Lücken, L., Rummel, J., Wagner, P., Wießner, E.: Microscopic traffic simulation using sumo. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), pp. 2575–2582 (2018). Ieee

[30] 3rd Generation Partnership Project (3GPP): Study on evaluation methodology of new vehicle-to-everything (v2x) use cases for lte and nr. Technical Report 37.885, 3rd Generation Partnership Project (3GPP). Available online. https://www.3gpp.org/ftp/Specs/archive/37_series/37.885/37885-f00.zip

[31] 3rd Generation Partnership Project (3GPP): Study on new radio access technology: Channel model for frequencies from 6 to 100 ghz. Technical Report 38.811, 3rd Generation Partnership Project (3GPP). Available online. https://www.3gpp.org/ftp/Specs/archive/38_series/38.811/38811-e00.zip

[32] International Telecommunication Union – Radiocommunication Sector (ITU-R): Propagation data and prediction methods required for the design of earth-space telecommunication systems. Recommendation P.618-13, International Telecommunication Union. ITU-R Recommendation P.618-13. https://www.itu.int/rec/R-REC-P.618/en

[33] Garcia, M.H.C., Molina-Galan, A., Boban, M., Gozalvez, J., Coll-Perales, B., Şahin, T., Kousaridas, A.: A tutorial on 5g nr v2x communications. IEEE Communications Surveys & Tutorials **23**(3), 1972–2026 (2021)

[34] Bazzi, A., Cecchini, G., Menarini, M., Masini, B.M., Zanella, A.: Survey and perspectives of vehicular wi-fi versus sidelink cellular-v2x in the 5g era. Future Internet **11**(6), 122 (2019)

[35] European Telecommunications Standards Institute (ETSI): Study on channel model for frequencies from 0.5 to 100 ghz: 3gpp tr 38.901 version 15.0.0 release 15. Technical report, ETSI (2018). Online; accessed 2025-05-26. https://www.etsi.org/deliver/etsi_tr/138900_138999/138901/15.00.00_60/tr_138901v150000p.pdf

[36] Yelure, B., Sonavane, S.: Analysis of connectivity duration in vehicular network. Turkish Journal of Computer and Mathematics Education **12**(13), 2318–2325 (2021)

[37] Nugraha, T.A., Mach, P., Becvar, Z.: Mobility-aware mode selection, relay selection, and resource allocation for d2d communication. IEEE Transactions on Network Science and Engineering (2025)