Distributive Justice for Fair Auto-adaptive Clusters of Connected Vehicles

Julian Garbiso*†, Ada Diaconescu*, Marceau Coupechoux*, Jeremy Pitt‡, Bertrand Leroy†
*LTCI, Telecom ParisTech, IMT, Paris-Saclay University - Paris, France
{name . surname}@telecom-paristech.fr
†Institut Vedecom - Versailles, France
{name . surname}@vedecom.fr
‡Imperial College London - London, United Kingdom
j.pitt@imperial.ac.uk

***Abstract*—Connected vehicles will likely use hybrid communication networks. Presumably a licence-free radio access technology (RAT) will be used for vehicle-to-vehicle (V2V) contact, complemented by a cellular network, with an associated usage cost. In previous work, we developed a self-adaptive clustering algorithm for reducing cellular access costs, while ensuring that clustering overheads do not saturate the V2V link. However, the vehicle in the role of Cluster Head (CH) is the only one to bear the communication costs in a cluster's lifetime. This means certain drivers may pay much more than others for the same service, which may in turn undermine the system's social acceptability. In this paper, we adopt the theory of distributive justice to ensure fairness over time, and hence make the system socially acceptable. We compare the proposed approach with our previous algorithm through simulations, analyzing network performance and specific fairness metrics. We show that the proposed approach improves fairness metrics significantly, while not affecting network performance.**

I. Introduction

Numerous services and applications for connected vehicles are being conceived, from road safety to entertainment, exchanging different volumes of data, from vehicle to vehicle (V2V) or between vehicles and the internet. Two standards emerge as favourites for delivering the connectivity performance needed, with the specific constraints of vehicular networks. These are, on the one hand, IEEE 802.11p (an adapted version of the well-known *WiFi*) and on the other hand, LTE (the fourth generation cellular network standard). While the former uses a licence-free spectrum, the latter requires a usage cost, which we aim to reduce. In [1] and [2] we presented an Auto-adaptive Clustering Algorithm, where clusters improve the usage efficiency of the cellular network traffic through data aggregation and off-loading on the V2V link. The algorithm also ensures a trade-off between reduced traffic in the cellular network and saturation with increased packet loss into the V2V link.

The problem addressed in this paper is as follows. In clustering algorithms where the Cluster Head (CH) acts as a gateway to the internet, being elected CH (the only vehicle in a cluster which uses the cellular network in this model) is a significant burden since each vehicle is linked to an individual's account with a mobile operator, with a limited traffic quota that it will consume for the benefit of the group. It becomes evident that improvements need to be made in order to ensure that, even though at a certain moment only one vehicle is a CH and will consume its quota, there *will* be a *fair* quota consumption across vehicles *over time*. The cellular quota consumption is visible to the drivers. If they perceive that it is consumed in an unfair way, they can, at any moment, leave the system.

In this paper we propose an approach that adopts the *theory of distributive justice* and applies it to the election of the Cluster Head in clustering algorithms. In particular, we apply it to the self-adaptive clustering algorithm we presented in [2]. At the same time, we believe that our approach is similarly applicable to other clustering algorithms.

We establish a correspondence between the variables in our clustering model and a set of *canons* of justice. The vehicles that are in the process of forming a cluster will cast their votes for their preferred and most convenient canon. The final elector of the Cluster Heads, which in our case is the cellular base station, takes all votes into account, together with the state of each vehicle, in order to make the best decision for both performance and fairness criteria. The proposed fairness-aware algorithm is compared, via simulation, with the one presented in [2] with respect to both short-term network performance (for which that algorithm was designed) and long-term fairness metrics[1]. The results show that the algorithm improves fairness criteria without adversely affecting network performance.

The article is organized as follows. Section II provides information about the background literature concerning the theory of commons and of distributive justice. We introduce our model and algorithms in Section III. Subsequently, the simulation results are presented and analyzed in Section IV. Conclusions are drawn in Section V, where hints for future work are also presented.

II. Background and Related work

Our application domain features different types of shared resources – e.g. aggregated information from clusters of

[1]Possible comparisons with other methods that can be adapted to this problem will be the object of future work.

vehicles; subscription-based access to cellular networks, at a cost; and unlicensed V2V access, prone to saturation if unregulated. For all of these resources, we have different degrees of ability for regulation (for preventing someone from using them), and the usage that someone makes of each of these resources affects in different degrees how much others can yield from it. This corresponds to the criteria of *excludability* and *subtractability* used by Elinor Ostrom [3] to classify goods in four categories (i.e. public goods, toll goods, private goods and common-pool resources).

After studying several cases of resource exploitation in different communities, and for very diverse types of goods, Ostrom presented eight design principles for governing commons, ensuring the endurance of the resources. In our work, we have always kept this design principles in mind.

Ostrom's principles for managing commons are there to guarantee that a resource *endures*, but they do not imply that the outcome is actually fair. In [4], the author summarizes some desirable properties of a *fair* result:

- *Proportional*: Each agent receives a resource allocation proportional to its participation in the system;
- *Envy free*: The allocation of a given agent should not be "more desirable" to another agent, compared to its own allocation;
- *Equitable*: No agent derives more utility from its received allocation than any other agent in the system from its own;
- *Efficient*: Deliver the greatest good for the greatest number;
- *Cost-effective*: The computational cost of the distribution is not disproportionate;
- *Timely*: The computation of the distribution can be completed fast enough to be useful.

There are, nevertheless, two problems: on the one hand, these properties do not always come together and there can even be a competition between those aspects (like a cost-effective distribution method that does not produce an equitable outcome). On the other hand, and most importantly, these properties are usually applied to static distribution methods. The proposal in [4] focuses on *fairness over time*, through a series of individually unfair allocations that eventually lead to a fair result on the long term, by following Rescher's idea of *legitimate claims*. In this respect, our interest matches their motivation: it will be impossible to produce a fair outcome when designating a cluster head, but we aim to achieve a globally fair designation outcome over the long run.

Rescher [5] introduces the concept of *social justice*, which consists in determining what legitimate claims an individual has, and treating each of these claims equally. His work does not rely on how those claims can emerge, but rather focuses on the concept of *distributive justice*, which deals with how far the legitimate claims of each individual should be met, given competing claims from other people or entities, and limited resources.

The analysis made by Rescher about the most usual ways that other authors had used to determine fairness could fit in what he called the *seven canons of distributive justice*: treatment as equals, according to each one's needs, according to each one's productive contribution, to efforts and sacrifices, to the social value of the services provided by the individual to the society, to supply and demand, or to merits and achievements. Rescher found that none of these canons, taken individually, could grant true distributive justice. Instead, he proposed to analyze every participant's legitimate claims following each of these aspects, and focus on how to balance them in case of conflict.

III. Fairness-Aware Election Algorithm

In our previous work, we proposed a clustering algorithm that aims to optimise the usage of the cellular network resources. The intra-cluster communication is ensured by a license-free V2V network, while the costly communication with the cellular network is done via the Cluster Head. In this algorithm, all vehicles in the network are monitored by the nearest cellular base station. When there are no clusters formed in a certain area, the base station, depending on the vehicular density, decides the maximum number of hops (affecting both cluster size and V2V network performance) and elects the vehicle that is closest to the geographic center of the sector as Cluster Head (CH). For clarity, we will call this one the **Fairness Agnostic Algorithm**. The problem with this approach is that, since the CH is the only one to bear the cost in terms of cellular quota consumption at a given moment, this can lead to unfair cost distribution, which may become socially unacceptable. Therefore, the presented contribution aims to provide a **Fairness-Aware Algorithm**, which addresses this issue.

We propose a practical mapping of our problem of distributing the cost of the cellular network access among the clustered vehicles in terms of Rescher's seven canons (since the seventh canon is not translatable to any term of merit or achievement in our case since there are no such concepts in our problem, we are focusing only on the first six, merging two of them that become equivalent). This way, we create five criteria that constitute five total orders of the vehicles in a cluster region, that are subsequently used for deciding who has to be the Cluster Head. According to each canon, these are the methods to determine the order in which vehicles should be elected as CH:

- **The canon of equality**: In decreasing order of volume of data previously received as a Cluster Member;
- **The canon of needs**: In increasing order of the inverse of the available cellular quota ($1/AvailableQuota$);
- **The canon of productivity**: In increasing order of volume of data previously shared as Cluster Head;
- **The canons of effort / social utility**: In increasing order of the number of times having previously served as Cluster Head;

- **The canon of supply and demand**: In decreasing order of non-compliance events (for instance, switching off the data connection while occupying the Cluster Head role).

The data needed to calculate all five total orders for a group of vehicles is instantly available for the cellular base station, which is, in our model, the one that has the ultimate decision on who will be proclaimed as CH.

Each vehicle, however, has the knowledge of all the specific metrics for itself only, and can thus deduce which criteria are more favorable to its situation. Each vehicle then has the opportunity to demand to be judged by the most urgent claim it has (for example, if it has been elected CH too many times, or if it has a very low available quota, it can demand these criteria to have a higher importance), when all the vehicles in a cluster region will cast a vote that will result in the *weights* of each canon for this group of vehicles in particular. The base station, aware of this result, establishes a weighed Borda [6] election: each canon acts as one Borda voter. For instance, according to the canon of Equality, the order of priority for being elected as CH could be *"1) Vehicle A; 2) Vehicle B; 3) Vehicle C"* while the canon of needs could determine *"1) Vehicle B; 2) Vehicle C; 3) Vehicle A"*. If equal weights are assumed here, *Vehicle B* would be elected as CH.

The final decision is then made this way:

1) The base station discovers a clustering region without Cluster Head and requests the vehicles in the area to cast a vote;
2) Vehicles estimate the canon that is most favorable to them and vote for it locally. Each vehicle can cast only one vote, for only one canon. Since we have five criteria, the vote of vehicle *i* will be a vector $\mathbf{v_i}$ of five elements, with one element equal to 1, and the rest equal to zero[2];
3) The outcome of the vote is the set of weights for each canon, determined by the cellular base station after receiving the votes. The weight of the canon *j* will be determined by the following formula:
$$w_j = \frac{\sum_{i=1}^{N_c} \mathbf{v_i}[j]}{N_c}, \tag{1}$$
where N_c is the number of vehicles in the clustering sector;
4) The cellular base station calculates the total orders associated to each canon;
5) The cellular base station runs the weighed Borda election, by applying the weight coefficient to each Borda vote emitted by each canon;
6) The cellular base station applies a final correction to the total order obtained from the Borda election: a normalization lowering the score for being elected as CH as the vehicle's distance from the center of the clustering region increases.
7) The notification is sent to the designated Cluster Head.

[2]Different voting systems could also be implemented. Another option could be assigning all integral values from 1 to 5 in each of the vector's position, resulting in a nested Borda vote.

It is important to note that, when passing from one base station to another, the clusters remain formed as they are. Only when there is no cluster head in a *clustering sector*, a new CH election will take place. And when this happens, any base station has access to the same up-to-date information about every user's account information in the mobile operator's system. Thus, there is no way that passing from one base station to another could affect fairness.

IV. Simulations and results

A. Simulation settings

Our algorithms are implemented in the Veins simulation framework, which synchronizes a traffic simulator (SUMO [7]) and a full-stack network simulator (OMNeT++ [8]). The experiment that we ran for testing the proposed algorithm consists of a group of 100 vehicles that pass through a 10 km highway segment, for approximately 100 times each, with randomized order of re-entrance. The simulated access protocol in the V2V network is IEEE 802.11p. The protocol for the network and transport layers is IEEE 1609.3. In the application layer we implement the Cooperative Awareness Messages (CAM) from the European ETSI ITS G5 standard. The path loss follows a Two-Ray Interference model [9].

In the curves that follow, where time is represented in the horizontal axis, it refers to the simulated time, where the 100 vehicles randomly re-enter with a precise inter-arrival time. This is thus the equivalent of the time that it would take for 10,000 vehicles to traverse the 10 km highway segment, with an uniform inter-arrival time.

B. Vehicles having served as CH

The most remarkable outcome of introducing distributive justice in the long term is that the role of Cluster Head, which seems like a temporary burden, will be taken by a greater proportion of the participants if we analyze successive samples over time.

In Figure 1 we can see how the curve of the number of vehicles having been Cluster Head rises much faster when using the *Fairness Aware Algorithm*: after **less than 4 simulated hours**, all of the vehicles have taken the CH role at least once. In contrast, with the *Fairness Agnostic Algorithm*, it is only after **more than 12 hours** that we reach the same result.

On the other hand, we have to analyze how many times the same vehicle has been elected as CH. Reducing the amount of times that a specific user takes this role is an important way for improving the user's perception of justice. In Figure 2 we can observe the evolution of the maximum number of times that a specific vehicle has been elected CH. We see that, for the maximum value, the *Fairness Agnostic Algorithm* is always in a situation of almost doubling its counterpart.

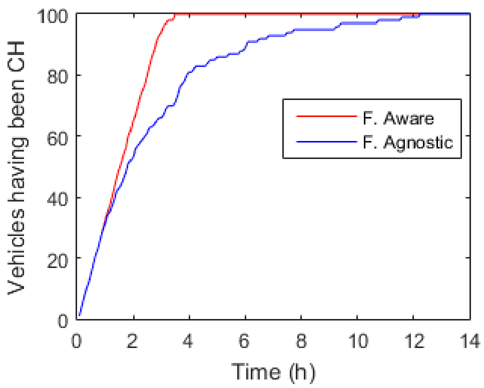

Fig. 1: Number of vehicles having served as CH *at least once*, in function of time.

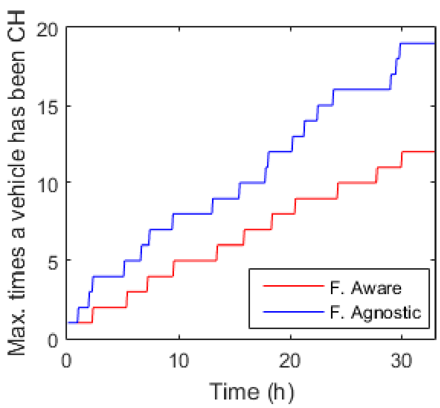

Fig. 2: Maximum number of times that any one vehicle has served as CH in function of time.

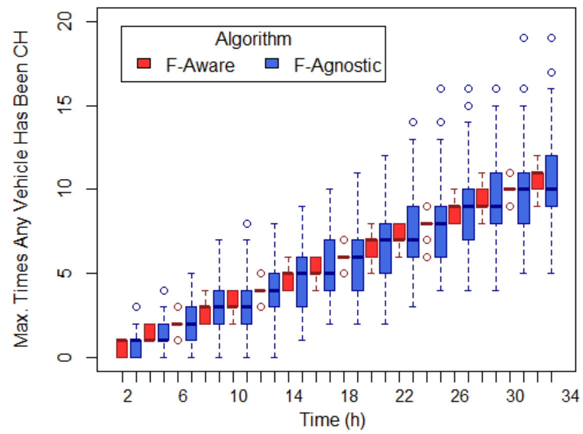

Fig. 3: Box plot showing the distribution of the number of times that every vehicle has served as CH, for both algorithms, over time.

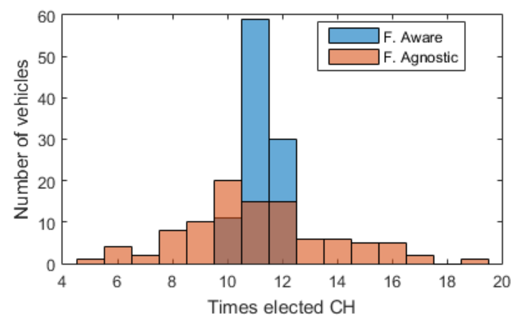

Fig. 4: Histogram of the final number of times that any one vehicle has served as CH.

The box-plot[3] in Figure 3 shows the distribution of CH assignments across vehicles, over time. Regardless of the time passing, the size of the boxes of the *Fairness Aware Algorithm* remain small and don't change their size. Their whiskers are usually tiny or non-existent. This means that, even if the average number of times that vehicles in general serve as CH increases with time (which is necessary), all of the vehicles serve approximately the same amount of times as CH. On the other hand, for the *Fairness Agnostic* boxes, their size increases, and the whiskers keep getting bigger, showing extreme disparities between the participants.

The effect in the long term can be seen in the *"final picture"* of Figure 4 where we can see the histogram of the number of times that every vehicle has served as CH when the end of the simulation is reached. We see an significant dispersion between values ranging from 4 to 20 for the *Fairness Agnostic Algorithm*, while for its *Aware* counterpart, almost 60% of the vehicles served exactly 11 times as CH, while the other 40% did so either 10 or 12 times.

[3]In the box plots in this article, the solid band represents the median (second quartile), while the box is delimited by the first and third quartiles. The whiskers mark the lowest and highest datum still within 1.5 Inter-quartile Range (IQR) of the lower and upper quartiles, respectively. The outliers are marked as circles.

C. *Cellular quota consumption*

The differences in the outcome of the Cluster Head election that we have just discussed in the previous point, has a direct impact on cellular quota consumption. We will now see some examples of the different distribution of the cellular quota usage with both algorithms, which translates almost directly into different economic costs.

1) A better worst case: We start by analyzing the case of our most unfavoured user for both algorithms in Figure 5. This curve shows the lowest individual available quota among the vehicles who have been CH, over time. We can see a pronounced difference in the slopes. This means that even the most unfavoured participant will always be much better off, in terms of economic cost, with the *Fairness Aware Algorithm* than with the agnostic one.

2) A fair distribution: Let us now see the impact on the available quotas of all vehicles (including those that have not been CH at a given moment). Figure 6 shows the box plot of the available quotas of all vehicles over time. Once again, the *Fairness Aware Algorithm*'s boxes and whiskers remain

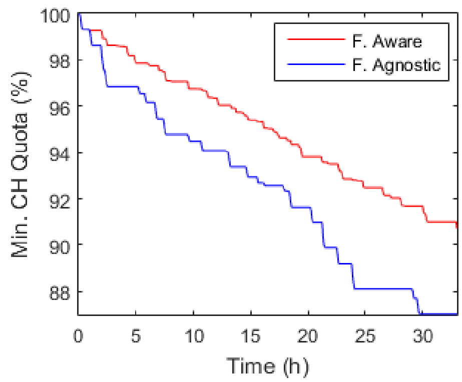

Fig. 5: Minimum available cellular quota among the vehicles having served as CH during the simulation.

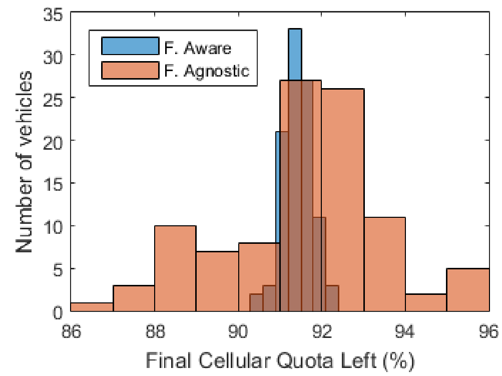

Fig. 7: Histogram of the final available cellular quotas among the vehicles having served as CH.

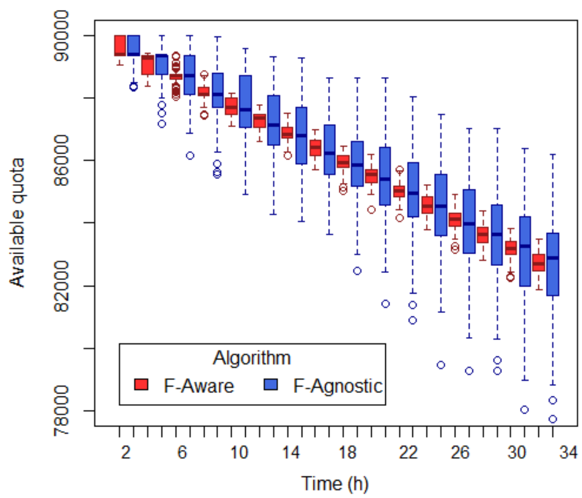

Fig. 6: Box plot showing the distribution of the available cellular network quota for every vehicle, for both algorithms, over time.

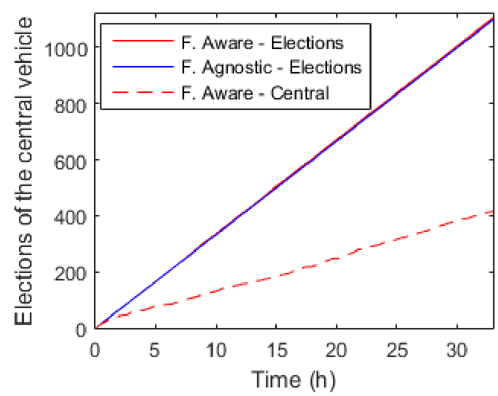

Fig. 8: Total number of CH elections in both algorithms, compared to the number of elections in which the *Fairness Aware Algorithm* and the *Fairness Agnostic Algorithm* make the same choice: the vehicle in the geometric center of the clustering sector.

always small, tightly following the straight line of the global average data consumption. On the other hand, the distribution of the quota consumption (and hence the economic cost) in the *Fairness Agnostic Algorithm* becomes more unfair with time, with major deviations from the median.

The clear final picture of the situation can be seen in the histogram of Figure 7, where we can see that for the *Fairness Aware Algorithm*, all 100 vehicles fit in a very tight range of final quotas, while for its *Agnostic* counterpart, they are distributed in a much broader range of possible final quotas, with just a few vehicles in each bin.

D. Similar elections

The *Fairness Agnostic Algorithm* has only one possible outcome for every CH election: the vehicle which is closer to the geometric center of the clustering sector. For the *Fairness Aware Algorithm*, this criterion is merely one amongst several choices (though it has a special weight, being considered in a second phase of the election process). Figure 8 shows the proportion of elections that share the same result in both algorithms: those where the "central" vehicle is elected. The solid line represents the number of CH elections taking place over time, while the dashed red line shows the number of elections in the *Fairness Aware Algorithm* that result in the central vehicle being elected. It is interesting to see that in both algorithms, the number of elections over time remains the same, and that the slight deviation from the original algorithm's geometry implied by not always electing the central vehicle does not have an impact on the number of elections. On the other hand, the proportion of elections having the outcome of electing the central vehicle for the *Fairness Agnostic Algorithm* is, trivially, equal to the number of elections.

E. Network performance

The original (*Fairness Agnostic*) clustering algorithm aimed to self-adapt the size of the clusters in order to keep an optimal balance in terms of network performance: maximizing data compression in the cellular network (and

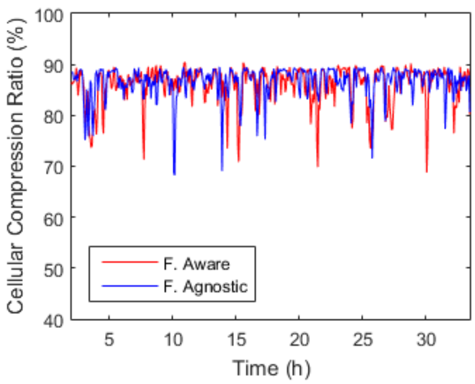

(a)

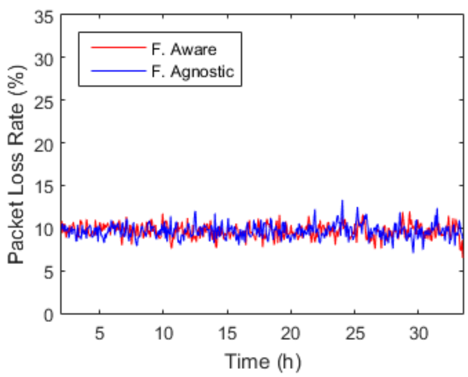

(b)

Fig. 9: Comparison between the *Fairness Aware* and *Fairness Agnostic* algorithms, regarding: (a) Data compression ratio on the cellular network link over (simulated) time. (b) Packet loss ratio (PLR) on the V2V network. The results show that incorporating fairness improvements does not have a cost in the network performance metrics that the *Fairness Agnostic* algorithm was designed for.

thus reducing economic cost) while limiting the Packet Loss Rate (PLR) in the V2V network, to a maximum acceptable threshold of 10%. We could expect that modifying the geometry and criteria of the election process could affect the metrics that this algorithm was designed for. Figures 9a and 9b compare both algorithms in terms of data compression and PLR, respectively. As we can clearly see, the improvements in distributive justice come at **no cost** in terms of network performance.

V. Conclusions and Future Work

We have presented a Fairness Aware Election Algorithm based on the theory of distributive justice. This solution aims to address the problem of social acceptability of clustering algorithms, where the individuals elected as Cluster Heads bear all the costs. We have applied this approach to our previous work on a specific self-adaptive clustering algorithm, conceived for reducing cellular access costs by only communicating aggregated data via the CH, while limiting overheads onto the V2V network. The previous algorithm was effective in terms of global data compression and network performance, but unfair with respect to CH elections and hence to the cost distribution across system users (and since a driver can keep track of its quota consumption, it could act as a motivation to leave the system, potentially leading to a collapse). In contrast, the proposed fairness-aware algorithm enables vehicles to influence CH elections by expressing their legitimate claims via voting (based on a Borda-type vote). Numerical simulations showed that this approach significantly improves fairness across all vehicles, over time (analysis based on several metrics that a regular user would most likely analyze when elaborating its own perception of justice). Meanwhile, simulations also showed that the fairness-aware algorithm preserves all network performance optimisations achieved by its fairness-agnostic counterpart.

We aim to design a system for cluster-based vehicular networks that follows Elinor Ostrom's [3] design principles for managing commons. These principles ensure that a resource endures, being governed equitably and sustainably. Our algorithm currently follows some of these principles. The vote of the participant vehicles (P1) for determining the weight of each legitimate claim in the CH election is a way for them to be involved in the rule-making process. The groups have clearly limited boundaries (P1). Even one of our criteria (taking into account non-compliance events) gets close to the gradual sanctioning required in Ostrom's principles (P6). In future work, it would be interesting to consolidate the monitoring and sanctioning required by Ostrom, and explore possible comparisons of the Fairness Aware Algorithm with other existing methods that may be adapted to this problem.

References

[1] J. Garbiso, A. Diaconescu, M. Coupechoux, and B. Leroy, "Dynamic cluster size optimization in hybrid cellular-vehicular networks," in *Intelligent Transportation Systems (ITSC), 2016 IEEE 19th International Conference on*. IEEE, 2016, pp. 557–563.

[2] ——, "Auto-adaptive multi-hop clustering for hybrid cellular-vehicular networks," in *Intelligent Transportation Systems (ITSC), 2017 IEEE 20th International Conference on*. IEEE, 2017.

[3] E. Ostrom, R. Gardner, and J. Walker, *Rules, games, and common-pool resources*. University of Michigan Press, 1994.

[4] J. Pitt, D. Busquets, and S. Macbeth, "Distributive justice for self-organised common-pool resource management," *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, vol. 9, no. 3, p. 14, 2014.

[5] N. Rescher, *Fairness: Theory and practice of distributive justice*. Transaction Publishers, 2002.

[6] J.-C. de Borda, "Mémoire sur les élections au scrutin, histoire de l'académie royale des sciences," *Paris, France*, 1781.

[7] M. Behrisch, L. Bieker, J. Erdmann, and D. Krajzewicz, "Sumo–simulation of urban mobility," in *The Third International Conference on Advances in System Simulation (SIMUL 2011), Barcelona, Spain*, 2011.

[8] A. Varga, "Discrete event simulation system," in *Proc. of the European Simulation Multiconference (ESM'2001)*, 2001.

[9] C. Sommer, S. Joerer, and F. Dressler, "On the Applicability of Two-Ray Path Loss Models for Vehicular Network Simulation," in *4th IEEE Vehicular Networking Conference (VNC 2012)*. Seoul, Korea: IEEE, November 2012, pp. 64–69.