

Cost-optimal V2X Service Placement in Distributed Cloud/Edge Environment

Abdallah Moubayed*, Abdallah Shami*, Parisa Heidari[†], Adel Larabi[‡], and Richard Brunner[†]

* Western University, London, Ontario, Canada; e-mails: {amoubaye, abdallah.shami}@uwo.ca

[†] Ericsson, Montreal, Quebec, Canada; e-mails: {parisa.heidari, richard.brunner}@ericsson.com

[‡] Edge Gravity by Ericsson, Montreal, Quebec, Canada; e-mail: ALarabi@edgegravity.ericsson.com

Abstract—Deploying V2X services has become a challenging task. This is mainly due to the fact that such services have strict latency requirements. To meet these requirements, one potential solution is adopting mobile edge computing (MEC). However, this presents new challenges including how to find a cost efficient placement that meets other requirements such as latency. In this work, the problem of cost-optimal V2X service placement (CO-VSP) in a distributed cloud/edge environment is formulated. Additionally, a cost-focused delay-aware V2X service placement (DA-VSP) heuristic algorithm is proposed. Simulation results show that both CO-VSP model and DA-VSP algorithm guarantee the QoS requirements of all such services and illustrates the trade-off between latency and deployment cost.

Index Terms—Multi-Access Edge Computing (MEC), Cloud Computing, Intelligent Transportation Systems (ITS), V2X Services, V2X Service Placement

I. INTRODUCTION

Vehicle-to-everything (V2X) communication, a cornerstone of intelligent transportation systems (ITSs), has garnered more attention from different entities including governmental agencies and automotive manufacturers due to the various projected merits of such systems. For example, such systems help reduce road accidents, introduce new business opportunities, and reduce the cost of managing vehicular fleets [1]. Additionally, V2X communication helps support new set of services such as traffic optimization and in-car entertainment [2]. However, offering such services is coupled with maintaining stringent performance requirements. More specifically, end-to-end (E2E) latency/delay is an important non-functional requirement for such services. For example, the minimum E2E latency requirement for pre-sense crash warning is 20-50 ms [3], [4].

To that end, mobile edge computing (MEC), the concept of providing computing and storage capabilities near sensors and mobile devices, has been proposed as a viable solution to reduce the E2E latency/delay. MEC has been discussed in the context of ITSs and V2X communication by hosting V2X services and applications on servers near the end-users to reduce the serving latency. As a result, providers can offer new services and applications.

Nonetheless, deploying a distributed computing environment consisting of core and edge nodes introduces fresh challenges [5], [6]. This includes the V2X services/applications' placement decisions. This is because of the limited computing and storage resources available at edge nodes that allow for lower latencies

when hosting such services on them. Another challenge is the cost associated with hosting such services on edge nodes, an issue that is of particular interest to service providers. This is due to the higher operational expenditures needed to maintain the edge nodes as compared to cloud nodes because of the larger number of physical structures to deploy and maintain. Therefore, this work focuses on formulating an optimization model that reduces the cost of hosting V2X services while guaranteeing delay, computational, and placement requirements.

The contribution of this work is modeling the cost-optimal V2X service placement problem in a distributed cloud/edge computing environment. This is formulated as a binary integer programming model. To the best of our knowledge, this is the first work that considers the problem of cost-optimal V2X service placement while satisfying the computational capabilities at the distributed computing nodes.

The organization of this paper is as follows: Section II briefly discusses a few previous research work in this area. Section III describes the system model. Section IV formulates the optimization problem. Section V presents the proposed heuristic. Section VI evaluates the performance and discusses the corresponding results. Section VII concludes the paper.

II. RELATED WORK

Few previous works in the literature have investigated the concept of MEC within the context of V2X communications. A hierarchical MEC architecture was described for adaptive video streaming in which the MEC platform changed the video stream quality according to the observed channel conditions [7]. Similarly, MEC technology was proposed to count and classify vehicles for real-time traffic surveillance [8]. Similarly, few previous literature works have studied the cost of service placement in MEC environment. For example, a novel virtual network function placement strategy that aims at minimizing a requirement, such as the end to end delay or the overall deployment cost was proposed in [9]. However, these works have several limitations as they mainly consider one V2X service or application at a time and investigate the impact of MEC on metrics such as latency while discarding other metrics such as deployment cost and computational capabilities of edge nodes. Although the previous work in [10] did study the problem of V2X service placement in such computing environments, but it only considered the delay requirement without studying the associated cost.

To the best of our knowledge, the problem of cost-optimal placement of V2X services while concurrently considering the availability of computational resources and delay requirements of such services/applications has not been previously investigated. Accordingly, this work formulates the cost-optimal V2X service placement problem in a distributed computing paradigm consisting of core and edge nodes while respecting other non-functional requirements, especially performance and delay of response time. We evaluate the performance of such a system and study the trade-off between the delay and cost of such a placement.

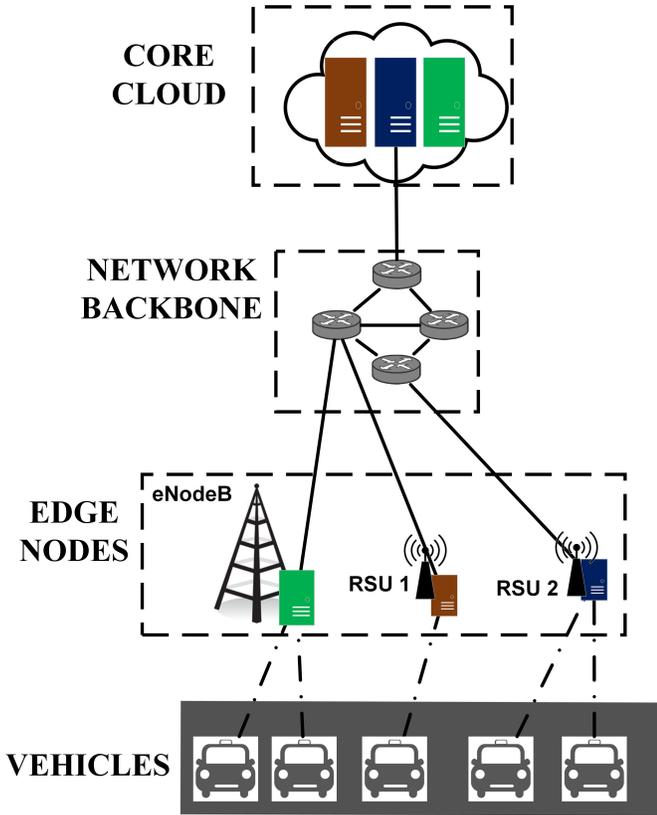


Fig. 1. System Model

III. SYSTEM MODEL

A. System Setup

Similar to [10], a uni-directional multiple-lane highway environment is assumed in this work. Furthermore, it is assumed that the vehicles are moving at constant speed with equal distance separation. This mimics vehicles traveling along a highway which on average travel at constant speed. Moreover, this work assumes that the highway segment is covered with LTE-A using multiple evolved NodeB (eNB) base stations. Additionally, multiple road side units (RSUs) are assumed to be placed on the side of the road along the highway. Each eNB or RSU has a MEC host with defined limited computational resources (CPU, memory, and storage). Accordingly, the network access edge is composed of these

eNBs and RSUs. Also, it is assumed that the network access edge is connected to a core cloud data center containing servers with larger computing capabilities via the network backbone. This is illustrated in Fig. 1. Note that V2X communication is supported by LTE-A through the Uu-based and PC5-based interfaces [11], [12].

B. Description of V2X

Three distinct V2X service types are considered in this work, each representing one use case of different V2X applications. A brief summary of these services and their requirements is given below.

1) Cooperative Awareness Basic Service:

It is defined by the Cooperative Awareness Message (CAM). The CAM message is routinely exchanged between vehicles and network road nodes. It contains information about the vehicles position, movement, and other sensor data [13], [14]. This service has a strict latency requirement given that it is one use case of traffic safety services (between 10-20 ms) [15].

2) Decentralized Environmental Notification Basic Service:

It is defined by the Decentralized Environmental Notification Message (DENM) [16]. The DENM message is a notification message alerting road users of a road-related event (e.g. traffic condition warning). Due to its nature, it is considered a hybrid use case as it belongs to the group of V2X traffic efficiency applications as well as to the group of V2X traffic safety applications. Therefore, it has a less strict latency requirement (tolerates up to 100 ms latency) [15].

3) Media Downloading and Streaming:

It is one of the V2X Infotainment applications' use cases [17]. Such a service provides vehicle users with on-demand information or entertainment through the Internet [17], [18]. Accordingly, this service has the least strict latency requirement (tolerates latencies up to 1 sec) [19], [20].

IV. COST-OPTIMAL V2X SERVICE PLACEMENT (CO-VSP)

This work formulates the cost-optimal placement problem of V2X services in a distributed cloud/edge environment. The goal is to place a set of V2X services in such a manner that minimizes the cost of placing these services on the available computing nodes while adhering to different delay, available computational resources, placement, and cost constraints. This is shown in the CO-VSP model described below.

A. Key Mathematical Notations:

Table I presents the key mathematical notations of this work. Moreover, the decision variable is:

$$X_s^c = \begin{cases} 1, & \text{if V2X service/application } s \text{ is placed on} \\ & \text{computing node } c. \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

B. Problem Formulation:

The formulation of the cost-optimal V2X service placement (CO-VSP) problem in a distributed computing environment made up of core and edge nodes is:

TABLE I
KEY MATHEMATICAL NOTATIONS

Set	Description
S	Set of V2X services' instances to be placed
U	Set of unique V2X services
S_u	Subset of V2X services of type $u \in U$
C	Set of computing nodes (core or edge node) available to host the V2X services
V	Set of vehicles accessing V2X services' instances
$d_{s,v}^c$	Delay experienced by vehicle v served by V2X service instance s if placed at computing node c
D_s^{th}	Delay/latency threshold of V2X service s
R_s^i	Computational resource requirement of V2X service s ($i \in \{CPU, memory, storage\}$)
Cap_c^i	Computational resource i available at computing node c
$cost_s^c$	Cost of placing V2X service s at computing node c
$cost^{th}$	Cost threshold allocated by the service provider

1) Objective Function:

- Equation (2a) aims to minimize the aggregate cost of placing the V2X services instances. The aggregate cost is determined by summing up all the costs associated with placing a service s at computing node c . When the decision variable is set to 1, only the associated cost will be computed as part of the aggregate cost.

$$\min \sum_{s \in S} \sum_{c \in C} X_s^c * cost_s^c \quad (2a)$$

2) Constraints:

- Delay Constraint:** Equation (2b) ensures that the maximum delay/latency threshold of the service is satisfied.

$$\sum_{c \in C} X_s^c \left(\frac{1}{|V|} \sum_{v \in V} d_{s,v}^c \right) \leq D_s^{th}; \forall s \in S \quad (2b)$$

- Computational Resource Availability Constraint:** Equation (2c) ensures that the aggregate computational resource requirement of the placed V2X service instances placed at computing node c does not exceed its available computational resources.

$$\sum_{s \in S} X_s^c R_s^i \leq Cap_c^i; \forall c \in C, \forall i \in \{CPU, memory, storage\} \quad (2c)$$

- Placement Constraint:** Equation (2d) guarantees that each V2X service instance s is only placed at one computing node.

$$\sum_{c \in C} X_s^c = 1; \forall s \in S \quad (2d)$$

- Unique Service Placement Constraint:** Equation (2e) guarantees that instances of unique V2X service of type u are hosted in varying nodes for redundancy purposes.

$$\sum_{s \in S_u} X_s^c \leq 1; \forall c \in C; \forall u \in U \quad (2e)$$

- Cost Constraint:** Equation (2f) specifies that the aggregate cost of placement is less than the maximum cost threshold allocated by the service provider.

$$\sum_{s \in S} \sum_{c \in C} X_s^c * cost_s^c \leq cost^{th}; \quad (2f)$$

Note that this work assumes that the number of vehicles is proportional to the number of instances of each service type. Accordingly, the number of vehicles considered can have a direct impact on the cost since additional V2X service instances would need to be deployed and hosted with the increased number of vehicles.

C. Complexity:

The problem is considered to be NP-complete since it is a binary integer linear programming problem similar to the facility location problem [21], [22]. This is further emphasized when determining the search space of the problem. The search space for this problem is $2^{|C||S|}$. For example, when $|C| = 10$ and $|S| = 5$, the size of the search space becomes 1.125×10^{15} . As such, it becomes computationally expensive to solve such a problem because of the search space's exponential growth. Therefore, it is crucial to develop a low-complexity heuristic algorithm as a viable solution.

V. DELAY-AWARE V2X SERVICE PLACEMENT (DA-VSP) ALGORITHM

The cost metric can be introduced to any general service placement strategy/algorithm. However, we select the placement algorithm of [10] to demonstrate our approach. The work presented in this manuscript proposes a low complexity heuristic, namely "Delay-aware V2X Service Placement" (DA-VSP) algorithm, as a solution to the V2X service placement problem. In other words, this algorithm is an adaptation of our previous G-VSPA heuristic algorithm proposed in [10] to consider the cost of deployment in a placement algorithm. Algorithm 1 provides the pseudo-code of the DA-VSP algorithm.

A. Description:

The proposed DA-VSP algorithm is organized as follows:

- Line 1: A mock variable is defined that determines if a V2X service instance was placed or not.
- Line 2: The unique services are sorted in ascending delay tolerance order.
- Lines 3-16: Algorithm iterates through the instances of each unique service type. For each instance to be hosted, the algorithm determines the computing node with the lowest deployment cost (line 7). If this node adheres to the delay constraint as well as the computational resource constraint, the decision variable is set to 1. Moreover, the

Algorithm 1 Delay-Aware V2X Service Placement (DA-VSP)

Input: $U = \{1, 2, \dots, |U|\}$, $S = \{1, 2, S_1, \dots, |S|\}$
 $C = \{1, 2, \dots, |C|\}$, $V = \{1, 2, \dots, |V|\}$
 $Cost_c = \{Cost_1, Cost_2, \dots, Cost_C\}$

Output: $X_{s_u}^c$, $Agg Cost = \sum_{s \in S} \sum_{c \in C} X_s^c * Cost_c$

- 1: **define** $X_{s_u} = \sum_{c \in C} X_{s_u}^c$
- 2: **set** $U_{sort} = Asc Sort(U)$
- 3: **for** $u \in U_{sort}$ **do**
- 4: **for** $s_u \in S_u$ **do**
- 5: **define** $d_{s_u}^c = \left(\frac{1}{|V|} \sum_{v \in V} d_{s,v}^c \right)$
- 6: **while** $X_{s_u} \neq 1$ **do**
- 7: **find** $cost_{s_u}^c = \min_{c \in C} Cost_c$
- 8: **if** $Rem Cap_c > R_{s_u} \ \& \ d_{s_u}^c < D_{s_u}^{th}$ **then**
- 9: **set** $X_{s_u}^c = 1$
- 10: **update** $Rem Cap_c = Rem Cap_c - R_{s_u}$
- 11: **update** $C = C \setminus c$
- 12: **else**
- 13: **update** $cost_{s_u}^c = \infty$
- 14: **end if**
- 15: **end while**
- 16: **end for**
- 17: **end for**
- 18: **return** $X_{s_u}^c$, $Agg Cost$

available computing resources at node c are updated and the node is removed from the set of available nodes for all other instances of the same type u (lines 8-11). If the node violates any of these constraints, the node is ignored and the algorithm checks the following node with the second lowest deployment cost (line 12-14). This process is repeated until all the service instances are placed.

B. Complexity:

The time complexity of the DA-VSP algorithm is $O(|C||S|)$. This is clear given that the algorithm examines $|C|$ potential nodes when trying to place a service instance s . Accordingly, using the same values for $|C|$ and $|S|$ as before, the algorithm would need around 50 operations to reach convergence.

The low-complexity nature of this algorithm makes it suitable for real-time deployment. Service providers can use it to determine a suitable balance between the anticipated cost and the desired performance by deploying the minimum number of RSUs based on the traffic pattern within a particular area of interest.

VI. PERFORMANCE EVALUATION

A MATLAB-based simulator was developed to evaluate the performance of the proposed CO-VSP model in terms

of deployment cost, average delay/latency, and computational resource utilization. Note that the CO-VSP model is used for benchmarking purposes with any real-life deployment using the proposed DA-VSP heuristic as the placement algorithm.

A. Simulation Setup:

An LTE-based V2X system is considered in this work that supports the three aforementioned services. The number of instances of each service is proportional to the number of vehicles. The delay tolerance threshold is set at 20 ms for the CAM service, 50 ms for the DENM service, and 150 ms for the media downloading and streaming service [15], [19], [20]. Additionally, CAM service is modeled as a medium sized virtual machine (VM) with low computational requirements due to the local update status of this service [23]. Similarly, DENM service is considered to be a large-sized VM [23]. This is because it processes more data about the traffic conditions, which results in a higher computational requirement. Lastly, the media downloading service is considered to be an extra-large VM to support and satisfy the different download requests [23].

For the physical environment, this work considers a highway scenario consisting of 2 lanes, each 2 km long. Moreover, it is assumed that there are 2 core cloud nodes, 3 eNodeBs, and 5 RSUs. Each core cloud node has 32 cores of CPU, 64 GB of memory, and 240 GB of storage. Similarly, each eNB and RSU has 8 cores of CPU, 16 GB of memory, and 240 GB of storage. Moreover, the delays/latencies are considered to be uniformly distributed with Vehicle-to-RSU delay between 1 ms - 10 ms [11], Vehicle-to-eNB delay between 20 ms - 40 ms [11], [24], and Vehicle-to-core delay between 60 ms - 130 ms [24]. The cost of hosting a service instance in the cloud is estimated to be 7.5k\$/month (assuming Amazon's web services are used) [25], while the cost of hosting a service instance at an edge node is estimated to be around 15k\$/month (almost a 100% increase in cost for edge nodes as compared to core cloud nodes) [26], [27]. The maximum threshold cost per month is assumed to be 500k\$/month [28].

B. Results and Discussion:

Fig. 2 shows the average monthly deployment cost while Fig. 3 shows the average delay/latency of the V2X services considered. Despite the fact that the proposed CO-VSP model guarantees the QoS requirements of the different V2X service, the trade-off between deployment cost and delay/latency is clearly evident. The average deployment cost for media services increases while its average delay/latency decreases with the increase in number of vehicles. This is expected since there are more media downloading instances being hosted at the edge with the increasing number of vehicles. In contrast, CAM and DENM services remain stable in terms of deployment cost and latency. This is due to the strict requirements in terms of delay for these services which force the model to always place them at the edge rather than the core cloud nodes. The same observations apply when the DA-VSP heuristic algorithm is considered. Moreover, it can be seen that the DA-VSP algorithm achieves similar performance to that of the CO-VSP model

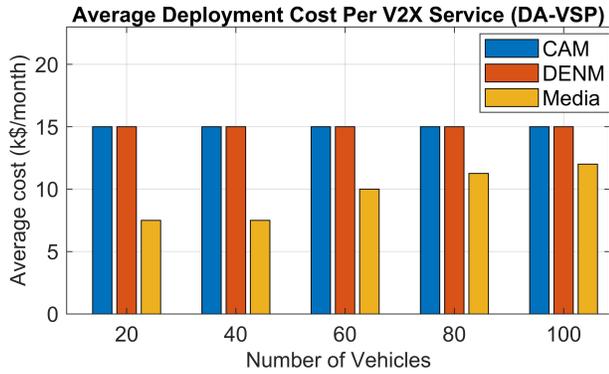
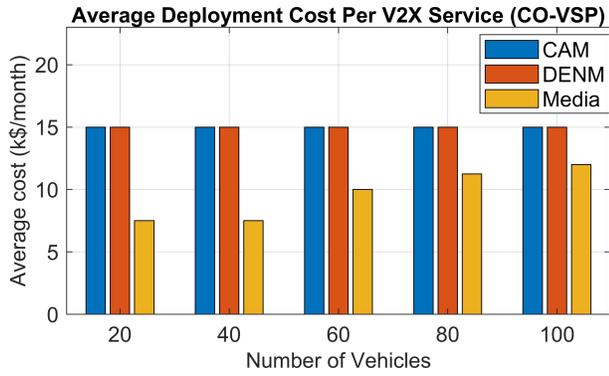


Fig. 2. Different V2X Services' Average Monthly Cost

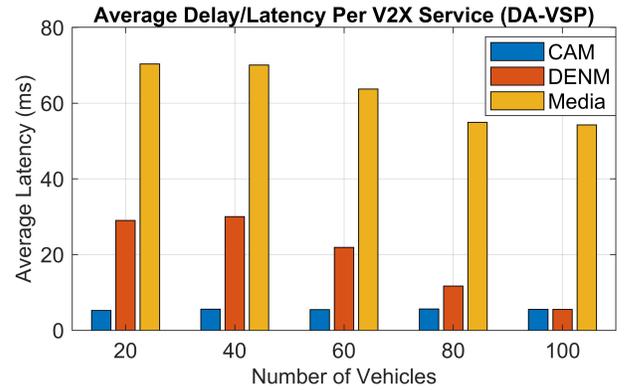
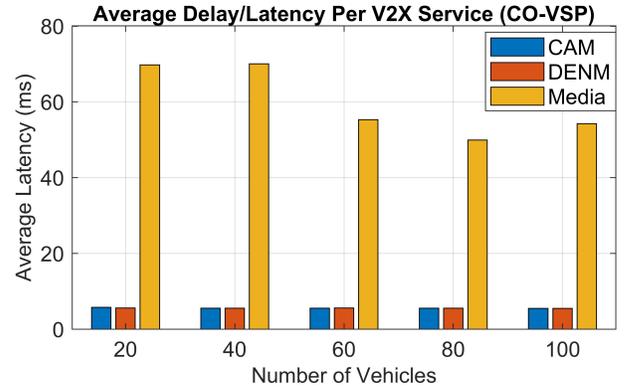


Fig. 3. Different V2X Services' Average Delay/Latency

while having lower computational complexity. This is because it also tries to find the computing node with the least cost that can satisfy the delay constraint.

Fig. 4 illustrates the average CPU utilization in both the cloud and edge nodes. As expected, the CPU utilization at the edge increases with the increased number of vehicles. Again, this is because of the greater number of media downloading service instances that move away from the core cloud and are hosted at the edge. In contrast, the cloud utilization becomes stable due to the limited number of instances it can host to satisfy the unique service placement constraint, which forces the model to place some of the delay-tolerant services closer to the user, thus increasing the edge's utilization.

Moreover, it was shown in [10] that the optimal formulation had an average runtime of around 70 ms while the heuristic had an average runtime of around 0.2 ms. Given that the CO-VSP optimization model is similar to the optimization model presented in [10] and the fact that the DA-VSP algorithm is an extension to the heuristic proposed in [10], the average runtime for the CO-VSP and DA-VSP are expected to be within the same ranges. This reiterates the real-time deployment potential of the DA-VSP algorithm with the CO-VSP model also being suitable to be used for benchmarking purposes.

VII. CONCLUSION

Vehicle-to-everything (V2X) communication, a cornerstone of intelligent transportation systems (ITSs), has garnered

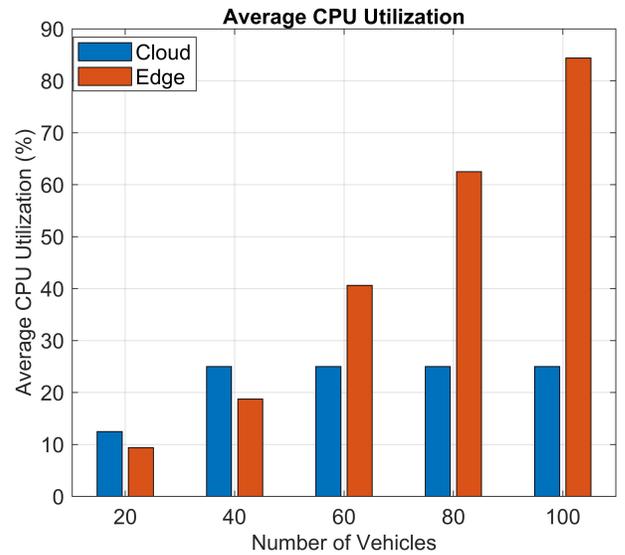


Fig. 4. Different Computing Nodes' Average CPU Utilization

more attention from different entities including governmental agencies and automotive manufacturers due to the many projected benefits of such systems. However, offering such services is coupled with maintaining stringent non-functional requirements such as E2E latency and performance. To that

end, mobile edge computing (MEC) has been posited as a viable solution to reduce the E2E latency/delay. Nonetheless, adopting a distributed computing environment introduces fresh challenges including the placement decisions of the V2X applications/services and the corresponding cost associated with hosting such services on edge nodes.

Therefore, this work formulated the cost-optimal V2X service placement (CO-VSP) problem while considering delay, computational, and placement requirements as a binary integer programming model. Moreover, this work developed a low-complexity heuristic titled “Delay-Aware V2X Service Placement (DA-VSP) Algorithm” as a solution to this problem. Through extensive simulation, it was shown that the proposed CO-VSP model successfully guaranteed the functional and non-functional requirements of the various V2X services. Moreover, the results highlighted the trade-off between deployment cost and latency in which delay-tolerant services tended to be placed at the cloud core to reduce the cost while delay-stringent services were placed at the edge to maintain their QoS requirements. Furthermore, it was observed that the proposed DA-VSP algorithm had comparable performance to the CO-VSP model while having lower computational complexity.

ACKNOWLEDGMENTS

This work is partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) [NSERC Strategic Partnership Grant STPGP - 521537].

REFERENCES

- [1] Z. MacHardy, A. Khan, K. Obana, and S. Iwashina, “V2X Access Technologies: Regulation, Research, and Remaining Challenges,” *IEEE Communications Surveys Tutorials*, vol. 20, no. 3, pp. 1858–1877, thirdquarter 2018.
- [2] 3rd Generation Partnership Project (3GPP), “Technical specification group services and system aspects; study on LTE support of Vehicle to Everything (V2X) services. Release 14,” 3GPP, Tech. Rep. TR 22.885, Dec. 2015.
- [3] C. V. S. C. Consortium *et al.*, “Vehicle safety communications project: Task 3 final report: identify intelligent vehicle safety applications enabled by dsrc,” *National Highway Traffic Safety Administration, US Department of Transportation, Washington DC*, 2005.
- [4] ETSI, “Intelligent transport systems (its); vehicular communications; basic set of applications; definitions,” Tech. Rep. ETSI TR 102 638, Tech. Rep., 2009.
- [5] A. Moubayed and A. Shami, “Softwarization, virtualization, & machine learning for intelligent & effective v2x communications,” *IEEE Intelligent Transportation Systems Magazine*, 2020.
- [6] L. Yang, A. Moubayed, I. Hamieh, and A. Shami, “Tree-based intelligent intrusion detection system in internet of vehicles,” in *2019 IEEE Global Communications Conference (GLOBECOM)*, 2019, pp. 1–6.
- [7] D. Sabella, N. Nikaiein, A. Huang, J. Xhembulla, G. Malnati, and S. Scarpina, “A Hierarchical MEC Architecture: Experimenting the RAVEN Use-Case,” in *2018 IEEE 87th Vehicular Technology Conference (VTC Spring)*, Jun. 2018, pp. 1–5.
- [8] W. Balid, H. Tafish, and H. H. Refai, “Intelligent vehicle counting and classification sensor for real-time traffic surveillance,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 6, pp. 1784–1794, 2018.
- [9] A. Leivadreas, M. Falkner, I. Lambadaris, M. Ibnkahla, and G. Kesidis, “Balancing delay and cost in virtual network function placement and chaining,” in *2018 4th IEEE Conference on Network Softwarization and Workshops (NetSoft)*, Jun. 2018, pp. 433–440.
- [10] A. Moubayed, A. Shami, P. Heidari, A. Larabi, and R. Brunner, “Edge-enabled V2X Service Placement for Intelligent Transportation Systems,” *IEEE Transactions on Mobile Computing*, pp. 1–1, 2020.
- [11] D. Martín-Sacristán, S. Roger, D. Garcia-Roger, J. F. Monserrat, A. Kousaridas, P. Spapis, S. Ayaz, and C. Zhou, “Evaluation of LTE-Advanced connectivity options for the provisioning of V2X services,” in *2018 IEEE Wireless Communications and Networking Conference (WCNC)*, Apr. 2018, pp. 1–6.
- [12] E. Aqeeli, A. Moubayed, and A. Shami, “Towards Intelligent LTE Mobility Management through MME Pooling,” in *2015 IEEE Global Communications Conference (GLOBECOM)*, 2015, pp. 1–6.
- [13] ETSI, “Intelligent Transport Systems (ITS) Vehicular Communications: Basic Set of Applications - Part 2: Specification of Cooperative Awareness Basic Service,” ETSI, 650 Route des Lucioles F-06921 Sophia Antipolis Cedex - FRANCE, Tech. Rep., 2011.
- [14] I. Shaer, A. Haque, and A. Shami, “Multi-Component V2X Applications Placement in Edge Computing Environment,” in *2020 IEEE International Conference on Communications (ICC)*, 2020, pp. 1–6.
- [15] Z. Amjad, A. Sikora, B. Hilt, and J. Lauffenburger, “Low Latency V2X Applications and Network Requirements: Performance Evaluation,” in *2018 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2018, pp. 220–225.
- [16] ETSI, “Intelligent Transport Systems (ITS) Vehicular Communications: Basic Set of Applications - Part 3: Specifications of Decentralized Environmental Notification Basic Service,” ETSI, 650 Route des Lucioles F-06921 Sophia Antipolis Cedex - FRANCE, Tech. Rep., 2010.
- [17] —, “Intelligent Transport Systems (ITS) Vehicular Communications: Basic Set of Applications - Definitions,” ETSI, 650 Route des Lucioles F-06921 Sophia Antipolis Cedex - FRANCE, Tech. Rep., 2009.
- [18] A. Moubayed, M. Injadat, A. Shami, and H. Lutfiyya, “DNS Typo-Squatting Domain Detection: A Data Analytics and Machine Learning Based Approach,” in *2018 IEEE Global Communications Conference (GLOBECOM)*, 2018, pp. 1–7.
- [19] ETSI, “LTE-Service Requirements for V2X Services,” ETSI, 650 Route des Lucioles F-06921 Sophia Antipolis Cedex - FRANCE, Tech. Rep., 2017.
- [20] MUX, “The low latency live streaming landscape in 2019,” 2019.
- [21] R. M. Karp, “Reducibility among combinatorial problems,” in *Complexity of computer computations*. Springer, 1972, pp. 85–103.
- [22] N. Megiddo and A. Tamir, “On the complexity of locating linear facilities in the plane,” *Operations research letters*, vol. 1, no. 5, pp. 194–197, 1982.
- [23] Microsoft, “How to: Change the size of a windows azure virtual machine,” Available at: [https://docs.microsoft.com/en-us/previous-versions/dynamicsnav-2013/dn168976\(v=nav.70\)](https://docs.microsoft.com/en-us/previous-versions/dynamicsnav-2013/dn168976(v=nav.70)), 2013.
- [24] K. Lee, J. Kim, Y. Park, H. Wang, and D. Hong, “Latency of Cellular-Based V2X: Perspectives on TTI-Proportional Latency and TTI-Independent Latency,” *IEEE Access*, vol. 5, pp. 15 800–15 809, 2017.
- [25] D. Dyachuk, “The true costs of running cloud native infrastructure,” Available at: https://sched.ws/hosted_files/kccncna17/5f/True%20Costs%20Dyachuk%20%282%29.pdf, 2017.
- [26] Federal Communications Commission (FCC), “A broadband network cost model,” FCC, Tech. Rep., May 2010.
- [27] N. Nikoorkaran, G. Karakostas, and T. D. Todd, “Combining capital and operating expenditure costs in vehicular roadside unit placement,” *IEEE Transactions on Vehicular Technology*, vol. 66, no. 8, pp. 7317–7331, Aug 2017.
- [28] J. Koomey, K. Brill, P. Turner, J. Stanley, and B. Taylor, “A simple model for determining true total cost of ownership for data centers,” *Uptime Institute White Paper, Version*, vol. 2, p. 2007, 2007.