Received 17 June 2024; accepted 9 July 2024. Date of publication 15 July 2024;
date of current version 26 July 2024. The review of this article was coordinated by Editor Koji Yamamoto.

Digital Object Identifier 10.1109/OJVT.2024.3427326

Data Offloading Over Vehicular DTNs: City-Wide Feasibility Study in Nagoya

TAKAMASA HIGUCHI [1] **(Member, IEEE), LEI ZHONG** [1] **(Member, IEEE), AND RYOKICHI ONISHI**[1]

(Invited Paper)

Toyota Motor Corporation, Tokyo 100-0004, Japan

CORRESPONDING AUTHOR: TAKAMASA HIGUCHI (e-mail: takamasa_higuchi@mail.toyota.co.jp).

ABSTRACT The increasing network traffic from connected vehicles is putting a strain on the limited bandwidth resources of cellular networks. Delay-tolerant networking (DTN) over vehicle-to-vehicle (V2V) communications has been considered as an effective means of offloading the cellular data traffic, while its quantitative performance in urban road traffic remains unclear in many aspects. In this paper, we unveil the benefits of data offloading over vehicular DTNs by city-scale network simulations in Nagoya, Japan. The simulation scenario embraces more than 8 million vehicle trips over five consecutive days. The vehicle routes are carefully calibrated against public statistics on the road traffic volume to enable realistic simulations of V2V communication opportunities between vehicles on the road. The results indicate the strong potential of vehicular DTNs in mixed urban road traffic, comprised of both public transport and privately owned vehicles – a large amount data traffic can be offloaded from cellular networks to V2V communication networks even with the limited ratio of vehicles participating the vehicular DTNs.

INDEX TERMS Cellular data offloading, vehicle-to-vehicle (V2V) communications, delay-tolerant networks.

I. INTRODUCTION

The proliferation of rich automotive data contents keeps increasing the communication demand from / to connected vehicles. Vehicles download large data files like software update packages and high-definition digital maps from a remote content distribution server, while uploading part of vehicle-generated sensor data to cloud computing platforms for data analytics. Some sources estimate that the total volume of data to be transferred between vehicles and the cloud will amount for 100 petabytes per month by 2025 [1].

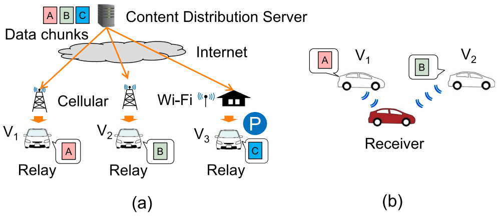

FIGURE 1. **Content distribution over vehicular DTNs.**

DTN over opportunistic V2V communications holds promise to offload the ever-increasing cellular network traffic from / to connected vehicles. When a content distribution server has a data content to be delivered to a certain group of receiver vehicles, the server may designate a small number of *relay vehicles* that help distribute the content. The server divides the original data content into smaller chunks and sends them to the relay vehicles either by way of cellular communications or WLAN hotspots as in Fig. 1(a). The relay vehicles cache the received data chunks in their on-board data storage device and later forward them over V2V communications when encountering any receiver vehicle needing that content (see Fig. 1(b)). The similar idea can also be applied to uplink communications from vehicles to the cloud. Many automotive applications (*e.g.,* high-definition mapping) require vehicle-generated sensor data to be collected to a cloud-based data repository. Instead of having a myriad of vehicles individually send their data to the cloud, a service operator may designate a set of relay vehicles that intermediate the data collection. As illustrated in Fig. 2(a), the relay vehicles collect sensor

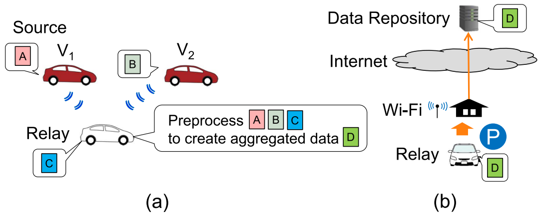

FIGURE 2. Sensor data aggregation over vehicular DTNs.

data from other vehicles (*i.e., source vehicles*) over V2V communications. The relay vehicles then aggregate the collected data to reduce the data volume before uploading them to cloud computing platforms (see Fig. 2(b)).

Data offloading over vehicular DTNs usually takes orders of magnitude longer delivery time than cellular networks, as receiver vehicles can download the contents only when they meet any relay vehicles holding the data chunks of interest. Although this can be a critical drawback for latency-sensitive network applications (*e.g.*, video streaming), there are various classes of automotive data contents that have much relaxed requirements on delivery latency (e.g., software update packages). Large data contents that can tolerate the delivery latency of several hours or days can be reasonably offloaded to vehicular DTNs without degrading the user experience. This paper focuses on such delay-tolerant network traffic in the automotive use cases.

Despite its promising benefits, the feasibility of data offloading over vehicular DTNs still remains unclear in many aspects. Sathiamoorthy et al. [2] were among the first to address this research question. They analyzed real vehicle traces of taxis and buses to estimate the latency of content distribution over V2V communications. They conducted network simulations using the vehicle traces as inputs, and concluded that 1 GB of data can be retrieved in about 16 hours. However, it should be noted that public transport vehicles often have unique characteristics in their trips. Buses travel along planned routes, following predefined time schedules, while taxis keep traveling in the city throughout a day, visiting lots of places. The simulation results based on the traces of public transport vehicles may not be representative of V2V communication networks, formed by privately owned vehicles, which constitute a majority of urban road traffic. Real traces of privately owned vehicles, however, are privacy-sensitive information, hence rarely made available for public use.

A possible alternative is to rely on road traffic simulations. The research community has made a significant effort to develop road traffic simulation scenarios suitable for vehicular network simulations [3], [4], [5], [6], [7], [8], [9], [10]. However, most of the publicly-available road traffic simulation scenarios use demographic statistics of the cities to generate synthetic travel demand, hence the number of vehicles along roads may not necessarily reflect the reality. This causes a challenge when simulating V2V-assisted cellular data offloading, where the frequency and duration of contacts between vehicles make a significant impact on the communication performance. The waiting time at traffic lights is also a critical factor, affecting the duration of V2V contacts, while few of the existing open traffic scenarios calibrate traffic light cycles with the real statistics. These limitations should be addressed to ensure the plausibility of the simulation results.

To enable realistic simulations of data offloading over vehicular DTNs, the authors have developed the Nagoya Urban Mobility (NUMo) scenario, modeling realistic vehicle mobility of the whole city of Nagoya, Japan [11]. The scenario can be executed in the open-sourced road traffic simulator SUMO, which has been widely used in the vehicular networking research community [12], [13], [14]. The NUMo scenario covers a 326 km^2 area spanning over the whole city of Nagoya. The road traffic is carefully calibrated with open traffic volume statistics from Japan Road Traffic Census [15], while traffic signal cycles are also aligned with a separate open dataset [16], released by the Japan Traffic Information Center (JARTIC). The scenario embraces 1.6 million vehicle trips for a 24-hour period, which can be repeated for long-term simulations across multiple days. The scenario is made available to the community at `https://github.com/ToyotaInfoTech/numo` under open-source licenses.

In this paper, we leverage the NUMo to unveil the feasibility of data offloading over vehicular DTNs in the dense urban road traffic in Nagoya. We conduct a series of network co-simulations, coupling the SUMO, the NUMo scenario and a network simulator, to evaluate key performance indices of V2V communication networks. In addition, we use the NUMo scenario to simulate V2V communication networks in Nagoya for five consecutive weekdays, and evaluate the latency of content distribution over vehicular DTNs. The results indicate that a large amount data traffic can be offloaded from cellular networks even in the early stage of deployment with a limited volume of vehicles contributing to the cooperative content distribution over V2V communication networks.

II. RELATED WORK AND CONTRIBUTION

Data offloading from cellular networks has been an active research domain in vehicular networking. A major approach is to have vehicles offload data traffic via roadside units (RSUs), supporting IEEE 802.11p radios. Ota et al. [17] proposed a method called Max-throughput Min-delay Cooperative Downloading (MMCD), which prioritizes content requests from vehicles based upon content delivery deadlines. This helps reduce the response time for the content requests, while offloading as much network traffic as possible to RSUs. Malandrino et al. [18] predict future paths of requesting vehicles and preload the requested contents on the RSUs along the predicted paths. The communication latency between the Internet-based content distribution server and RSUs is often much longer than the radio access network between vehicles and the RSUs. The content preloading on RSUs resolves this performance bottleneck and increases the amount of data that can be transferred within a short contact duration between requesting vehicles and RSUs. Shen et al. [19] improve the efficiency of traffic offloading by adjusting the order of data contents to be distributed from RSUs to vehicles. Although traffic offloading via RSUs is a promising solution to mitigate cellular network load, its effectiveness is highly dependent on the density of RSUs along roads. RSUs supporting IEEE 802.11p or C-V2X communications are still in the early stage of deployment, hence vehicles would have limited opportunities to find them on the road. In addition, ITS bands typically have limited bandwidth (i.e., 30 MHz band in the U.S. and Europe, while 10 MHz band in Japan). It is questionable whether these limited channel resources can support the distribution of large data contents. A possible alternative is to use public WLAN hotspots along roads. However, the communication performance may be limited for vehicles traveling on the roads, since many of the WLAN hotspots are deployed indoors, hence radio signals are often attenuated by buildings and other obstacles.

V2V-assisted content distribution would have strong potential to address the challenges, associated with infrastructure-based traffic offloading. A seminal work in this line of research is Epidemic Routing [20]. Whenever a vehicle encounters another vehicle on the road, it uses V2V communications to send all the data contents that are not available on the other vehicle. In this way, the contents are quickly replicated and spread over the V2V communication network. Spray-and-Wait [21] reduces the traffic load on V2V networks by limiting the number of replicas each vehicle creates per content. Some of the recent works employ coding techniques to improve the efficiency of content delivery. Jain et al. [22] and Wang et al. [23] use erasure coding to encode data contents on a content distribution server with the aim of reducing content delivery latency. Lee et al. [24] additionally re-encode multiple coded chunks on intermediate network nodes while they are delivered from the content distribution server to receiving vehicles (*i.e.*, network coding). The state-of-the-art solutions above have vehicles proactively copy data contents in V2V communication networks. As a result, many vehicles in the network will keep replicas of the same data content. This distribution model works well if the size of data contents is moderate, and a majority of vehicles in the V2V communication network are interested in the same contents. Otherwise, vehicles end up keeping large data contents that they are not interested in, leading to suboptimal usage of their limited data storage space. The latter is often the case for some classes of automotive data contents - for example, a software update package may have large data volume and can be targeted for just a certain model of vehicles. A request-based distribution model [2], [25] would be a better fit for such large and/or targeted data contents. A content distribution server elects a limited number of vehicles as relay vehicles, and send data contents to the selected vehicles. Unlike the push-based content distribution model, relay vehicles do not proactively replicate data contents to other vehicles. Instead, it is the receiver vehicle's responsibility to request interested contents when encountering any relay vehicle. Since the data contents are cached only on a limited number of relay vehicles, it can save communication and data storage overhead for content distribution. The most relevant to our work is a study done by Sathiamoorthy et al. [2]. They analyze vehicle trace data from a taxi fleet in Beijing and a bus fleet in Chicago, and show that erasure coding at a content distribution server significantly reduces the content delivery latency. It is worth noting, though, that the mobility characteristics (*e.g.*, the number and length of trips) of public transport vehicles are significantly different from privately owned vehicles. Despite the promising results with the commercial vehicles, it is still unclear if the similar mechanism is also applicable for privately owned vehicles, which constitute the majority of road traffic.

The contribution of this paper can be summarized in two-fold:

- Leveraging the realistic city-scale road traffic simulations by the NUMo [11], we conduct a series of network simulations to unveil the frequency and duration of contacts between vehicles on the road. These are the key performance factors characterizing the effective network bandwidth of vehicular DTNs, hence play a fundamental role in understanding the feasibility of data offloading over V2V communication networks.
- In addition, we simulate content distribution over vehicular DTNs to unveil the networking latency under the mixed traffic with both public transport and privately owned vehicles. To the best of our knowledge, we are the first to identify the feasibility of data offloading over vehicular DTNs in such a city-scale mixed traffic.

III. NAGOYA URBAN MOBILITY SCENARIO

This section highlights the key characteristics of the NUMo scenario, which will be used as the basis of our vehicular DTN simulations. Note that the design and implementation of the NUMo scenario are beyond the primary scope of this paper. Interested readers are referred to our prior conference publication [11], which discusses the detailed design considerations made in the scenario development process.

A. SCENARIO OVERVIEW

NUMo is a road traffic simulation scenario, designed for the microscopic road traffic simulator SUMO. Its road network model consists of 11,520 road segments, covering most of the highways and major arterial roads across the whole city of Nagoya. The road traffic volume is thoroughly calibrated against public traffic volume statistics on 1,618 road segments across the city [15], while traffic light cycles are also aligned with the historical data [16] wherever possible. The scenario embraces more than 1.6 million vehicle trips per day, making it one of the largest SUMO traffic scenarios open to the research community.

B. SCENARIO CHARACTERISTICS

Fig. 3 visualizes the number of vehicles passing each road segment for 1-hour windows, starting at 5 am, 7 am and 12 pm. The heatmap indicates that the traffic volume is highly non-uniform across different parts of the municipality, and the

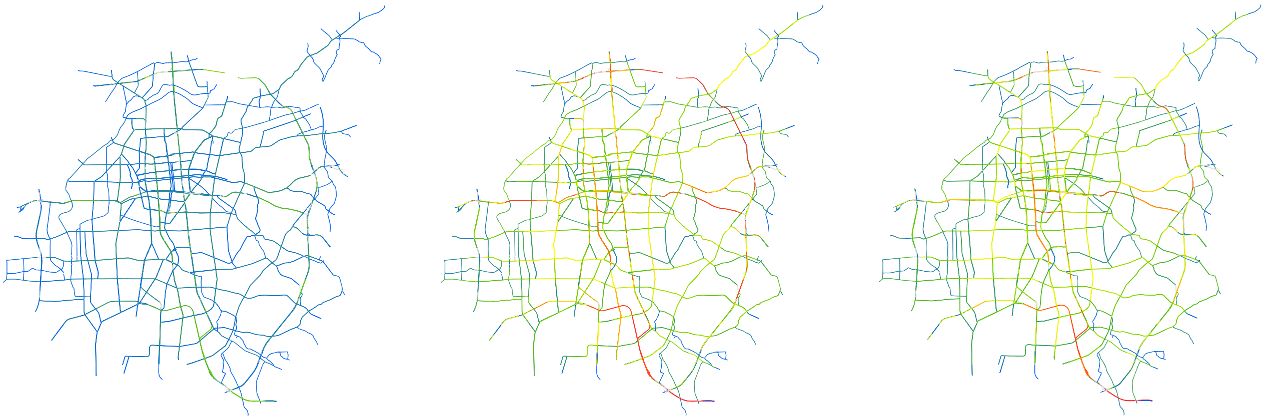

(a) 5-6 am (b) 7-8 am (c) 12-1 pm

FIGURE 3. **Hourly traffic volume simulated by NUMo (blue:≈0 veh/h, green: ≈500 veh/h, yellow: ≈1,000 veh/h, orange: ≈1,500 veh/h, red: > 2,000 veh/h).**

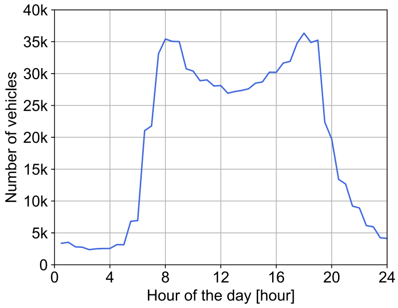

FIGURE 4. **Number of vehicles (Nagoya).**

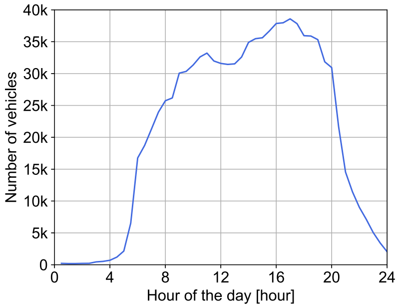

FIGURE 6. **Number of vehicles (Berlin [26]).**

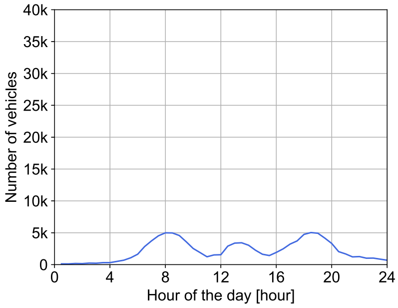

FIGURE 5. **Number of vehicles (Luxembourg [3]).**

traffic tends to be heavier at the city center as well as the major roads along the perimeter of the city.

Fig. 4 shows the total number of vehicles running in the simulation at each time of the day. For comparison, we have also run road traffic simulations with a couple of other open SUMO traffic scenarios, LuST [3] and BeST [26], each modeling the vehicle traffic in Luxembourg and Berlin, respectively. Figs. 5 and 6 show the number of vehicles simulated by these two existing scenarios. The NUMo scenario exhibits clear peaks in the morning and evening rush hours, with more than 35 thousand vehicles traveling in the city. While the commuter traffic fades away during the offpeak hours around noon, the gross traffic volume remains 75% of the peak volume because of the non-commuter traffic such as logistics vehicles. These are in contrast to the other two SUMO traffic scenarios: the traffic volume in the LuST significantly drops in offpeak hours with the minimal volume being less than 40% of the peak traffic, whereas the BeST scenario does not exhibit clear peaks in the traffic volume for morning and evening rush hours. These are partly because both scenarios mainly focus on commuter traffic within the cities of interest. The through traffic (*i.e.*, the trips that start and/or end outside the city) and non-commuter traffic (*i.e.*, commercial vehicles, tourist traffic, etc.) are not considered in the traffic demand modeling. The NUMo addresses these limitations by calibrating the traffic demand based upon the actual hourly vehicle counts on major road segments, preventing the simulated traffic volume from being underestimated.

Last but not least, Fig. 7 shows the cumulative distribution of the relative errors in the traffic volume. We counted the hourly traffic volume at 1,618 observation points in the simulation, and calculated the relative error with respect to the ground-truth in the Traffic Census dataset [15]. The median error in the edge-wise traffic volume was 22%, giving a good approximation of the real road traffic patterns.

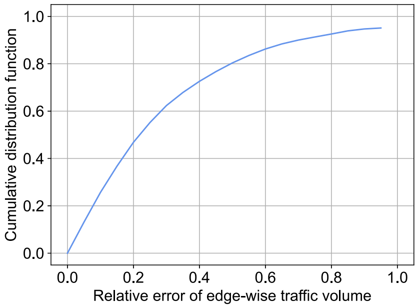

FIGURE 7. Errors in edge-wise traffic volume.

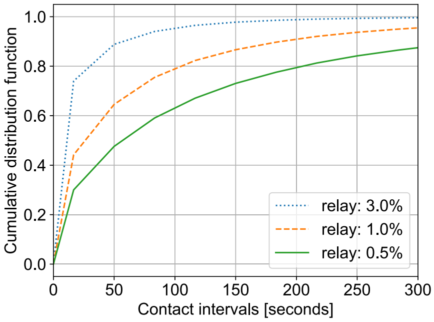

FIGURE 8. CDF of contact intervals.

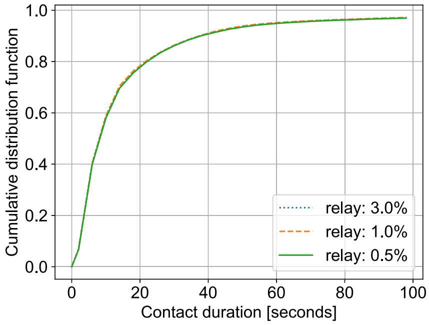

FIGURE 9. CDF of contact duration.

IV. CHARACTERIZING V2V COMMUNICATION PERFORMANCE

Leveraging the realistic road traffic simulation by the NUMo scenario, we conduct a series of network simulations to benchmark the fundamental characteristics of V2V communication networks in Nagoya. The performance of V2V communications can be characterized by a couple of key metrics: *contact intervals* and *contact duration*. A contact between a pair of vehicles starts when they come within the V2V communication range of each other, and ends when they go beyond the communication range. A contact interval is the time intervals between two consecutive contacts, while the contact duration is the metric signifying the length in seconds of each contact. The contact intervals indicate the frequency of data communication opportunities between vehicles, whereas the volume of data that can be transferred in each contact is largely dependent on the contact duration.

The SUMO simulator provides an Application Programming Interface called TraCI, which allows external programs to read the positions and status of simulated vehicles as well as controlling their mobility. We use the TraCI to couple the SUMO simulator with a custom-built network simulator to simulate wireless communications among vehicles.

Because of the scale of the simulations, we opt for a simplifying assumption that vehicles have a constant communication range (*i.e.*, unit disk model). Although it results in omitting physical layer details such as signal attenuation by obstacles and multipath fading, it in return enables significant reduction of computational overhead for network simulations. This allows V2V networking simulations, involving the whole city, which is orders of magnitude larger in size than the typical scale of full-stack V2X network simulations.

A small ratio of vehicles traveling in the city serve as relay vehicles that cache the data chunks in on-board data storage and distribute them to other vehicles over V2V communications. We varied the ratio of relay vehicles over all the vehicles on the road from 0.5% to 3% to unveil the performance characteristics of V2V communications in different phases of deployment. Other vehicles that are not designated as relay vehicles can be potential receivers, which request data chunks when encountering the relay vehicles on the road. We run a series of simulations with different participation rates, and investigate the frequency and duration of contacts between the relay vehicles and receiver vehicles throughout the 24-hour period.

Fig. 8 shows the cumulative distribution function of contact intervals among vehicles. The contact intervals tend to become shorter as a more number of relay vehicles are involved in the content distribution. The median contact interval is 72 seconds when the ratio of relay vehicles is 0.5%, while it decreases to 41 seconds with the ratio of 1.0% and 13 seconds for 3%. The results show that the volume of relay vehicles makes significant impact on the frequency of communication opportunities, while contacts between vehicles can still be observed at about 1 minute intervals even with the smallest participation rate of 0.5%.

Fig. 9 shows the cumulative distribution function of contact duration between relay and receiver vehicles. The result shows that a majority of the contacts last less than 10 seconds. This means that a limited amount of data can be transferred during each single contact, hence receiver vehicles must encounter multiple relay vehicles to complete downloads of large data contents. Note, however, that some of the contacts continue for more than a few tens of seconds. This typically happens when a pair of vehicles both stop at a traffic light, or when they travel in the same direction along the same street. If relay vehicles can allocate larger cache storage space for each data content, the relay vehicles can transfer more data chunks in such occasional long-lasting contacts, which helps further reduce the content distribution latency.

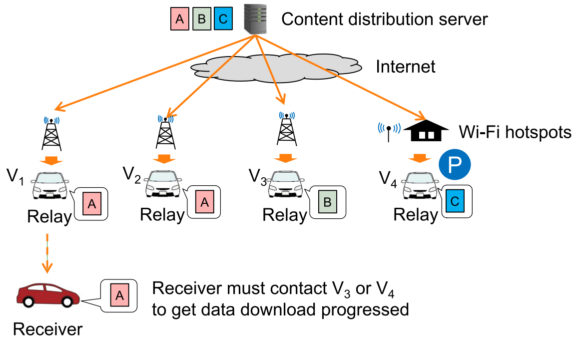

FIGURE 10. **Example benefit of data encoding.**

V. V2V-ASSISTED CONTENT DISTRIBUTION IN NAGOYA

A. SYSTEM MODEL

We further conduct an extensive simulation study using the NUMo scenario to evaluate the performance of V2V-assisted content distribution in a mixed urban traffic in Nagoya. Our system model follows the basic architecture of V2V-assisted content distribution, illustrated in Fig. 1. A content distribution server delivers data content(s) to receiver vehicles via V number of relay vehicles. When delivering a content with the size of L bytes, the server divides the original content into k number of data chunks with the size of $l = L/k$ bytes each. These data chunks are sent to the relay vehicles over cellular and/or WLAN networks. Each relay vehicle allocates the storage space of m bytes ($m \leq L$) per content to save the data chunks received from the content distribution server. If individual relay vehicles cannot save all of the data chunks, a different subset of chunks may be distributed to multiple relay vehicles. In this case, a receiver vehicle must contact more than one relay vehicles to collect all the data chunks required to reconstruct the original data content. Note that the chunk size l is determined such that each chunk becomes small enough to be sent over a V2V communication link during a typical contact duration of a pair of vehicles.

The content distribution server then uses cellular networks to send receiver vehicles the identifiers of the data chunks to download, along with a deadline to complete data download. When a receiver vehicle encounters a relay vehicle, it requests missing data chunks that are available on the relay vehicle's data storage. For a receiver vehicle to reconstruct the original data content, it must download all the k data chunks constituting the original content. In case the required number of data chunks have not been received by the deadline, the remaining chunks will be downloaded directly from the content distribution server by way of cellular networks.

A challenge in this scheme is that the progress of data download tends to slow down as a receiver vehicle obtains more chunks of the content of interest. Let us consider a simple example scenario in Fig. 10, where an original data content is divided into three data chunks A, B, and C. To reconstruct the content of interest, a receiver vehicle must collect all the three chunks through a series of contacts with multiple relay vehicles. Initially, the receiver vehicle does not have any data chunk of the content of interest. Therefore, any relay vehicle V_1 through V_4 can offer a new chunk to this receiver, helping the progress of data download. Once the receiver obtains chunk A from V_1, however, a contact with V_2 no longer helps download a new chunk, hence the receiver must meet either V_3 or V_4 to get it progressed. Thus, the chance of finding a missing piece from V2V networks becomes narrower as the receiver collects more data chunks.

This problem can be addressed by an erasure coding technique [2]. The content distribution server encodes the data chunks before sending them to relay vehicles. Like in the uncoded case, the original content is divided into k data chunks. Then the server converts them to n ($n \gg k$) coded chunks by an erasure code. The coding scheme is designed in a way that the original data content can be reconstructed from any combination of k (out of n) coded chunks. By setting the parameter n to a number much larger than k, the content distribution server can generate non-overlapping sets of coded chunks for every relay vehicle. This allows a receiver vehicle to obtain new coded chunks upon a contact with any relay vehicle, removing the need to find a vehicle caching particular data chunks. Unless otherwise noted, we enable erasure coding in the following simulation study, but we also compare the performance with the uncoded case to identify the quantitative benefit of data encoding.

The existing research on V2V communications have usually assumed the IEEE 802.11p or C-V2X radios in the ITS bands (i.e., the U.S. and Europe allocate a 30 MHz bandwidth in the 5.9 GHz frequency band, while Japan uses a 10 MHz bandwidth in the 760 MHz band). However, the limitation in the channel bandwidth makes it challenging to transfer large data files which we assume as a main target of data offloading. Therefore, we rather opt for regular WLANs in the unlicensed bands in our simulation study. More specifically, we assume that vehicles use an IEEE 802.11n radio for V2V communications. We choose the IEEE 802.11n partly because it is operable on the 2.4 GHz band. In many regions, the outdoor use of WLANs on the 5 GHz and 6 GHz bands are often more restrictive than the 2.4 GHz band. Also, it has extensive performance measurements with vehicle-mounted radios done in the literature [27], which can be used to calibrate the simulation parameters. We assume that relay vehicles serve as WLAN access points, while receiver vehicles connect to them as WLAN clients. Based on the typical data rate of vehicle-mounted IEEE 802.11n radios, reported by Meireles et al. [27], we set the communication range to 100 meters and the data rate to 50 Mbps. A WLAN client can connect to up to one access point at a time, and link setup between a vehicular WLAN access point and a vehicular WLAN client is assumed to take 1 s to complete in the default parameter configuration. The default link setup time is determined based on the real-world measurement campaign reported in the prior art in the literature [28], [29].

B. SIMULATION SCENARIO

In the simulation study, we randomly select a designated ratio of vehicles from the NUMo scenario and elect them as relay vehicles. We varied the participation rate from 1% to 5% to investigate the sensitivity to the volume of vehicles involved in the V2V-assisted content distribution. The content size L is set to either 500 MBytes or 1 GByte in the default configurations, while we also tested the performance varying the content size from 500 MBytes to 3 GBytes. Without loss of generality, we assume that relay vehicles cache data chunks originating from a single data content at a time. If the content distribution server concurrently distributes multiple data contents to different sets of receiver vehicles, data chunks from different contents may coexist on their on-board data storage. We assume that data contents must be delivered to receiver vehicles within five days. Part of the content that has not been downloaded by the deadline must be downloaded directly from the cloud-based content distribution server.

A content distribution server divides an original data content into k data chunks. The number k of data chunks is determined in a way that the size of each chunk will be small enough to be transferred during a typical contact duration between a pair of vehicles. We set the chunk size to 6.25 MBytes, which is the amount of data that can be transferred in a second at the data rate of 50 Mbps. The parameter k is determined by dividing the size of the original data content by the chunk size. Every time a relay vehicle starts a new trip, the content distribution server assigns a new set of data chunks. The relay vehicle then keeps them in the local data storage and forwards to receiver vehicles until the end of the current trip.

For uncoded content delivery, relay vehicles cache overlapping subsets of the k data chunks. Each relay vehicle assigns the variable amount of on-board data storage space to cache data chunks of the content. The content distribution server distributes the data chunks to relay vehicles in a round-robin manner to fill the relay vehicles storage space allocated to this content. The coded content delivery uses the same amount of cache storage space as the uncoded case. The key difference is that the content distribution server uses an erasure code to generate m/l number of coded chunks for each relay vehicle.

The NUMo scenario simulates urban road traffic for a 24-hour period, whereas it often takes multiple days for a receiver vehicle to complete downloading a large data content over V2V communications. To fill the gap, we repeat the NUMo simulations 5 times with different random seeds. This enables a V2V network simulation for 5 consecutive weekdays, allowing us to observe the long-term progress of content distribution. Another aspect to note is that each receiver vehicle may make multiple trips a day. The NUMo scenario, however, does not model repeated trips by the same vehicle, hence different trips are associated with different vehicle IDs. To track data download progress across multiple trips on the same day, we added 300 extra vehicles in the NUMo scenario that follows typical travel patterns of commuters as illustrated in Fig. 11. These extra vehicles are designated as receivers, and request data chunks when contacting relay vehicles on the road. A receiver vehicle makes two trips a day between home and work. Since trip length is an important factor affecting the volume of data a receiver can download over V2V communications, the length distribution of the receiver trips should approximate the reality to ensure the plausibility of simulation results. To achieve this, we generated 10,000 candidate trips from random origins to random destinations, and calculated the length of a route connecting the two locations. Then we sampled 300 of the candidate trips such that the distribution of their length matches the one in NUMo. The origin of the sampled trip is considered the home location of this vehicle, while the destination of the sampled trip is considered the work location. For each receiver vehicle, the same pair of home / work locations are used throughout the 5-day duration of network simulations. A receiver vehicle leaves home during the morning rush hour. The mean departure time of a vehicle is sampled from a uniform distribution between 7-9am, while departure time on each day has small variations, following a Gaussian distribution with the standard deviation of 30 minutes. Likewise, the mean departure time of a return trip from work to home follows a uniform distribution ranging from 5-7pm, while daily fluctuations in the departure time is modeled by a Gaussian distribution with the 30-minute standard deviation. We assume that vehicle on-board computers are turned off while being parked, hence parked vehicles cannot communicate with other vehicles passing by. While receiver vehicles add 600 trips a day in the simulation, it accounts for just 0.04% of the 1.6 million existing trips in the NUMo scenario, hence its impact on the overall road traffic would be negligible.

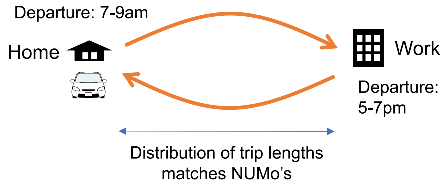

FIGURE 11. Travel patterns of receiver vehicles.

All the simulations (including road traffic simulations) are run on a Linux machine with a 6-core 12-thread CPU @ 4.4 GHz clock frequency, 64 GB DDR4 RAM and 500 GB SSD connected by a PCIe Gen3 interface. While the simulation speeds depend on a variety of factors (*e.g.*, the ratio of relay vehicles, etc.), all the scenarios we present in this section run faster than the wall-clock time.

C. PERFORMANCE METRICS

Throughout the simulation study in this section, we use *offloading rates* as a key performance indicator. Let $\delta_{i,j}$ be a binary variable that takes 1 if a receiver vehicle j has received a data chunk i by the content delivery deadline (*i.e.*, within five days in this simulation study). Otherwise, $\delta_{i,j}$ becomes 0.

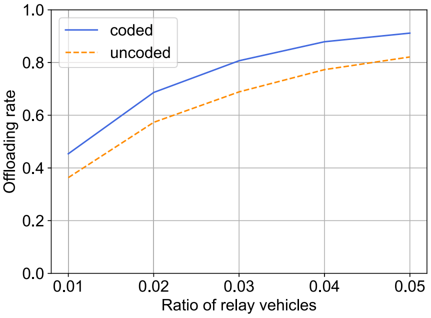

FIGURE 12. **Impact of participation rates**

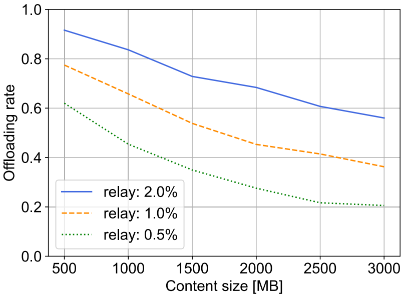

FIGURE 13. **Impact of content size.**

The offloading rate f is defined by:

$$f = \frac{\sum_{i \in \mathcal{C}} \sum_{j \in \mathcal{R}} \delta_{i,j}}{k|\mathcal{R}|} \tag{1}$$

where $\mathcal{C}$ is a set of all the data chunks, while $\mathcal{R}$ denotes a set of all the receiver vehicles. Note that a receiver vehicle stops receiving once it obtains k (coded) data chunks, required to reconstruct the original data content. Hence $\sum_{i \in \mathcal{C}} \delta_{i,j} \leq k$ for any receiver $j \in \mathcal{R}$. The offloading rate f becomes 1 when all the receiver vehicles have received k unique (coded) data chunks by the deadline.

In general, data offloading over vehicular DTNs have both pros and cons, and higher offloading rates may not necessarily mean better communication performance. On the positive side, higher data offloading rates would help mitigate the peak traffic load on cellular networks, which in turn reduces the hardware resources, communication bandwidth and energy budget to be provisioned in the network infrastructure. On the negative side, the content delivery over vehicular DTNs often takes much longer latency than direct distribution from content distribution servers. However, as discussed in Section I, some types of automotive data contents (e.g., software update packages) are not as latency-critical as other types of data (e.g., video streaming), hence need not be delivered within a short period of time in the order of seconds or minutes. For such delay-tolerant network traffic, the benefit of higher offloading rates would typically surpass the downside (i.e., delivery latency), as long as the content is delivered within a designated deadline. In case part of the data content has not been downloaded by the deadline, the remaining part shall be downloaded directly from the content distribution server. This fallback mechanism helps prevent the longer content delivery latency from making any negative impact on applications.

D. SIMULATION RESULTS

1) RATIO OF RELAY VEHICLES AND BENEFIT OF CODING

Fig. 12 shows the offloading rates when the ratio of relay vehicles contributing to the V2V-assisted content distribution is varied from 1% to 5%. To evaluate the benefit of erasure coding, we compare the performance under two different configurations: the *coded* scenario encodes data chunks on the content distribution server before sending them to relay vehicles, while the *uncoded* scenario disables this feature. We set the content size to 1 GByte and the cache size to 100 MBytes per relay vehicle per content for this set of experiments. The offloading rate improves as a more number of relay vehicles are involved in the content distribution. This is a natural consequence, as receiver vehicles have more frequent contacts with relay vehicles, getting plenty of opportunities to obtain data chunks over V2V communications. The offloading rate eventually reaches 91% when 5% of vehicles traveling in the city serve as relay vehicles and coding is enabled. It is also worth noting that as much as 45% (*i.e.*, 450 MBytes) of the data chunks can be obtained via V2V communications even if the ratio of relay vehicles is as low as 1%. This is an encouraging observation, as it indicates V2V-assisted content distribution starts bringing offloading benefits even in the early stage of deployment in the market. The results also indicate that coding on the content distribution server consistently improves the offloading rates by about 10%. Since both the coded and uncoded scenarios use almost the same amount of cache storage space on relay vehicles (aside from the small headers to be attached to coded chunks), this clearly shows the benefit of data coding in the V2V-assisted content distribution.

2) SIZE OF DATA CONTENT

Another important aspect is the size of data content that can be distributed over V2V communication networks within a reasonable amount of latency. We varied the content size from 500 MBytes to 3 GBytes, and evaluated the offloading rates at the end of the 5-day simulation period. The performance was tested with various ratios of relay vehicles, ranging from 0.5% to 2.0%, to clarify the offloading benefits expected in different stages of deployment. In this set of experiments, we assume that relay vehicles can allocate enough storage space to cache the entire data content. This allows us to rule out the impact of cache capacity limits, which will be separately investigated in the following section. Fig. 13 shows the offloading rates for the varying content size. If the content is less than 1 GByte in data size, a majority of data chunks can be collected over V2V communications within the 5-day time window even if the ratio of relay vehicles is limited to no more than 1%. While the larger data contents tend to take more time to complete distribution, at least part of the data chunks

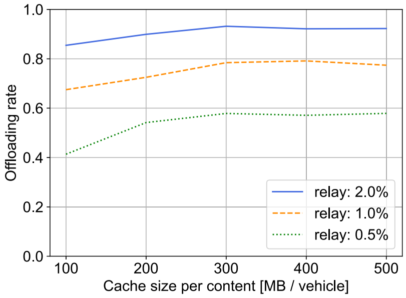

FIGURE 14. **Impact of cache size**

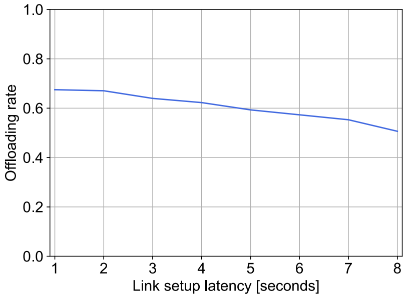

FIGURE 15. **Impact of link setup latency**

can be obtained over V2V communications. If the content distribution is not completed within the latency budget set by the content distribution server, the remaining chunks may be directly distributed from the server to individual receiver vehicles via cellular networks to avoid service disruptions.

3) SIZE OF CACHE STORAGE CAPACITY ON RELAY VEHICLES

Vehicles often have a limited data storage capacity on their on-board computing units. Hence the cache storage space on relay vehicles shall be saved as much as possible, while the shortage of cache storage space may limit the amount of data that can be transferred in long-lasting contacts between vehicles. Fig. 14 shows the offloading rates with the varying size of cache storage space allocated by each relay vehicle per content. The content size is set to 500 MBytes in this set of experiments, while we vary the ratio of relay vehicles from 0.5% to 2.0% like in the experiments in the previous section. The offloading rates tend to improve when the cache size is less than 300 MBytes (*i.e.*, 60% of the content size). The larger cache size helps relay vehicles to keep sending data chunks during long contacts with receiver vehicles (*e.g.*, when stopping at traffic lights) without running out of the cached data to be sent. The performance benefit from the increased cache size is more significant when the ratio of relay vehicles on the road is limited, where the data transfer during occasional long-lasting contacts play an important role in supplementing the fewer frequency of contacts. The offloading rates saturate once the cache size exceeds 300 MBytes. This happens because most of the contacts are not long enough to transfer more than 300 MBytes of data under the road traffic condition in Nagoya. This indicates that the cache size of 300 Mbytes would give the best trade off between savings of data storage resources and the latency of content distribution over V2V communication networks. Note, however, that the optimal value would depend on the characteristics of road traffic as well as the performance of the underlying V2V communication technology, calling for careful parameter tuning through simulation studies.

4) LINK SETUP LATENCY

The link setup over wireless LAN usually takes a certain amount of latency for discovery of access points, authentication and IP address assignment [29]. In the V2V-assisted content distribution where vehicles transfer data in short duration of contacts, the link setup latency may make a critical impact on the overall content distribution latency. To quantify the impact, we evaluate the offloading rates while varying the link set up latency from 1 second to 8 seconds. We set the content size to 500 MBytes, the cache size to 100 MBytes and the ratio of relay vehicles to 1.0 % in this set of experiments. The results in Fig. 15 indicate that the offloading rates linearly degrades as the link setup latency increases. This is a natural consequence of losing the effective duration of V2V communications at the beginning of every contact between relay and receiver vehicles. The prior art in the literature [29] showcases that the link setup time can be reduced to less than a second by carefully tuning the WLAN system parameters and/or leveraging the Fast Initial Link Setup feature defined in the latest IEEE 802.11 standard. The countermeasures to limit the link setup latency is vital to maximize the offloading benefits.

VI. CONCLUSION

In this paper, we have investigated the feasibility of V2V-assisted cellular data offloading through city-scale network simulations in Nagoya. Leveraging the NUMo scenario, modeling the realistic vehicle traffic in the city, we have conducted extensive network simulations to test the data offloading performance under various combinations of parameters, including the ratios of relay vehicles, the size of data contents, the capacity of cache storage, and link setup latency. The results indicate that vehicular DTNs hold promise to offload large amount data traffic from cellular networks, even in the early stage of deployment with a limited volume of relay vehicles on the road.

As the next step of this project, we are preparing a proof-of-concept implementation of the content distribution over vehicular DTNs, aiming to conduct small-scale experiments using real wireless devices mounted on vehicles on the move. The results from the PoC experiments will be compared against the simulation results to validate the feasibility of the concept as well as the simulation models used. Some part of the network stack (*e.g.*, transport layer protocols) may be optimized for V2V communications to achieve further gain in the offloading performance. These potential improvements shall also be covered in future work.

REFERENCES

[1] "Automotive edge computing consortium, general principle and vision white paper version 4.0.2," 2023. [Online]. https://aecc.org/resources/publications/

[2] M. Sathiamoorthy, A. G. Dimakis, B. Krishnamachari, and F. Bai, "Distributed storage codes reduce latency in vehicular networks," in *Proc. IEEE 31st Annu. Int. Conf. Comput. Commun.*, 2012, pp. 2646–2650.

[3] L. Codeca, R. Frank, S. Faye, and T. Engel, "Luxembourg SUMO traffic (LuST) scenario: Traffic demand evaluation," *IEEE Intell. Transp. Syst. Mag.*, vol. 9, no. 2, pp. 52–63, Summer 2017.

[4] L. Codecá and J. Härri, "Towards multimodal mobility simulation of C-ITS: The Monaco SUMO traffic scenario," in *Proc. IEEE Veh. Netw. Conf.*, 2017, pp. 97–100.

[5] S. C. Lobo, S. Neumeier, E. M. G. Fernández, and C. Facchi, "InTAS - The ingolstadt traffic scenario for SUMO," in *Proc. SUMO User Conf. Proc.*, 2020, pp. 73–92.

[6] Y. Yamazaki, Y. Tamura, X. Défago, E. Javanmardi, and M. Tsukada, "ToST: Tokyo SUMO traffic scenario," in *Proc. IEEE 26th Int. Conf. Intell. Transp. Syst.*, 2023, pp. 3597–3604.

[7] D. Förster, H. Löhr, A. Grätz, J. Petit, and F. Kargl, "An evaluation of pseudonym changes for vehicular networks in large-scale, realistic traffic scenarios," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 10, pp. 3400–3405, Oct. 2018.

[8] M. Guériau and I. Dusparic, "Quantifying the impact of connected and autonomous vehicles on traffic efficiency and safety in mixed traffic," in *Proc. IEEE 23rd Int. Conf. Intell. Transp. Syst.*, 2020, pp. 1–8.

[9] K. Kušić, R. Schumann, and E. Ivanjko, "Building a motorway digital twin in SUMO: Real-time simulation of continuous data stream from traffic counters," in *Proc. IEEE Int. Symp.*, 2022, pp. 71–76.

[10] K. Kušić, R. Schumann, and E. Ivanjko, "A digital twin in transportation: Real-time synergy of traffic data streams and simulation for virtualizing motorway dynamics," *Adv. Eng. Inform.*, vol. 55, 2023, Art. no. 101858.

[11] T. Higuchi, L. Zhong, and R. Onishi, "NUMo: Nagoya urban mobility scenario for city-scale V2X simulations," in *Proc. IEEE 2024 Veh. Netw. Conf.*, 2024, pp. 17–24.

[12] C. Sommer, R. German, and F. Dressler, "Bidirectionally coupled network and road traffic simulation for improved IVC analysis," *IEEE Trans. Mobile Comput.*, vol. 10, no. 1, pp. 3–15, 2011.

[13] R. Riebl, H.-J. Günther, C. Facchi, and L. Wolf, "Artery: Extending veins for VANET applications," in *Proc. IEEE 2015 Int. Conf. Models Technol. Intell. Transp. Syst.*, 2015, pp. 450–456.

[14] M. Segata, S. Joerer, B. Bloessl, C. Sommer, F. Dressler, and R. L. Cigno, "Plexe: A platooning extension for veins," in *Proc. IEEE 2014 Veh. Netw. Conf.*, 2014, pp. 53–60.

[15] Japanese Ministry of Land, Infrastructure, Transport and Tourism (MLIT), "Nation-wide road and street traffic survey - general road traffic volume survey," (in Japanese), 2023. [Online]. Available: https://www.mlit.go.jp/road/census/r3/

[16] Japan Road Traffic Inf. Center (JARTIC), "Open data," (in Japanese), 2023. [Online]. Available: https://www.jartic.or.jp/service/opendata/

[17] K. Ota, M. Dong, S. Chang, and H. Zhu, "MMCD: Max-throughput and min-delay cooperative downloading for drive-thru internet systems," in *Proc. IEEE Int. Conf. Commun.*, 2014, pp. 83–87.

[18] F. Malandrino, C. Casetti, C.-F. Chiasserini, and M. Fiore, "Offloading cellular networks through ITS content download," in *Proc. IEEE 9th Annu. Commun. Soc. Conf. Sensor Mesh Ad Hoc Commun. Netw.*, 2012, pp. 263–271.

[19] X. Shen, X. Cheng, L. Yang, R. Zhang, and B. Jiao, "Data dissemination in VANETs: A scheduling approach," *IEEE Trans. Intell. Transp. Syst.*, vol. 15, no. 5, pp. 2213–2223, Oct. 2014.

[20] A. Vahdat and D. Becker, "Epidemic routing for partially-connected ad hoc networks," Duke Univ., Tech. Rep., 2000.

[21] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Spray and wait: An efficient routing scheme for intermittently connected mobile networks," in *Proc. ACM SIGCOMM Workshop Delay-Tolerant Netw.*, 2005, pp. 252–259.

[22] S. Jain, M. Demmer, R. Patra, and K. Fall, "Using Redundancy to cope with failures in a delay tolerant network," in *Proc. Conf. Appl., Technol., Architectures, Aprotocols Comput. Communi.*, 2005, pp. 109–120.

[23] Y. Wang, S. Jain, M. Martonosi, and K. Fall, "Erasure-coding based routing for opportunistic networks," in *Proc. ACM SIGCOMM Workshop Delay-Tolerant Netw.*, 2005, pp. 229–236.

[24] U. Lee, J.-S. Park, J. Yeh, G. Pau, and M. Gerla, "Code torrent: Content distribution using network coding in VANET," in *Proc. 1st Int. Workshop Decentralized Resource Sharing Mobile Comput. Netw.*, 2006, pp. 1–5.

[25] S. Kapadia, B. Krishnamachari, and S. Ghandeharizadeh, "Static replication strategies for content availability in vehicular ad-hoc networks," *Mobile Netw. Appl.*, vol. 14, no. 5, pp. 590–610, 2009.

[26] K. Schrab, R. Protzmann, and I. Radusch, "A Large-scale traffic scenario of berlin for evaluating smart mobility applications," in *Proc. Smart Energy Smart Transport*, 2023, pp. 276–287.

[27] R. Meireles, A. Rodrigues, A. Stanciu, A. Aguiar, and P. Steenkiste, "Exploring Wi-Fi network diversity for vehicle-to-infrastructure communication," in *Proc. IEEE Veh. Netw. Conf.*, 2020, pp. 1–8.

[28] A. Kato et al., "Link setup time reduction by FILS on IEEE 802.11-based inter-vehicular communications," *IEEE Access*, vol. 9, pp. 159796–159808, 2021.

[29] T. Higuchi, S. Ucar, C. Lee, and O. Altintas, "Link setup latency in vehicle-to-vehicle WLANs: A comparative study," in *Proc. IEEE Consum. Commun. Netw. Conf.*, 2024, pp. 765–770.

TAKAMASA HIGUCHI (Member, IEEE) received the B.E., M.E., and Ph.D. degrees from Osaka University, Suita, Japan, in 2010, 2012, and 2014, respectively. He started his professional carrier as an Assistant Professor with Osaka University, where he led several research projects on mobile computing and networking. During his appointment with Osaka University, he was also affiliated with the University of California, Los Angeles, CA, USA, as a Visiting Researcher. In 2016, he joined Toyota and has been leading R&D projects on vehicular communications, vehicular cloud computing, and network simulations. He is currently a Senior Researcher with Toyota Motor Corporation, Tokyo, Japan. He was a TPC Co-Chair of the IEEE Vehicular Networking Conference in 2021 and 2023.

LEI ZHONG (Member, IEEE) received the M.E. degree from Tongji University, Shanghai, China, in 2008, and the Ph.D. degree from The Graduate University for Advances Studies, Hayama, Japan, in 2011. He is a chief Delegate and Chair in several related standardization organizations such as 3GPP, IEEE, Wi-Fi Alliance, and AECC. He is currently a Group Manager and Principal Researcher with Information and Communication Planning Division, Toyota Motor Corporation, Tokyo, Japan, where he is leading the communication network group on the research and standardization for connected vehicles. He has authored or coauthored more than 60 technical papers and applied more than 20 patents in his research areas which include wireless networking, edge computing, machine learning, connected vehicles, Internet-of-Things, and Big Data. He is also the Vice-Chair of IEEE TCGCC Special Interest Group on Green Internet of Vehicles and the Editor of *Elsevier Vehicular Communications* and IEEE OPEN JOURNAL OF THE COMPUTER SOCIETY.

RYOKICHI ONISHI received the B.S. and M.S. degrees in information and communication engineering and the Ph.D. degree in electrical engineering and information systems from The University of Tokyo Bunkyo, Japan. Since 2001, he has been with Toyota InfoTechnology Center and Toyota Motor Corporation, where he is currently a General Manager of InfoTech-IS Deptartment, Information and Communication Planning Division. Hisresearch interests include end-to-end digital infrastructure combining communication, computation, and data storage, which is geared for emerging software-defined user-centric services, intelligent driving system with "greener" generative AI, and remote monitoring and control of automated driving system. He has been awarded 31 patents by the U.S. Patent Office in this domain.