

B-AWARE: Blockage Aware RSU Scheduling for 5G Enabled Autonomous Vehicles

MATTHEW SZETO, Arizona State University, USA

EDWARD ANDERT, Arizona State University, USA

AVIRAL SHRIVASTAVA, Arizona State University, USA

MARTIN REISSLEIN, Arizona State University, USA

CHUNG-WEI LIN, National Taiwan University, Taiwan

CHRIST RICHMOND, Duke University, USA

5G Millimeter Wave (mmWave) technology holds great promise for Connected Autonomous Vehicles (CAVs) due to its ability to achieve data rates in the Gbps range. However, mmWave suffers from a high beamforming overhead and requirement of line of sight (LOS) to maintain a strong connection. For Vehicle-to-Infrastructure (V2I) scenarios, where CAVs connect to roadside units (RSUs), these drawbacks become apparent. Because vehicles are dynamic, there is a large potential for link blockages. These blockages are detrimental to the connected applications running on the vehicle, such as cooperative perception and remote driver takeover. Existing RSU selection schemes base their decisions on signal strength and vehicle trajectory alone, which is not enough to prevent the blockage of links. Many modern CAVs motion planning algorithms routinely use other vehicle's near-future path plans, either by explicit communication among vehicles, or by prediction. In this paper, we make use of the knowledge of other vehicle's near future path plans to further improve the RSU association mechanism for CAVs. We solve the RSU association algorithm by converting it to a shortest path problem with the objective to maximize the total communication bandwidth. We evaluate our approach, titled B-AWARE, in simulation using Simulation of Urban Mobility (SUMO) and Digital twin for self-driving Intelligent Vehicles (DRIVE) on 12 highway and city street scenarios with varying traffic density and RSU placements. Simulations show B-AWARE results in a 1.05x improvement of the potential data rate in the average case and 1.28x in the best case vs. the state-of-the-art. But more impressively, B-AWARE reduces the time spent with no connection by 42% in the average case and 60% in the best case as compared to the state-of-the-art methods. This is a result of B-AWARE reducing nearly 100% of blockage occurrences.

CCS Concepts: • **Networks** → *Mobile networks*; **Network resources allocation**; • **Computer systems organization** → **Embedded and cyber-physical systems**; **Availability**.

Additional Key Words and Phrases: mmWave, V2I, CAV, RSU, selection, user association, autonomous vehicles, 5G, vehicular networks

ACM Reference Format:

Matthew Szeto, Edward Andert, Aviral Shrivastava, Martin Reisslein, Chung-Wei Lin, and Christ Richmond. 2023. B-AWARE: Blockage Aware RSU Scheduling for 5G Enabled Autonomous Vehicles. 1, 1 (March 2023), 23 pages. <https://doi.org/XXXXXXX.XXXXXXX>

Authors' addresses: [Matthew Szeto](mailto:maszeto@asu.edu), maszeto@asu.edu, Arizona State University, Tempe, Arizona, USA; [Edward Andert](mailto:edward.andert@asu.edu), Arizona State University, Tempe, Arizona, USA, edward.andert@asu.edu; [Aviral Shrivastava](mailto:aviral.shrivastava@asu.edu), Arizona State University, Tempe, Arizona, USA, aviral.shrivastava@asu.edu; [Martin Reisslein](mailto:reisslein@asu.edu), Arizona State University, Tempe, Arizona, USA, reisslein@asu.edu; [Chung-Wei Lin](mailto:chungwei@ntu.edu.tw), National Taiwan University, Taipei, Taiwan; [Christ Richmond](mailto:christ.richmond@duke.edu), Duke University, Durham, North Carolina, USA, christ.richmond@duke.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

XXXX-XXXX/2023/3-ART \$15.00

<https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Autonomous vehicles (AVs) are already present in today's society. Even so, many of these AVs still require an individual at the wheel to take over in the case of an emergency. To improve the safety of AVs and alleviate the need for physically present safety drivers, the paradigm of connected autonomous vehicles (CAVs) arose. CAVs have the ability to share sensor data with one another to perform sensor fusion or cooperative perception. For example, a vehicle which spots an obstruction in the road can transmit this information to the vehicles behind it, allowing other vehicles to safely avoid the hazard. Another feature of CAVs is that they allow for remote safety driver takeovers. This means that a single human safety driver can service 10s to even 100s of vehicles at the same time depending on the rate of incidents [29]. This idea of remote takeovers has been adopted among currently operating AV companies as a stop gap in order to reduce operator costs while they push their level 3 autonomy vehicles towards level 4/5 [31]. Regardless of the exact application, CAVs typically require a high percentage of network up-time and large amounts of network bandwidth. There are two main infrastructure schemes for connecting CAVs, Vehicle-to-Infrastructure (V2I) and Vehicle-to-Vehicle (V2V). In V2I, autonomous vehicles communicate with road side units (RSUs), which can then relay information to other autonomous vehicles on the road. In V2V, vehicles communicate directly with one another. 5G mmWave has been considered for use in V2X (V2I + V2V) communication due to its high bandwidth and low latency characteristics. 802.11ad, a 60 GHz mmWave based protocol, can achieve a theoretical maximum throughput of 6.75 Gbps [26]. This high transmission speed makes it a good choice for use with self-driving vehicles, as current commercially available connected autonomous vehicles can consume up to 750 Mb of data per second [43]. Future vehicles may operate by consuming even more data, and a larger fraction of that may be coming from communication channels.

MmWave connections are highly directional, unlike other technologies such as WI-FI. For a strong connection, the transmitter and receiver must maintain line-of-sight (LOS), as non line-of-sight (NLOS) connections will not be as strong [17]. In a dynamic environment LOS paths can be obstructed by other large vehicles [8]. Such blockages could potentially be minimized by strategically placing RSUs, but there is still a chance that the links will become blocked [39]. This will cause significant loss in signal strength, as blockages can cause a 20–30 dB loss at 60 GHz [7, 45]. If a CAV were to lose its LOS with a currently connected RSU, a handover would occur to find an RSU with a stronger connection. A *handover* is when a CAV moves from a serving RSU to a target RSU, and in a vehicular environment these handovers can occur frequently [36]. Handovers have a large overhead, due to the lengthy beamforming procedure needed to form mmWave connections [6]. Beamforming is the process of finding the beam parameters (e.g., direction, spread) that will provide the antenna on the vehicle the best connection for the longest time by some metric. The duration of the beamforming procedure can be reduced based on the implemented beamforming algorithm, but in general, it still takes a long time – often hundreds of milliseconds [28]. Thus blockages, or more precisely blockage-induced handovers, present one of the big challenges to the success of mmWave technology, and to reduce that, several RSU assignment algorithms are being proposed.

Figure 1 illustrates the problem of blockages, explains the existing RSU assignment approaches, and provides the key intuition behind our approach. Figure 1 (a) shows a driving scenario in which two vehicles, V_{ego} and V_B , are traveling north on a 4-lane road, and there are two nearby RSUs, RSU_A on the left side of the road, and RSU_B on the right side of the road, and the ego vehicle V_{ego} needs to connect to an RSU. The previous algorithm [1] selects the RSU that has the best RSS (Received Signal Strength) to connect with. In this scenario, without loss of generality, assume that the RSU_B is higher power than RSU_A . Then the approach [1] will chose RSU_B . However, this may not be the




Fig. 1. (a) Consider a scenario where vehicle V_{ego} wants to connect to an RSU, and there are two RSUs around, RSU_A and RSU_B . A previous popular approach, which we call Baseline[1], may connect with RSU_B if the signal quality is better from RSU_B . However, the LOS to RSU_B will soon be blocked by the vehicle V_B . (b) Another previous approach SMART[35] considers the vehicle's own trajectory and evaluates that for its own trajectory it will achieve higher total potential bandwidth by connecting to RSU_B , and therefore will connect to RSU_B . (c) Our Approach, B-AWARE, also considers the trajectory of other vehicles – V_B in this case, and determines that connection to RSU_B will be blocked and therefore it is beneficial to connect to RSU_A .

best choice, since the vehicle V_B may soon block the connection of V_{ego} to RSU_B , and result in an expensive handover. The more recently proposed RSU assignment algorithm, SMART[35], makes a decision for RSU assignment based on the total bandwidth it will achieve in the near future. Figure 1 (b) shows that given its own (V_{ego}) path plan, SMART will compare the total bandwidth achieved if it connects to RSU_A vs. RSU_B , and based on that it may also choose the RSU_B . The problem with both the previous algorithms is that they did not consider the impending blockage caused by the vehicle V_B . Figure 1 (c) summarizes our approach. In our approach V_{ego} considers the near-future plan of the vehicle V_B (which it can either receive directly from vehicle V_B through communication, or by prediction) to conclude that the RSU_B will be blocked, and therefore it may be better to connect to RSU_A .

In this paper, we define an optimization problem to maximize the potential datarate from received signal strength (RSS) based on RSU selections, given the near-future path plan of the ego and the other vehicles. Our solution is a framework, titled B-AWARE, which utilizes vehicle dimensions, trajectory, and map data to design an *RSU selection schedule* that avoids blockages and prevents unnecessary handovers. The RSU selection schedule tells the vehicle which RSU to associate with and when (the table below figure 1 (c)). The B-AWARE framework avoids unnecessary handovers and unexpected outages from blockages, thereby improving the maximum potential raw datarate for each vehicle, allowing safe and reliable CAV network applications, such as remote driver takeover and blind spot data sharing. The main contribution of this paper is as follows: We develop a blockage-aware RSU selection scheme, B-AWARE, which reduces unexpected outages and decreases unnecessary handovers whilst maintaining strong signal quality in a V2I environment by leveraging CAV path plans and sensing information.

We test our approach in both city and highway environments, varying both traffic density and RSU placements. The results of our comparison show that B-AWARE results in a overall 1.05x improvement in potential data transmission in the average case and 1.28x in the best case in both highway and city cases. Results also show our approach is resilient to changes in vehicle paths and can reduce the amount of time a vehicle is disconnected from the network by 42% in the average case and 60% in the best case. This is partly a result of B-AWARE reducing almost 100% of blockage occurrences in simulation. B-AWARE therefore increases the potential datarate and

exhibits a significantly larger resilience to complete connection loss caused by blockages over existing approaches in multiple scenarios.

This paper is organized as follows: Section 2 further explains the impact of handovers on network performance and related works, Section 3 describes our system model, Section 4 is the definition of our maximization problem, Section 5 explains our RSU scheduling strategy and schedule optimization algorithm, section 6 explains our experimental setup, Section 7 explains the results comparing our RSU scheduling algorithm with two other RSU selection approaches, and Section 8 summarizes our findings and next steps.

2 BACKGROUND AND RELATED WORK

This section gives some background as to why handovers and blockages are detrimental to mmWave links. Understanding this can help realize an RSU association policy to avoid such events.

2.1 mmWave Handover Overhead

In [18], researchers found the average time of a blockage induced hand off to be 2.749 seconds. Researchers in [11] reported similar results, finding the outage time of a blockage induced handover to be up to 2.8 seconds. This outage time is a combination of disassociation time with the current RSU and re-association time with a new RSU. These high handover times are largely due to a lack of intelligent handover mechanism in the current 802.11ad standard. Current consumer off-the-shelf hardware only chooses to find a new RSU when there has been poor link performance for an extended period of time. In 802.11ad, links are created by using a brute-force beam refinement procedure. During this procedure, a sector-level-sweep (SLS) is performed to locate the general direction of the receiver. Then, during the beam refinement procedure, different beams are tested until the one with the strongest strength is found. The actual association with a new RSU takes less than a second. [18] reports this to be an average of 240 milliseconds and [11] reports a slightly higher 440 milliseconds. A value of 700 milliseconds is used as the handover delay for experiments in [22]. This procedure is captured in Figure 5 summarised from [11] results. The occurrence of a blockage causes a significant drop in throughput and the link quality has not fully recovered until the searching and re-association with a new RSU is complete. The cost of handovers has also been modeled as a percentage of resources that needs to be used for signaling [27, 35]. Due to such high overhead, it is clear handovers have a direct impact on system throughput and too many handovers can be detrimental to system performance. For mmWave frequencies where LOS is imperative, if a CAV selects a base station which is then quickly blocked, there will be a drop in link performance and the CAV will need to find a new base station. Because even intermediary blockages can trigger handovers, avoiding blockages can help improve network performance. Research has shown even a single car moving down the street can lose LOS with an RSU every couple seconds [36].

2.2 Handover Time Reduction

There are works which aim to reduce the time it takes to perform a handover which may be caused by blockage. This is not to be confused with approaches that reduce the *number* of handovers. These papers reduce handover overhead by improving the beamforming or beam selection process. [30] uses signal to interference and noise ratio (SINR) to predict the direction of currently connected user equipment (UE). The current serving RSU then notifies the other RSU in the UE's line of travel which beams to use to reduce the beamforming overhead. [19] forms a similar approach, where contextual information from the current link is transmitted to the next target RSU, which uses the information to reduce the beam search space. [41] looked at handovers for CAV moving indoors and between indoors and outdoors. The authors predicted the movement of the CAV using a Markov chain and used that information to reduce the handover overhead. These works show with proper

mobility prediction, handover time can be reduced. [28] and [6] both demonstrate the advantages of leveraging positional data to reduce beamforming overhead. [28] studies the impact of GPS position data on beamforming overhead in a real world scenario. The authors find they can predict the best beam to select using GPS data and machine learning. [6] provides a machine learning based framework which utilizes geo-location data to schedule UEs to cells/beams and reduce beam selection overhead. Using their framework, which they apply to a 28 GHz 5G NR network, they are able to reduce UE's initial access to a base station by up to 38%. Their approach is also functional for speeds of up to 150 km/h, meaning it can be applied to vehicular networks. [24] also shows the benefit of location information on network performance. [9] uses deep reinforcement learning to detect when a blockage has occurred and switch links. It detects blockages based on channel state. In [3], authors use machine learning to predict when a mobile agent will perform a handover. Because they can predict the handover time, the RSU can preemptively handover to a new RSU. Their experiments show they can reduce handover time and they consider static blockages. Other approaches opt to use a backup link if the current link becomes blocked. [44] uses pre-connection method to reduce handover time. Some antenna elements to connect to potential RSUs, so when a blockage breaks the link, the vehicle can handover quickly. [32] uses dual-band LTE with mmWave to quickly switch when a drop in quality is detected.

2.3 RSU Selection




Fig. 2. Example of RSU selection using SMART [35]. The vehicle calculates which RSU will yield the highest data transmission before the threshold-based handoff condition is met and selects accordingly. This approach is path aware, but not dynamic blockage aware.




Fig. 3. B-AWARE system model: Blue car has NLOS to r_1 due to blockage from vehicle. Planned position for the blue car is shown. From the blue car to r_2 and r_3 , two distances are given, and the RSS is dependent on the distance. The motion plan and predicted RSS are used to define the optimal RSU selection schedule.

The following related works attempt to design intelligent RSU selection schemes. Due to the high overhead of handovers, most selection schemes are designed to reduce the number of handovers or blocked links. Handovers can be prevented by avoiding blockages and maintaining connection to an RSU for as long as possible. Most literature focuses on using some form of machine learning or statistical model to predict and learn which beams will be blocked by static blockages [9] [2]. [34] tackles the problem of frequent blockage induced handovers in dense mmWave networks. They define a policy where mobile agents associate with base stations by selecting the one with the longest predicted unobstructed LOS time. Their association is performed based on the UE's expected velocity; however, they only consider stationary blockages. [35] implements two reinforcement

learning based RSU selection algorithms which aim to reduce the number of handovers while maintaining a minimum QoS. In their approach, they use machine learning to predict which RSU will yield the best data and select accordingly. This approach applied to vehicles is shown in Figure 2. [27] uses a Markov decision process (MDP) model for RSU selection. The states are defined based on the probability of links being in an outage, LOS, or NLOS. [40] takes a similar approach, using Markov chains to obtain analytical expressions for the blockage probability, average blockage duration, and the SINR distribution. [42] implements a centralized learner which makes association and handover decisions for vehicles. Their algorithm is able to learn the presence of static blockages using a Semi-Markov Decision Process (SMDP). Other papers attempt to reduce the impact of blockage induced handovers; if the handover overhead can be reduced, there is less impact of blockages and base stations can be selected on signal quality alone. [17] proposes a methodology which leverages re-configurable intelligent surfaces (RIS). Where when a blockage is detected, the RSU switches beams to reflect off an intelligent surface and route around blockages to the UE, thus preventing a handover. Their approach does consider mobile blockages, but requires additional infrastructure. [21] solves the blockage induced handover problem through intelligent beam selection. They frame the problem as a contextual multi-armed bandit problem. To reduce handover time, the currently connected RSU makes beam alignment predictions for the target RSU. [2] also defines an intelligent beam selection scheme to avoid blockages, they use past beam observation to predict the next best beam when the primary beam becomes blocked. Although experiments in their work show the low latency of switching to the next best beam, their approach is still reactive – in the sense that it does not predict the impending blockages. In contrast, our approach, B-AWARE, utilizes the near-future path plan of the nearby vehicles to identify the dynamic blockages and plans a better RSU assignment.

3 SYSTEM MODEL

This section will detail the vehicle model, environment model, communication model, and a model to estimate the potential data rate in a channel, considering the static and dynamic blockages between the vehicle and an RSU. Some of the parameters of our system model are captured in Figure 3.

3.1 Vehicle Characteristics

Autonomous vehicles regularly plan their own trajectories and also estimate the near-future trajectories of their nearby vehicles to drive safely. The vehicle's planned position over time is defined as a series of coordinate pairs with time, $((x_1, y_1, t_1), (x_2, y_2, t_2), (x_3, y_3, t_3), \dots)$. Each coordinate is an x and y position associated with a timestep of length Δt . We define the length of time that the vehicle knows its future position as the planning horizon, L_p . Vehicles are equipped with high bandwidth 60 Ghz antennas, allowing them to connect to a single RSU at a time within range S_{range} . These antennas are mounted in the center of the vehicle atop the roof to have the best connection with RSUs. Vehicles also have Lidar and cameras, allowing them to sense objects within a certain radius. Vehicles have 3D environment maps which include the building models and RSU locations. The duration of a vehicle's trip is notated as t_{trip} . Vehicles sizes are given by $W_v \times L_v \times H_v$. Our scenario assumes vehicles of varying sizes, such as sedans, semi-trucks, and buses.

3.2 RSU Characteristics and Deployment

The set of RSUs in a region are notated as $R = \{r_1, r_2, r_3, \dots, r_j\}$, where r_j is the RSU at the j^{th} index. RSUs are equipped with Multiple-In-Multiple-Out (MIMO) 60 Ghz antennas. These antennas can communicate with vehicles within S_{range} . RSUs are connected with one another via a high bandwidth back-haul link. RSUs know the location of other nearby RSUs, as well as their own

location, mounting height, and a 3D environment map. RSUs are mounted along roadsides atop streetlights and in the road median, as shown in Figure 4. Since we are considering dynamic blockages, placing RSUs to simply provide coverage is not enough [14][16]. If there is only a single RSU which a vehicle can connect to, and that link becomes blocked, then the vehicle will lose connection to the network. Such a scenario is shown in Figure 4. In the figure, the large truck has blocked V_{ego} 's LOS with the RSU it was previously connected to. We design our RSU placement such that at any given time, a vehicle has at *minimum* two RSUs in S_{range} . In order for our algorithm to avoid blockages, there needs to be another RSU to connect to if a link becomes blocked. While it is still possible for all potential links to be blocked, the chance of outage is less likely. If the autonomous vehicle is not-connected to the network, then RSUs have to predict the motion of the vehicle [23]. However, if the vehicle is connected to the network, then it may receive the near-future trajectory of other vehicles through communication channels [20].




Fig. 4. Large vehicles can block the connection to other vehicles in the network. V_{ego} is trying to connect to the RSU, but the blocking vehicle has obstructed LOS. This causes V_{ego} to disconnect to the current RSU and connect to a different RSU in LOS.




Fig. 5. Detriment of blockages and handovers on potential throughput as shown by [11]. After the occurrence of a blockage, throughput drops significantly to a point where it is useless for high bandwidth CAV applications. Throughput is only rectified after the CAV disconnects from the previous blocked RSU and connects to a new RSU. The disconnection time is measured as t_D and the time to connect to a new RSU is t_H .

3.3 Association Frequency

The set of timesteps is defined as $T = (t_1, t_2, t_3, \dots, t_n)$ where t_n is the n th timestep and Δt is the duration of the timestep. All timesteps have equal duration. The length of a timestep is chosen to encapsulate the occurrences of network events, like blockage and handover. Vehicles associate to RSUs periodically at each timestep. If a vehicle becomes blocked within a timestep, the blockage will not result in handover until the next timestep. The set of potential RSUs a vehicle can connect to at time t_n is written as $R_{pot,t_n} = \{r_1, r_2, r_3, \dots, r_j \mid \forall r_j \in R\}$. An *association* between the RSU r_j and vehicle V at timestep t_n is expressed as $r_j^{t_n}$. A vehicle can only connect to one primary RSU during a timestep, while RSUs can connect to many vehicles simultaneously during a timestep.

3.4 Channel Model and LOS Categorization

An antenna and receiver have LOS if a line between them with points $P_1 = (X_1, Y_1, Z_1)$ and $P_2 = (X_2, Y_2, Z_2)$ does not intersect with buildings or surrounding vehicle bounding boxes. Vehicles are simplified into rectangular prisms of size $W_V \times L_V \times H_V$, where W_V is the max width of the vehicle, L_V is the length of the vehicle, and H_V is the height of the vehicle, and the radio is at the center of the top of the rectangular prism.

Table 1. AD Table: MCS and Achievable Datarate

Sensitivity Level (K_{MCS})	-47	-51	-54	-63	-64	-68	-78
Datarate (Gbps)	6.75	5.19	4.158	1.386	.8663	.385	.027

For LOS conditions, we apply the free-space path loss model, shown in Equation 1 [45], where d is the distance between the transmitter and receiver in meters, which in our case is the distance between an RSU r_j and the vehicle, V . f is the carrier frequency in Hz and c is the speed of light constant. The resulting path loss is given in dB:

$$PL_{LOS} = 10 \log_{10} \left[\left(\frac{4\pi df}{c} \right)^2 \right]. \quad (1)$$

If there is NLOS, we assume an additional loss of size P to be added to PL_{LOS} from Equation (1). This loss is significant enough to cause disconnections in our system model. The NLOS path loss is given as:

$$PL_{NLOS} = PL_{LOS} + P. \quad (2)$$

$$PL = \begin{cases} PL_{LOS} & \text{if LOS} \\ PL_{NLOS} & \text{otherwise.} \end{cases} \quad (3)$$

The received signal is calculated following Equation (4), where G_{bs} is the antenna gain, G_m is the mobile gain, and P_{tx} is the transmission power in dBm. PL is the pathloss, which is dependent on LOS status and is found from (3). The RSS for a link, $r_j^{t_n}$, can then be defined as $RSS(r_j^{t_n})$:

$$RSS = G_{bs} + G_m + P_{tx} - PL. \quad (4)$$

The potential datarate, D , is found using the RSS with the Modulation and Coding Schemes (MCS) defined in the 802.11ad standard [33]. 802.11ad gives sensitivity thresholds, K_{MCS} , for the minimum RSS required to use a MCS. The MCS that is used affects the datarate achieved. The achievable datarate and the associated thresholds K_{MCS} are in Table 1. We call this table the achievable datarate (AD) table. We denote the datarate of link $r_j^{t_n}$ at timestep t_n as $D(r_j^{t_n})$ and it is chosen with Table 1. K_{MCS} is the minimum RSS required to receive the associated datarate. Thus, the achievable datarate is calculated using AD as a lookup table.

$$D(r_j^{t_n}) = AD[RSS(r_j^{t_n})]. \quad (5)$$

3.5 Data Transmission Calculation

The potential datarate for a timestep can be calculated based on RSS, the RSS is shown in Equation 4. Following equation 3, RSS and consequently datarate, are both functions of RSU and vehicle positions. Vehicles only associate with RSUs at the beginning of a timestep and timesteps are synchronized between all vehicles in the network. As shown in Figure 5, the achievable datarate is also affected by the occurrence of handovers and blockages. Because blockages cause a large drop in RSS, they trigger a handover, which we call a *blockage induced handover*. As shown in, [11], it takes some time for the CAV to attempt recovery and decide to disconnect from the current RSU, t_D . t_D is the additional penalty from the blockage on the handover time. The time to actually connect to the new RSU is labelled as t_H . We assume there is reduced time in a scheduled association. Since the vehicle knows the location of the target RSU, a preemptive connection or GPS location based method can be implemented to reduce the beamforming overhead [6][44]. From the delays described, we

realize 3 handover overheads, O_H , O_{PH} , and O_{BH} . These overheads are the length of time caused by the type of handover where meaningful data cannot be transmitted. The calculation for O_{BH} is shown in Equation 6. Since a standard handover is not triggered by blockage, $O_H = t_H$. We also assume planned handovers take less time, due to an intelligent beamforming implementation, $O_{PH} \leq O_H$.

$$O_{BH} = t_D + t_H. \quad (6)$$

Throughout the duration of a trip, we can compile the total time spent in handover events as a set of intervals consisting of the start time of the handover event t_x and the end time of the handover event $t_x + O_x$ shown in Equation 7. E_x is the set of intervals in a vehicle's trip for which the handover events are occurring and no data can be transmitted. Each interval has length O_x and the start time of each handover event is t_x^i where i is the i th occurrence of handover event x . x is restricted to the three types of handover events, standard handover - H, planned handover - PH, and blocked handover - BH. Similarly, we can compose the same style sets for blockage induced handovers and planned handovers in Equations (9) and (8), respectively.

We can calculate the potential data during a trip as the sum of data, $D(r_j^{t_n})$, at each timestep over that period, minus the occurrence count of each penalty event, multiplied by their respective penalties. We assume that during ANY event regardless of whether it is a handover, blockage induced handover, or planned handover the data rate is effectively zero for our high bandwidth applications [11]. Therefore, the time lost during each timestep, Δt , that would otherwise be transmitting at $D(r_j^{t_n})$ is the length or $\mu()$ of the intersection of all handover event intervals with the current time interval $[t, t + \Delta t)$. Thus, the total potential data transmission rate for a vehicle trip can be found using Equation (11):

$$E_H = \{[t_H^0, t_H^0 + O_H), [t_H^1, t_H^1 + O_H), \dots, [t_H^i, t_H^i + O_H)\} \quad (7)$$

$$E_{PH} = \{[t_{PH}^0, t_{PH}^0 + O_{PH}), [t_{PH}^1, t_{PH}^1 + O_{PH}), \dots, [t_{PH}^i, t_{PH}^i + O_{PH})\} \quad (8)$$

$$E_{BH} = \{[t_{BH}^0, t_{BH}^0 + O_{BH}), [t_{BH}^1, t_{BH}^1 + O_{BH}), \dots, [t_{BH}^i, t_{BH}^i + O_{BH})\} \quad (9)$$

$$G = E_H \cup E_{PH} \cup E_{BH} \quad (10)$$

$$D_{trip} = \sum_{t=0}^n D(r_j^{t_n})(\Delta t - \mu([t, t + \Delta t) \cap G)). \quad (11)$$

From Equation 11, the datarate achieved during a trip of length t_{trip} is dependent on the RSU selection $r_j^{t_n}$ at each timestep. The datarate also depends on the previous RSU selection, $r_j^{t_{n-1}}$, as achievable data is affected by handover overhead. The ordered sequence of RSU selections for a vehicle is given as

$$S^{t_n, t_n + t_{trip}} = (r_j^{t_n}, r_j^{t_{n+1}}, r_j^{t_{n+2}}, \dots, r_j^{t_n + t_{trip}}). \quad (12)$$

4 PROBLEM FORMULATION

We form a maximization problem with Equation (14) as our objective function we seek to maximize. Because the achievable datarate is dependent on the RSU selections, we aim to find the ordered set of vehicle RSU associations, $S^{t_n, t_n + t_{trip}}$, which maximizes the potential datarate for a time interval $(t_n, t_n + t_{trip})$. This makes our problem an RSU, or RSU, association/selection problem. We choose to maximize achievable datarate because it is directly related to signal strength, which is what most

approaches use to select RSUs. The potential data rate of a trip for V from t_n to $t_n + t_{trip}$, shown in Equation (11), is dependent on $S^{t_n, t_n + t_{trip}}$ from Equation (12). Therefore we can rewrite Equation 11 as:

$$D_{trip} = \sum_{t=0}^t D(S^{t_n, t_n + t_{trip}}[n])(\Delta t - \mu([t, t + \Delta t] \cap G)). \quad (13)$$

Our objective function is:

$$\arg \max_{S^{t_n, t_n + t_{trip}}} D_{trip} \quad (14)$$

5 PROPOSED APPROACH

In this section, we define how we utilize vehicle position data and environment information to maximize the total potential bandwidth, D_{trip} , as calculated in Equation (14). As shown in Figure 1, signal strength and trajectory is not enough information to make an optimal RSU selection. In order to select the optimal RSU, we also must consider the trajectory of other vehicles on the road. Using the path planning module, an autonomous vehicle has its planned position up to a certain time, L_P [12, 38]. In addition, vehicles have access to RSU positions along with map data. Using this data, we propose the creation of an *association schedule*, $S^{t_n, t_n + L_P}$, that defines the sequence of RSU selections, $r_j^{t_n}$, for each vehicle to follow. The schedule tells the vehicle which RSU to connect to and when. Once a schedule has been created, an estimation of the achieved potential datarate can be calculated by applying Equation (11). The occurrence of blockages can be predicted by checking if a blocking vehicle's bounding box along its planned path intersects a link between a vehicle and an RSU. Being able to predict if a blockage will occur within L_P for a schedule $S^{t_n, t_n + L_P}$ can help maximize the datarate, as scheduling RSU connections to avoid blockages can reduce link downtime. Since handovers are triggered following the schedule, a vehicle may also choose to connect to an RSU with a stronger connection when other approaches would not.

At a high level, the scheduled association policy can be explained in the following steps, which occur at each timestep on an RSU.

- (1) Each vehicle connects to the RSU according to its connection schedule
- (2) Each vehicle then shares its position, and size with the currently connected RSU.
- (3) Using the knowledge of vehicle trajectories and sizes, the RSU finds the schedule for each vehicle which maximizes the potential datarate.
- (4) Each RSU then sends the optimal blockage aware association schedules back to each vehicle connected to it.

5.1 Vehicle Algorithm

We show the execution context of our RSU selection algorithm with the different software modules running on the CAV in figure 6. The solution for a non-connected RSU will be slightly (but not substantially) different. Our approach relies on the output of the planning module, like the one found in Apollo, a popular open source autonomous driving stack [38]. In hybrid mode, the Apollo planning module takes the map information, obstacles, and current status of the AV to calculate the planned trajectory. This trajectory is found using a combination of a machine learning based planning model and set of scenario algorithms. The RSU Communication module, shown in green, is a new addition which will control the radio to create the connection to the RSU from the schedule and receive information from the RSU. The planned trajectory from the planning module is what would be transmitted to the currently connected RSU. The algorithm on the vehicle needs to




Fig. 6. Integration of our B-AWARE framework integration with Autoware AV stack: B-AWARE primarily relies on the output of the planning module, as we plan RSU connections based on vehicle's future positions.

perform the following new operations at each timestep: At the start of each timestep, the vehicle connects to the RSU from its schedule. If the vehicle is unable to connect to the RSU in its schedule or it is making its initial connection, it will connect to the RSU with the best signal strength.

5.2 Schedule Creation Algorithm on RSU

The RSU schedules of each vehicle are modified by their currently connected RSU. Because RSUs have access to other vehicle's trajectories, they can better predict the schedule which will yield the best potential data. While we are not considering the computation power of the RSU, we assume the correction algorithm execution time must be reasonably small, such that a vehicle can receive an accurate schedule before the start of the next timestep. The selection of an RSU at each time for L_P is a series of actions that can be represented as a directed acyclic graph (DAG). Because the vehicle trajectory is known, the currently connected RSU is also able to find RSUs in the range of the vehicle along its path. At any given time, a vehicle will have a set of RSUs within its transmission range, we label these as potential RSUs, R_{pot,t_n} . With R_{pot,t_n} at each timestep acting as nodes, a DAG is formed to describe all the potential vehicle to RSU connection schedules. Nodes are RSU selections $r_j^{t_n}$ and edges represent the handover from one RSU to another. Edge weights are calculated from Equation 11 with a trip time of 1 timestep.

An example DAG construction is shown in Figure 7. In this example, there are 3 available RSUs at each timestep; however, this may vary. The RSUs which are in range of the CAV change as the vehicle moves through the environment. Node A_0 represents RSU A at the current timestep, t_0 . All nodes have a subscript, which is the future time at which they are available. For example, A is an available RSU for the current vehicle at t_1 and t_2 . Note that the amount of edges leaving any node at t_n is equal to the number of available RSUs at t_{n+1} . This is so all possible schedule combinations are captured. The edges are directed toward the t_{n+1} nodes, so no edges connect nodes within the same timestep to one-another. Following, 11, the potential datarate is dependent on the previously connected RSU. If a vehicle is switching between RSUs, there will be handover overhead. Thus, edge weights entering a node will depend on the node they are exiting. We define the edge weights between nodes N_1 and N_2 as the potential datarate a vehicle would have for a timestep by connecting to the RSU represented by N_2 at t_{n+1} from the RSU represented by N_1 at t_n . The edge weight is calculated from 11. For example, the edge weight from A_0 to B_1 is calculated as




Fig. 7. DAG constructed from RSUs in range at each timestep, which is then used to find optimal schedule. Edge weights are calculated as the expected data achieved by connecting to the RSU represented by the nodes.




Fig. 8. RSU selection schedule construction and integration with path plan.

the potential datarate a vehicle would achieve at t_1 if it were to connect to B at t_1 from A at t_0 . Edge weights, with the exception of those directed toward the *END* node, are excluded for cleanliness.

To find the RSU selection schedule which maximizes the datarate, we add an additional node to the DAG. The start node, N_{start} is set to be the RSU the vehicle is currently connected to. In Figure 7 this is A_0 . The nodes at the final timestep in L_p all connect to an ending node, N_{end} , by edges with weight 0. Edge weights are assigned based on datarate. If any RSUs are predicted be in NLOS at a timestep, all edges directed toward that node will have a weight of an infinitely large negative number. This is because we assume the datarate achieved from a blocked RSU to be zero. The blockage prediction is performed by using the vehicle's planned position in combination with the other vehicle's planned positions and dimensions. The edge weights are then all negated, so that we can find the shortest path from N_{start} to N_{end} in $O(N+E)$, where N is the number of nodes and E is the number of edges[37]. The resulting nodes in the shortest path will be the RSU connection schedule, S^{t_n, t_n+L_p} , which maximizes the datarate in Equation 14.

The complete schedule construction and optimization is illustrated in Figure 8. Each column represents a timestep. At every timestep for L_t , a vehicle has its planned position, pos , and set of potential RSUs. Using the potential RSUs at each timestep in L_p a DAG is constructed and the path leading to the maximum datarate is calculated. This path is highlighted in green and results in the schedule shown in the bottom row. Note the connection schedule is at the granularity of the path planner. The length of the schedule is also the length of the path planner, L_p . At each

Table 2. Experimental Scenarios, environment, RSU count, vehicle count, and RSU placements are all varied.

#	Roadway	RSUs	Vehicles	RSU Placement
1	Highway	12	26	1
2	Highway	12	101	1
3	Highway	12	128	1
4	Highway	24	26	2
5	Highway	24	101	2
6	Highway	24	128	2
7	City	16	101	3 & 4
8	City	16	177	3 & 4
9	City	16	297	3 & 4
10	City	32	101	3 & 4
11	City	32	177	3 & 4
12	City	32	297	3 & 4

timestep, the RSU calculates the best selection schedule for each vehicle connected to it. Because the RSU is connected to the network, it has access to the 3D map of the area and other nearby vehicle trajectories. It uses this information to create the DAG and calculate the best schedule.

6 EXPERIMENT SETUP

To evaluate the effectiveness of our approach, and compare it with previous approaches, we perform experiments in multiple highway and city street traffic scenarios with varying density of vehicles and placement of RSUs. Scenarios are created by importing a 2 kilometer stretch of an interstate highway from Open Street Map (OSM) and generating traffic with Simulation of Urban Mobility (SUMO). The summary of all experimental scenarios is captured in Table 2. The first 6 scenarios are in a highway environment, where RSUs are deployed following the RSU placements shown in Figure 9. For each RSU deployment, the number of vehicles in the simulation is set at 3 traffic levels. The final 6 scenarios are in a city environment. For these scenarios we also vary the RSU deployments and 3 levels of traffic. Specifics on the environment, RSU placement, and traffic characteristics are explained in the following sections.

6.1 City and Highway Environments

We evaluate our approach on two different maps. The first being a 2km long single direction 7 lane highway, which is a segment of the I-10 West taken from OSM. The next map is a 2km x 2km road network consisting of two connected intersections. Lane count ranges from 2-3 per direction and lanes of differing direction are divided by a median. Buildings which can act as blockages are placed along the roadside. The improvement of B-AWARE over other approaches is dependent on dynamic and static blockage occurrence. These events are fully captured in the city and highway environments we use.

6.2 RSU Deployment Strategies

We deploy RSUs in both dense and sparse fashion using the different strategies shown in Figure 9. The arrows represent the direction of traffic flow. Dual side deployments are shown in 1 and 2, where the RSUs can be placed on both sides of the road. In a denser deployment like 2, the vehicles will be more prone to excessive handovers [18], but less prone to blockage. The converse holds true for strategy 1. Strategies 3 and 4 are single side deployments, where they are set along one




Fig. 9. Different RSU deployment strategies. We consider both sparse (1,3,4) and dense (2) deployments. In dense RSU deployments, there are more options for a vehicle if a link becomes blocked.

roadside or along the median. Following the system model presented in Section 3, we deploy RSUs such that at any time a vehicle on the road has at least 2 RSUs within range. This is because our approach is an association policy, which tells the vehicle which RSU to connect to. If there is only 1 RSU available, then such an approach is not needed.

6.3 Traffic Density and Blocking Vehicle Count

The ratio of vehicles to non-blocking vehicles in all scenarios is roughly 3:1. For both the highway and city environments, we vary the traffic density by increasing the amount of vehicles in the SUMO simulation. In the highway scenario, the number of vehicles ranges from 26 to 128. In the city scenario, the number of vehicles ranges from 101-297. We keep the number of vehicles large to ensure blockages occur in the simulation.

6.4 System Model Parameters

The values of the parameters defined in Section 3 for our experiments are shown in Table 3. We apply the channel model described in Section 3, using the freespace pathloss model with $P = 20$, $f = 60$ GHz, and $P_{tx} = 20$ dBm. RSUs and CAVs have a communication range, S_{range} of 200 meters. The dimensions for blocking and non-blocking vehicles are given in the table. We set a timestep length, Δt , of 1 second. We model initial access to RSUs, blockages, and handovers between RSUs as a duration of time when no data can be transmitted [14]. We assume association with an RSU takes 320 milliseconds, whether it be initial access or handover; this was calculated as the average of the two latency's found in [18] and [11]. We treat this as the standard handover time, O_H . Since our proposed method generates a schedule and also shares positional data, we assume an approach similar to [6] is implemented. [6] demonstrated the impact of scheduling RSU selection and geo-location on the beamforming overhead, finding a 38% reduction in initial access time. Other works also show the positive impact of using location information on achievable data rate [13, 24]. Using the percentage from [6] and the average association time, we calculate the RSU association time when connection is scheduled, O_{PH} , to be 198.4 milliseconds. From the real world experiments in [11], there is some penalty when a CAV becomes blocked and must select a new RSU. For our experiments, when a CAV's current connection to an RSU becomes blocked, we assume there is zero data transmitted for that timestep, or a length of 1 second. Therefore, the overhead of a blockage induced handover, O_{BH} , is given a value of 1. This latency sums the time for the CAV to realize there is a blockage and re-associate with the next best RSU. We use these numbers in conjunction with the occurrence count of handover and blockage events to calculate the potential max data rate for each vehicle during the simulation. For simplicity, we assume perfect beam tracking, meaning

Table 3. System Model Parameters

Parameter	Description	Value
PL	Pathloss Model	Freespace
P	Additional NLOS Loss	20
f	Carrier Frequency	60 GHz
P_{tx}	tx Power	20 dBm
S_{range}	Max. tx Distance	200 m
RSU_{thresh}	Baseline and SMART threshold	-54 dBm
$W_c \times L_c \times H_c$	Car Dimensions	21.9 x 2.6 x 4.1 m
$W_B \times L_B \times H_B$	Blocking Vehicle Dimen.	5 x 1.6 x 1.5 m
Δt	Timestep length	1 s
L_P	Length of path plan & sel. sched.	10 s
O_H	Assoc. Time	320 ms
O_{PH}	Scheduled Assoc. Time	198.4 ms
O_{BH}	Blocked Assoc. Time	1 s

there is no overhead induced by a connection being mobile. The length of the path plan, L_P , is set to 10 seconds, which is the time used in current open source AV software stacks [12]. The threshold used by the baseline and SMART algorithms for triggering handoff, RSU_{thresh} , is selected to be -54 dBm. We are assuming CAVs require multiple Gbps at each timestep and -54 dbm is the minimum threshold to achieve such a datarate in our system model, as shown in Table 1.

6.5 SUMO + DRIVE Simulation Testbed

We implement our algorithm in Digital Twin for Self-driving Intelligent Vehicles (DRIVE), which is a MATLAB based simulator that communicates with SUMO via SUMO's Traffic Control Interface (TraCI) [25]. This simulator was chosen so we can leverage the built-in support for modelling 5G mmWave links. Path uncertainty is introduced into the simulator to get closer to a real world scenario. When a vehicle shares its planned position with the serving RSU, there is a 25% chance the position sent to the RSU toward the end of the plan is not the exact position. We assume the CAVs used in the simulation are equipped with an off-the-shelf GNSS receiver and that time synchronization between the vehicles is maintained using the GNSS time sync. Hasan *et. al* have shown between two vehicles to be capable of synchronizing to within 30 nanoseconds, which is far more precise than required for our application [15]. We propose that the vehicles can also conform to the gPTP standard defined in the IEEE 802.1AS standard, which provides for microsecond accurate synchronization. This allows us to have a multi modal fallback system to maintain the highest available timesync. If gPTP is available from any peer vehicles or RSU within range of the ego CAV, that is used first. If that fails, GNSS is used as a fallback. Finally, if neither is available we rely on the clock drift to be corrected and low within the CAV systems such that the drift during these relatively short lapses will be minimized. As soon as GNSS or gPTP is available, those are used again. The error from either of these methods should be on the order of in the worst case 5 ms and our system can handle the error with relatively small downside. An error of 5 ms or smaller will be easily contained in our sensor error buffer with little loss, which we define next.

Sensing and timing error is modeled using the method outlined by Andert *et. al* by adding a buffer around the vehicle within SUMO that is equivalent to the expected sensor error as we cannot be certain where the vehicle is within that zone [5]. This model shown in figure 10 is similar to the sensing error model used by Andert *et. al* in their time-sensitive autonomous intersection model

[4]. Positioning error over time can increase. However, this increase in positioning error can be recorded over time by the B-AWARE system and added back into future trajectories as a steadily increasing position error buffer, shown in Figure 11. This increasing position buffer accounts for reasonable trajectory uncertainty in a computationally efficient manner, as it only affects the width and length of the vehicle in B-AWARE's computation, thus adding negligible overhead.




Fig. 10. To account for position error, a static buffer can be added around the vehicle to account for the position uncertainty. B-AWARE simply considers the vehicles area as larger. Typical position localization (position) uncertainty is less 10cm for an autonomous vehicle so this buffer is not large.




Fig. 11. Throughout the path plan of the vehicle, typical position and plan uncertainty can be applied as an increase in the position error buffer over time. This allows B-AWARE to account for growing uncertainty in the path plan as the time horizon increases.

6.6 Baseline and SMART Comparison Algorithms

For each scenario, we assess the performance of our approach against two other RSU handoff and selection algorithms. The first being a threshold-based handoff algorithm, like the one found in 802.11ad, which we call baseline. If the RSS drops below a threshold, a handoff will occur. The CAV then selects the RSU with the highest RSS to associate with. This can also be described as rate-based handoff (RBH), since the highest RSS will yield the best potential data rate [35]. This selection policy was chosen for comparison as it most closely resembles the behaviour of current consumer off-the-shelf 5G hardware. We also compare our approach to an idealized version of the SMART-S algorithm proposed in [35]. SMART-S is a state-of-the-art selection algorithm applicable to vehicular networks for solving the 5G matching problem. To select the next RSU, SMART implements a reward function which minimizes the number of handoffs by selecting the RSU which can transmit the most amount of data from time t to t_n^k when a vehicle switches to RSU k . t_n^k is the time at which the next handoff occurs based on the handoff condition. Because the time before the next handoff and the data rate achieved at each time are unknown, researchers implemented a reinforcement learning algorithm to estimate these parameters. For our experiments, we have access to vehicle's paths and RSU placement before hand, so we can calculate the future volume of transmitted data for each potential RSU selection. The handoff condition will be based off of event A2 defined in 3GPP, but instead of using reference signal receive power (RSRP), we will handoff based on RSS which is available in DRIVE [1, 25]. Therefore SMART uses the same handoff trigger as the one in baseline. In our experiments, handoffs are triggered when the RSS falls below a prescribed threshold, RSU_{thresh} , see Table 3.

6.7 BAWARE Performance Analysis

In addition to comparing our method to existing approaches, we also examine the behaviour of our algorithm through the following two experiments. First, we measure the execution time of the primary RSU schedule optimization algorithm described in Section 5.2. The DAG is constructed based on the RSUs available to a vehicle at each future timestep for the length of the planning




Fig. 12. Average potential datarate improvement from using B-AWARE over SMART and Baseline for all vehicles in each scenario. B-AWARE consistently provides a higher potential datarate on average in all scenarios.




Fig. 13. Best and worst case potential datarate improvement from using B-AWARE over SMART and Baseline for a single vehicle in each scenario. In the best case, B-AWARE can improve the datarate of a single vehicle by 25% over existing approaches. In the worst case, there may be some loss in datarate, but this is much lower than the maximum improvement in all scenarios.

horizon. Finding available RSUs (i.e. ones that are not blocked by other vehicles) can be done through ray tracing, which modern GPUs can perform in a matter of milliseconds [10]. We focus our performance analysis on the run time of the schedule optimization algorithm, which occurs after the DAG has been constructed. For this experiment, we consider a scenario where a single RSU is calculating the optimal schedule for a single vehicle, with 4 potential RSUs at each timestep for a length of 10 timesteps. This creates a DAG with 42 nodes and 168 edges. Our optimization algorithm and example DAG were implemented in C++ using the Boost Graph Library (BGL). The average execution time over 20 independent runs was measured on a PC equipped with an AMD Ryzen 7 3700X and 16GB of RAM. In the second experiment, we look at the impact of the planning horizon on our algorithm's improvement over other approaches. We do this by running scenario 6 multiple times while varying the length of the planning horizon between 3 and 7.

7 RESULTS DISCUSSION

This section covers our experimental results and related insights.




Fig. 14. Best case improvement in average datarate for a vehicle against the existing approaches. As the traffic count increases, the potential benefit of B-AWARE does as well. In higher traffic densities, blockages may occur for a vehicle more frequently. Since other approaches do not account for blockage and B-AWARE does, we see a larger increase in average datarate




Fig. 15. B-AWARE can reduce almost 100% of blockage occurrences over Baseline and SMART in the 12 scenarios.

7.1 Better Potential Data over Previous Approaches

Figure 12 shows the average improvement in datarate for all vehicles using B-AWARE over Baseline and SMART. Total potential bandwidth is calculated for each vehicle using equation 11, e.g. summing the available bandwidth for each time-step according to equation 5 when there is no blockage or handover occurring. This is then divided by the vehicle's trip time to get the average datarate. Each scenario, with the same traffic traces, is ran 3 times, one time for each algorithm, the average datarate of each vehicle is then calculated and compared across the 3 simulations. The percent improvement of B-AWARE over each other approach is calculated and then averaged for all the vehicles in the simulation. In all scenarios, B-AWARE is able to maintain a consistent improvement in average datarate over the two existing approaches, ranging from 2%-6%. SMART achieves a better potential datarate over the normal approach because it selects RSUs based on vehicle trajectory. Although it is *path aware*, SMART is not *blockage aware* like B-AWARE. This is why B-AWARE can achieve a better average potential datarate over the baseline approach and SMART. Since




Fig. 16. B-AWARE can reduce the average disconnection time by as much as 75% in the highway scenarios. Even as traffic density increased, the benefit of using B-AWARE over existing approaches remains constant.




Fig. 17. B-AWARE can reduce the average disconnection time by as much as 80% in the city scenarios. Even as traffic density increases, the benefit of using B-AWARE over existing approaches remains constant.




Fig. 18. The length of the planning horizon impacts the improvement of our approach over the baseline RSU association policy. For scenario 6, the planning horizon length is varied from 3 to 5 to 7. As length of the planning horizon increases, BAWARE can form a more optimal schedule, leading to a larger improvement over the baseline approach. Even if the planning horizon is short, we still see an improvement over the baseline RSU association policy.

the baseline approach and SMART are both threshold-based, they will wait to handover until the threshold has been broken, even if there is an RSU with a better signal strength available. B-AWARE does not have this limitation.

Although B-AWARE is able to improve the datarate on average, it is not a very large improvement. This is because the main benefit of B-AWARE is its ability to avoid blockages. Avoiding a single blockage or handing over to gain an additional Gigabit of throughput does not have much effect on the total average potential data. Because on average some vehicles may not experience a high amount of blockages. To see the true impact of B-AWARE, we look at which vehicle had the largest percent improvement in potential datarate by using B-AWARE over the other approaches in each scenario. Since the same traffic traces are run for each algorithm and scenario, we chart the vehicle which saw the maximum and minimum percent improvement in average potential datarate over Baseline and SMART by using B-AWARE, as shown in Figure 13. In the best case, B-AWARE can improve the average datarate for a vehicle by as much as 26% over the existing approaches. This

large improvement is due to a vehicle travelling alongside multiple blocking vehicles. Using a selection policy like Baseline, the vehicle has no knowledge of its surroundings and will continue attempting to connect to RSUs which will be shortly blocked. In the worst case, the degradation remains around 0%. There are some cases, like in scenario 6, where the worst case reaches a 7% degradation in average potential throughput for B-AWARE. This is a result of some vehicles having short trips, where B-AWARE is not able to fully optimize from the initial RSU selection. However, it can be seen the max improvement is always greater than the maximum degradation and using B-AWARE results in a higher datarate on average.

7.2 Significant Outage Reduction over Previous Approaches

The increased potential datarate is primarily a result of the decreased connection time, allowing more time for data transmission. The reduction in connection time is a result of reduced handover overhead and blockage avoidance, as the time without connection is calculated as the sum of time a handover is occurring and when a link is blocked. A blockage event is logged any time a vehicle losses LOS with the RSU it is currently connected to. Figure 15 shows the percent reduction in total blockage occurrences by using B-AWARE compared to the existing approaches. Each scenario is ran 3 times using each approach; Baseline, SMART, and B-AWARE. The percent decrease in blockage count during the simulation between the existing approaches and B-AWARE are charted for each scenario. From Figure 15, it is clear B-AWARE can avoid almost 100% of the blockages which occur in all experiment scenarios. In one scenario B-AWARE reduced the total number of blockages in the simulation by as much as 573. If we assume a larger blockage penalty, like the 2.7 seconds from [18] that is 25 minutes of cumulative avoided outage time in a 200 second long period. The percent decrease in average disconnection time for all scenarios are captured in Figures 16 and 17. For each scenario, we calculate the average disconnection time among the vehicles for Baseline, SMART, and B-AWARE. We then calculate the percent improvement of B-AWARE over the other approaches. The highway scenarios are shown in Figure 16. Each RSU deployment strategy is represented by a separate line. It is shown in all scenarios, B-AWARE is able to reduce the average vehicle disconnection by a minimum of 50%. Figure 17 is constructed in similar fashion and also shows that on average, B-AWARE can decrease the link downtime by about 50% over existing approaches.

7.3 Consistent Performance with Traffic Scaling

We can also see from Figures 16 and 17 that our approach still succeeds in higher traffic density relative to the other approaches on average. In addition, the potential benefit of using B-AWARE increases as the traffic density increases; this is shown in Figure 14. Using the 12 scenarios, we take the maximum percent datarate improvement, shown in Figure 14, at each traffic density and average them together. As the traffic density increases, the probability of a vehicle experiencing more blockages also increases. Since other approaches are not blockage aware, B-AWARE's improvement over them increases as the blockage rate increases. B-AWARE's improvement is also shown to be consistently better in the average case by Figures 16 and 17, regardless of traffic density.

7.4 Robustness to Environment and Path Changes

Even with path uncertainty set to 25%, our approach is still able to predict blockages and significantly reduce the number of link outages for a vehicle. This is in addition to unexpected path changes that the SUMO lane changing and lane following models themselves introduce as our transmitted paths don't account for unexpected lane changes and speed deviations caused by other vehicles until they actually happen. We also show our approach can adapt to different traffic densities and RSU deployment strategies, as shown in Figure 15. Regardless of RSU deployment, traffic,

and environment, B-AWARE performs consistently well over existing approaches. The average improvement in datarate and link uptime for B-AWARE also remains constant, as shown in Figures 14, 17, and 16 despite differences in scenarios, traffic density, and RSU deployments. Figure 15 shows that even with a short planning horizon, our algorithm still improves upon the baseline RSU association policy. As the planning horizon is longer, BAWARE is able to optimize the RSU associations further, leading to a larger increase in total potential datarate.

7.5 Fast Execution Time

In order for our approach to be implemented, it must be able to execute at high speed and finish the optimal schedule computation before the start of the next timestep. Modern GPUs can perform blockage detection through ray tracing in a matter of milliseconds [10]. Once the available RSUs are found through blockage detection and path prediction, the schedule optimization algorithm runs on the constructed DAG. From the experiment described in Section 6.7, we found it takes an average of 336550 nanoseconds or 0.00033655 seconds for an RSU to compute the optimal schedule for a single vehicle. These results are from a "worst-case" scenario, where the planning horizon is 10 seconds and there are 4 potential RSUs at each timestep. In a dense traffic scenario where a single RSU is serving 200 vehicles, it would take an RSU approximately only 0.06731 seconds to find the optimal schedule for all vehicles. This leaves plenty of time for the RSU to perform the raytracing and communicate the optimal RSU selection schedule to the connected vehicles.

8 CONCLUSION

In this paper, we have addressed the current challenges of applying 5G mmWave to autonomous vehicles, namely maintaining line of sight with the serving RSU and minimizing the occurrence of handovers. We developed the B-AWARE framework, which defines an optimization problem to maximize the potential datarate for a vehicle's trip. B-AWARE defines an RSU selection schedule, which is created for a vehicle by the current serving RSU from the most current environment and vehicle trajectory information. This schedule is optimized to produce the maximum potential datarate for a vehicle's planning horizon by avoiding blockages and unnecessary handovers. Our simulation evaluations of B-AWARE in the SUMO and DRIVE test-bed results in a 1.05x improvement of the potential datarate in the average case and 1.28x in the best case vs. the state-of-the-art algorithm SMART. B-AWARE reduces close to 100% of blockage occurrences resulting in an impressive reduction in time spent with no connection - 42% in the average case and 60% in the best case as compared to SMART. Our approach is shown to give CAVs a significant increase in link up-time as well as a higher average potential datarate over traditional and state-of-the-art RSU selection strategies. In addition, B-AWARE is shown to be robust to trajectory prediction failures, traffic induced chaos, RSU deployment differences, and different traffic scenarios. In the future we would like to explore the application of B-AWARE to full-size Autonomous Vehicles so that we can further improve the B-AWARE approach by applying it in the physical world.

REFERENCES

- [1] 3GPP TS 36.331. 2010. E-UTRA Radio Resource Control (RRC); Protocol specification (Release 9).
- [2] Adel Aldalbahi, Farzad Shahabi, and Mohammed Jasim. 2021. Instantaneous beam prediction scheme against link blockage in mmwave communications. *Applied Sciences* 11, 12 (2021), Art. no. 5601.
- [3] Ahmed Alkhateeb, Iz Beltagy, and Sam Alex. 2018. Machine learning for reliable mmWave systems: Blockage prediction and proactive handoff. In *Proc. IEEE GlobalSIP*. IEEE, New York, NY, USA, 1055–1059.
- [4] Edward Andert, Mohammad Khayatian, and Aviral Shrivastava. 2017. Crossroads: Time-sensitive autonomous intersection management technique. In *Proceedings of the 54th Annual Design Automation Conference 2017*. 1–6.
- [5] Edward Andert and Aviral Shrivastava. 2022. Accurate Cooperative Sensor Fusion by Parameterized Covariance Generation for Sensing and Localization Pipelines in CAVs. In *2022 IEEE 25th International Conference on Intelligent*

- Transportation Systems (ITSC)*. IEEE, 3595–3602.
- [6] Marius Arvinte, Marcos Tavares, and Dragan Samardzija. 2019. Beam management in 5G NR using geolocation side information. In *Proc. IEEE CISS*. IEEE, New York, NY, USA, 1–6.
 - [7] Mate Boban, Diego Dupleich, Naveed Iqbal, Jian Luo, Christian Schneider, Robert Müller, Ziming Yu, David Steer, Tommi Jämsä, Jian Li, et al. 2019. Multi-band vehicle-to-vehicle channel characterization in the presence of vehicle blockage. *IEEE Access* 7 (2019), 9724–9735.
 - [8] Ananya Chattopadhyay, Aniruddha Chandra, and Chayanika Bose. 2021. Impact of RSU Height on 60 GHz mmWave V2I LOS Communication in Multi-lane Highways. In *Proc. IEEE VTC2021-Spring*. IEEE, New York, NY, USA, 1–5.
 - [9] Sheng Chen, Kien Vu, Sheng Zhou, Zhisheng Niu, Mehdi Bennis, and Matti Latva-Aho. 2020. 1 A deep reinforcement learning framework to combat dynamic blockage in mmWave V2X networks. In *Proc. IEEE 6G SUMMIT*. IEEE, New York, NY, USA, 1–5.
 - [10] Petrik Clarberg, Simon Kallweit, Craig Kolb, Pawel Kozlowski, Yong He, Lifan Wu, Edward Liu, Benedikt Bitterli, and Matt Pharr. 2022. Real-Time Path Tracing and Beyond. HPG 2022 Keynote.
 - [11] Mohammed Dahhani, Gentian Jakllari, and André-Luc Beylot. 2019. Association and reliability in 802.11 ad networks: An experimental study. In *Proc. IEEE LCN*. IEEE, New York, NY, USA, 398–405.
 - [12] Autoware Foundation. 2023. Autoware. <https://github.com/autowarefoundation/autoware>
 - [13] Nil Garcia, Henk Wymeersch, Erik G Ström, and Dirk Slock. 2016. Location-aided mm-wave channel estimation for vehicular communication. In *Proc. IEEE SPAWC*. IEEE, New York, NY, USA, 1–5.
 - [14] Sanjay Goyal, Marco Mezzavilla, Sundeep Rangan, Shivendra Panwar, and Michele Zorzi. 2017. User association in 5G mmWave networks. In *Proc. IEEE WCNC*. IEEE, New York, NY, USA, 1–6.
 - [15] Khondokar Fida Hasan, Yanming Feng, and Yu-Chu Tian. 2018. GNSS time synchronization in vehicular ad-hoc networks: Benefits and feasibility. *IEEE Transactions on Intelligent Transportation Systems* 19, 12 (2018), 3915–3924.
 - [16] Ish Kumar Jain, Rajeev Kumar, and Shivendra S Panwar. 2019. The impact of mobile blockers on millimeter wave cellular systems. *IEEE Journal on Selected Areas in Communications* 37, 4 (2019), 854–868.
 - [17] Long Jiao, Pu Wang, Amir Alipour-Fanid, Huacheng Zeng, and Kai Zeng. 2021. Enabling efficient blockage-aware handover in RIS-assisted mmWave cellular networks. *IEEE TWC* 21, 4 (2021), 2243–2257.
 - [18] Kishor Chandra Joshi, Rizqi Hersyandika, and R Venkatesha Prasad. 2019. Association, Blockage, and Handoffs in IEEE 802.11 ad-Based 60-GHz Picocells—A Closer Look. *IEEE Systems Journal* 14, 2 (2019), 2144–2153.
 - [19] Wang Junsheng, Chen Yawen, Lu Zhaoming, Wen Xiangming, and Wang Zifan. 2019. A low-complexity beam searching method for fast handover in mmWave vehicular networks. In *Proc. IEEE WCNCW*. IEEE, New York, NY, USA, 1–6.
 - [20] Mohammad Khayatian, Mohammadreza Mehrabian, Harshith Allamsetti, Kai Wei, Po Yu, Chung-Wei Lin, and Aviral Shrivastava. 2021. Cooperative Driving of Connected Autonomous Vehicles Using Responsibility-Sensitive Safety (RSS) Rules, In Proceedings of the 12th ACM/IEEE International Conference on Cyber-Physical Systems (ICCPs) (2021-04-10). ICCPS, 11–20. <https://mpslab-asu.github.io/publications/papers/Khayatian2021ICCPs.pdf>, paper <https://mpslab-asu.github.io/publications/slides/Khayatian2021ICCPs.pptx>, slide
 - [21] Xiaotong Li, Ruiting Zhou, Ying-Jun Angela Zhang, Lei Jiao, and Zongpeng Li. 2020. Smart vehicular communication via 5G mmWaves. *Computer Networks* 172 (2020), 107173.
 - [22] Zhiyuan Li and Weidong Wang. 2018. Handover performance in dense mmWave cellular networks. In *Proc. IEEE WCSP*. IEEE, New York, NY, USA, 1–7.
 - [23] Jianbang Liu, Xinyu Mao, Yuqi Fang, Delong Zhu, and Max Q.-H. Meng. 2022. A Survey on Deep-Learning Approaches for Vehicle Trajectory Prediction in Autonomous Driving. In *Proc. IEEE Int. Conf. on Robotics and Biomimetics (ROBIO)*. IEEE Press, Sanya, China, 978–985. <https://doi.org/10.1109/ROBIO54168.2021.9739407>
 - [24] Yi Lu, Mikhail Gerasimenko, Roman Kovalchukov, Martin Stusek, Jani Urama, Jiri Hosek, Mikko Valkama, and Elena Simona Lohan. 2020. Feasibility of location-aware handover for autonomous vehicles in industrial multi-radio environments. *Sensors* 20, 21 (2020), Art. no. 6290.
 - [25] Ioannis Mavromatis, Robert J Piechocki, Mahesh Sooriyabandara, and Arjun Parekh. 2020. Drive: A digital network oracle for cooperative intelligent transportation systems. In *Proc. IEEE ISCC*. IEEE, New York, NY, USA, 1–7.
 - [26] Ioannis Mavromatis, Andrea Tassi, Robert J Piechocki, and Andrew Nix. 2018. Efficient V2V communication scheme for 5G mmWave hyper-connected CAVs. In *Proc. IEEE ICC Workshops*. IEEE, New York, NY, USA, 1–6.
 - [27] Marco Mezzavilla, Sanjay Goyal, Shivendra Panwar, Sundeep Rangan, and Michele Zorzi. 2016. An MDP model for optimal handover decisions in mmWave cellular networks. In *Proc. IEEE EuCNC*. IEEE, New York, NY, USA, 100–105.
 - [28] João Morais, Arash Behboodi, Hamed Pezeshki, and Ahmed Alkhateeb. 2022. Position aided beam prediction in the real world: How useful GPS locations actually are? *arXiv preprint arXiv:2205.09054* (2022).
 - [29] Clare Mutzenich, Szonya Durant, Shaun Helman, and Polly Dalton. 2021. Updating our understanding of situation awareness in relation to remote operators of autonomous vehicles. *Cognitive Research: Principles and Implications* 6, 1 (2021), 1–17.

- [30] Raul Parada and Michele Zorzi. 2018. Context-aware handover in mmWave 5G using UE's direction of pass. In *Proc. EW*. VDE, Catania, Italy, 1–6.
- [31] Hannah Parr, Catherine Harvey, and Gary Burnett. 2023. Investigating Levels of Remote Operation in High-Level On-Road Autonomous Vehicles using Operator Sequence Diagrams. DOI: 10.21203/rs.3.rs-2510863/v1 (2023).
- [32] Michele Polese, Marco Giordani, Marco Mezzavilla, Sundeeep Rangan, and Michele Zorzi. 2017. Improved handover through dual connectivity in 5G mmWave mobile networks. *IEEE JSAC* 35, 9 (2017), 2069–2084.
- [33] Bernhard Schulz. 2017. 802.11 ad-WLAN at 60 GHz—A Technology Introduction. White Paper.
- [34] Li Sun, Jing Hou, and Tao Shu. 2020. Spatial and temporal contextual multi-armed bandit handovers in ultra-dense mmWave cellular networks. *IEEE TMC* 20, 12 (2020), 3423–3438.
- [35] Yao Sun, Gang Feng, Shuang Qin, Ying-Chang Liang, and Tak-Shing Peter Yum. 2017. The SMART handoff policy for millimeter wave heterogeneous cellular networks. *IEEE TMC* 17, 6 (2017), 1456–1468.
- [36] Anup Talukdar, Mark Cudak, and Amitava Ghosh. 2014. Handoff rates for millimeterwave 5G systems. In *Proc. IEEE VTC Spring*. IEEE, New York, NY, USA, 1–5.
- [37] Robert Endre Tarjan. 1983. *Data Structures and Network Algorithms*. SIAM, Philadelphia, PA, USA.
- [38] Baidu Apollo team (2017). 2023. Apollo: Open Source Autonomous Driving. <https://github.com/ApolloAuto/apollo>
- [39] Caglar Tunc, Mustafa F Özkoç, Fraida Fund, and Shivendra S Panwar. 2020. The blind side: Latency challenges in millimeter wave networks for connected vehicle applications. *IEEE TVT* 70, 1 (2020), 529–542.
- [40] Caglar Tunc and Shivendra S Panwar. 2021. Mitigating the impact of blockages in millimeter-wave vehicular networks through vehicular relays. *IEEE Open J. Intelligent Transp. Sys.* 2 (2021), 225–239.
- [41] Aysenur Turkmen, Shuja Ansari, Paulo Valente Klaine, Lei Zhang, and Muhammad Ali Imran. 2021. IMPRESS: Indoor Mobility Prediction Framework for Pre-Emptive Indoor-Outdoor Handover for mmWave Networks. *IEEE Open J. Commun. Soc.* 2 (2021), 2714–2724.
- [42] Nguyen Van Huynh, Diep N Nguyen, Dinh Thai Hoang, and Eryk Dutkiewicz. 2021. Optimal beam association for high mobility mmwave vehicular networks: Lightweight parallel reinforcement learning approach. *IEEE Transactions on Communications* 69, 9 (2021), 5948–5961.
- [43] Song Wang, Jingqi Huang, and Xinyu Zhang. 2020. Demystifying millimeter-wave V2X: Towards robust and efficient directional connectivity under high mobility. In *Proc. ACM MobiCom*. ACM, New York, NY, USA.
- [44] Xiong Wang, Linghe Kong, Jintao Wu, Xiaofeng Gao, Hang Wang, and Guihai Chen. 2019. mmHandover: A pre-connection based handover protocol for 5G millimeter wave vehicular networks. In *Proc. IEEE IWQoS*. IEEE, New York, NY, USA, 1–10.
- [45] Atsushi Yamamoto, Koichi Ogawa, Tetsuo Horimatsu, Akihito Kato, and Masayuki Fujise. 2008. Path-loss prediction models for intervehicle communication at 60 GHz. *IEEE Transactions on Vehicular Technology* 57, 1 (2008), 65–78.