Global On-line Optimization for Charging Station Allocation

Jean-Marc Lasgouttes

Global On-line Optimization for Charging Station Allocation

Jean-Marc Lasgouttes

Inria

France

Email: jean-marc.lasgouttes@inria.fr

Abstract—The goal of the Mobility2.0 project was to provide full electric vehicles users with a set of tools that reduce range anxiety and favor a partial modal shift towards public transportation. As part of this work, we have designed a so-called global optimization criterion for selecting charging stations where change of mode can occur. The idea is to minimize the mean quadratic travel time of all users, in a way that can be used for on-line allocation with good performance. We show through simulation that this leads to a sizable improvement with respect to a "greedy" station selection.

I. Introduction

The benefits of electric vehicles are well-known today: better energy-efficiency, clean air, better utilization of renewable energy sources, etc. However, the widespread adoption of full electric vehicles (FEVs) still faces many hurdles, which are yet to be overcome. We focus herein on the following issues:

- Limited FEV range may lead to *range anxiety* for drivers in the absence of a procedure which ensures that they comfortably reach their destination in any circumstance;

- Public parking spaces in urban areas are generally a limited resource and, especially during the initial deployment stages, the electrification of parking spaces may not keep up with the number of FEVs;

- The scarcity of urban road space leads to general traffic congestion which can be alleviated by an at least partial modal shift towards public transport.

The Mobility2.0 project [1]–[5] has been completed in February, 2015. It has developed and tested an in-vehicle commuting assistant for FEV mobility, enabling more reliable and energy-efficient electro-mobility by controlling available range and proposing as needed a *park and ride* type of trip. A suitable traffic management system for FEVs must take each of the above bottlenecks simultaneously into account in order to achieve a journey optimization which is acceptable and favorable for all drivers. Such integrated approach necessitates the use of co-operative systems to co-ordinate between vehicles and the infrastructure. The Mobility2.0 server is accessed from nomadic device applications for iOS and Android. It comprises 3 core components:

- the *Range Estimator* computes the FEV battery energy consumption for a given leg of a journey (see [5]);

- the *FEV Demand Predictor* handles all interactions related to charging;

- the *Multimodal Journey Planner* relies on the two other components to propose possibly multimodal trips that satisfy range constraints.

The Multimodal Journey Planner (MJP) is a layer over OpenTripPlanner (OTP), an open source platform for multimodal and multi-agency journey planning [6]. Beyond standard OTP features—finding a route using either the FEV or public transportation, subject to user preferences—it provides:

- Bi-modal journeys (FEV and public transportation) through the selection of a charging or parking station;

- Mitigation of range anxiety through selection of routes that are feasible, with or without leaving the FEV in a station, given the current battery state;

Since the bulk of automotive transportation consists of daily commutes to the office and back, Mobility2.0 has focused on this scenario. The MJP suggests to use a parking station when it is faster to drive the FEV to a station and then use public transportation to reach the final destination (typically when the workplace in city center). When the trip is not feasible without recharging the battery, the MJP selects a charging slot and makes sure that the charging period is long enough to be able to finish the trip. This last case is an important part of our strategy to mitigate range anxiety. Indeed, it guarantees the feasibility of the trip in the current conditions.

The possible journeys returned by the algorithm are sorted by increasing travel time, but the computation of these times using OTP can be expensive. To limit the computation time to a reasonable limit like e.g. 10 seconds, the implementation uses a two-tier computation, first using coarse heuristics based on historical data to rank the solutions, then a more precise prediction using OTP. The technical description of this algorithm is out of the scope of this paper. Instead, we discuss another contribution, which is a new type of journey selection which goes beyond fastest trip: journeys can be sorted according to the so-called *global optimization* criterion, which aims to minimize globally the travel time of all users. It is believed that this global approach to resource allocation will help to both reduce further range anxiety and favor use of public transportation. This idea is related to the theory developed by Wardrop [7], where two main types of traffic assignment equilibria have been identified:

- user optimal: each user non-cooperatively seeks to minimize his cost of transportation;

- system optimal: each user behaves cooperatively in choosing his own route to ensure the most efficient

use of the whole system; in general economists recommend to use pricing to attain this equilibrium.

The paper is organized as follows: The model parameters are given in Section II. The new optimization criterion is defined in Section III, and an approximate on-line algorithm is derived, both with limited and unlimited range; complexity issues are also discussed. Section IV is devoted to the evaluation of the algorithm via simulation; it shows that this optimization criterion can lead to a sizable performance improvement.

II. Model Setting

Before proceeding to describe the algorithm, it is necessary to describe the model at hand in more details.

Let us first consider the demand. It is common to define users through an Origin/Destination (O/D) matrix: the city is split into several zones, and the number of trips between these zones is counted. This is usually done through the so-called Trip Distribution, the second component (after trip generation, but before mode choice and route assignment) in the traditional four-step transportation forecasting model [8]. We will only characterize a user by his type t, that is his origin and destination. As mentioned in previous section, we chose to restrict the demand to the simplest type of daily travels of the form origin/destination/origin. The data can be collected by using the Mobility2.0 system in a given city for a long enough period. Each user of type t (for a given O/D/O scheme) has a probability of occurrence and a travel time for the different journeys that is estimated using OpenTripPlanner.

We turn now to the problem of FEV range. The probability $\mathbb{P}(R \leq r)$ that the initial battery range for a user is less than some value r is supposed to be independent from the user type. The empirical distribution can be estimated together with the computation of the O/D matrices above. The energy necessary for a given trip is obtained through the Range Estimator service of the Mobility2.0 server. Note that no energy is needed for a direct trip using public transportation; this implies that every user will get at least one feasible route, and also provides an upper bound on travel times.

The expected travel times are supposed to be known. In practice, they will be estimated using the results of previous similar OTP requests done by users, since no better data is easily available. The charging or parking stations contain a number of slots; slots are assumed to be allocated for the duration of the period, i.e. it is not possible to have one slot used successively by two users. Currently the algorithm does not make a difference between a charging station and a mere parking area. Handling this would be trivial, but make the description of the algorithm more involved.

To sum up, the parameters of our model are the number of stations S, the number of types T and, for each user type t:

- n_t: number of users expected during the session;

- τ_{ts}: expected travel time for a trip via station s;

- τ_{td} (resp. τ_{tp}): expected travel time for a direct trip using a FEV (resp. public transportation);

- $\mathbb{P}(R \leq r)$: cumulative distribution function of the FEV initial range;

- r_{ts} (resp. r_{td}): energy required for a trip via station s (resp. for a direct FEV trip).

III. The Global Optimization Criterion

The aim of the global optimization algorithm is to provide a cost function that measures the performance experienced by the whole community of users. The cost is defined as the mean squared travel time, which is more sensitive to large values; using the mean travel time would have the inconvenience of "sacrificing" some users by giving them very long travel times.

Our setting can be seen as a resource allocation problem, known as the Transportation Problem in Operations Research literature [9]. It is solvable using several algorithms, among which the simplex algorithm or the Hungarian algorithm [10], [11]. Unfortunately, these algorithms are not well-adapted here for two reasons:

- The allocation of slots to users is done on-line, when the user does a request. It is not possible to wait until all the users are known before doing the allocation;

- The complexity of these algorithms is very high, especially since, due to the effect of range limitations, each request has different characteristics, which is equivalent to having TS types of customers.

Another way to present the problem is in terms of Markov Decision Process [12], which leads to linear programming approaches, presumably too expensive here.

We therefore present a simple heuristic approach, where the allocation cost for one customer scales as TS (Section III-C). This is fast enough for systems with thousands of stations. Its principle is to penalize the cost for the user with an approximation of the extra cost incurred to future users who compete for the same resource (a charging or parking slot). Since the implications can be intricate, we only consider a first order effect.

A. The Case of Unlimited Range

Let us explain first the principle when the range of the FEV is unlimited. In this case, all journeys are feasible and the only task is to select the best one. The simple case is when user A considers a direct journey using either a FEV or public transportation. This choice will not affect the possibilities offered to future users and the penalized cost is therefore equal to the normal cost.

Assume now that user A considers using a station s. This choice will only cause a problem if, at some time in the future, s minimizes the travel time for some other user B, but B finds it already full at arrival. To compute the effect of this conflict, we have to compute (1) its probability and (2) how it influences the travel time of user B. Of course, the modification of behavior of B will have consequences itself. We chose to ignore these second order effects, which makes our computation a bit optimistic.

To compute the probability of a conflict between two users, we use the model of the demand. The probability that the number of potential users of a given station is larger than the number of available slots is obtained using a binomial formula.

To compute the effect of a conflict, one has to compare the cost of station s for user B with the second best choice he can do. Unfortunately, this second choice is not known yet, since the situation will be different at this point of time. It is though possible to propose two simple bounds. A lower bound is given by the station (or direct trip) that is the second-best choice at the time when we make the decision for user A. This value is optimistic, since the situation will worsen with time. Conversely, an upper bound is given by a direct trip using FEV or public transportation. This value is pessimistic, but it compensates the optimistic nature of our algorithm remarked before. This is the bound we are going to use thereafter.

Assume that users of types $t \in \mathcal{T}_{\mathbb{A}}$ are in concurrence with user A at station s and that their weight (probability of occurrence) is w_t. Using the upper bound above, the respective cost increases they can experience are

$$\Delta C_{ts} = \min\left(\tau_{td}^2, \tau_{tp}^2\right) - \tau_{ts}^2.$$

Note that $\Delta C_{ts} \geq 0$, since s is the available station that minimizes travel time for users of type t and is therefore better than a direct trip. If $w = \sum_{t \in \mathcal{T}_{\mathbb{A}}} w_t$ is their total weight, then the expected extra cost in case of conflict is

$$\Delta C_s = \mathbb{P}(B_{w,n_{\mathrm{rem}}} \geq n_s - 1) \times \frac{1}{w} \sum_{t \in \mathcal{T}_{\mathbb{A}}} w_t \Delta C_{ts}, \quad (1)$$

where n_{rem} is the remaining number of customers to come, n_s the number of remaining slots in station s, and $B_{w,n}$ a binomial random variable counting events of probability w after n draws.

B. Introducing Range Effects

When the range is limited, the formula has to be modified to take into account the accessibility of station s for user A given his FEV range, which is trivial. However, it is also necessary to assess the possibility for future users to reach stations for their journeys. In order to do that, one has to modify the above algorithm as follows.

- Each user B that can potentially be in conflict with A is weighted by the probability that its range is large enough to use station s: $w_{ts}^r = w_t \mathbb{P}(R > r_{ts})$;

- The direct journey used as an upper bound for user B is a weighted mean of the FEV-only journey cost and of the public transportation trip, since the journey could be too long for the FEV range of user B:

$$\Delta C_{ts}^r = \mathbb{P}(R \leq r_{td} | R > r_{ts}) \tau_{tp}^2$$
$$+ \mathbb{P}(R > r_{td} | R > r_{ts}) \min\left(\tau_{td}^2, \tau_{tp}^2\right) - \tau_{ts}^2.$$

As a consequence, equation (1) is still valid, except that the values of w_t and ΔC_{ts} have to be replaced:

$$\Delta C_s^r = \mathbb{P}(B_{w',n_{\mathrm{rem}}} \geq n_s - 1) \times \frac{1}{w^r} \sum_{t \in \mathcal{T}_{\mathbb{A}}} w_{ts}^r \Delta C_{ts}^r$$

C. Implementation Issues

One important concern about the algorithm is scalability. Using the usual big-O notation[1], the complexity of a greedy

---
[1] $u = O(v)$ as $v \to \infty$ if, for $v > v_0$ there exists M such that $|u| \leq M|v|$.

allocation is $O(S)$ (a single loop over all the stations). A global allocation requires in theory $O(TS^2)$ steps: a loop over the stations and, for each type of user B, a greedy allocation to search for conflicts.

The dependency on the number of types is a problem because, as will be shown in next section, a precise model improves the quality of global allocation. There is not much that can be done to alleviate this, though. The dependency in S^2 is more worrying at first sight. The problem is mitigated in the implementation by the fact that the stations are sorted by expected travel times for each user type. The code then maintains a pointer to the best available station for type t users, which avoids to walk through the whole list of stations at each allocation. The complexity of greedy search for a total of N users thus becomes $O(N + S)$, so that the complexity of doing each global allocation is $O\big(TS(1 + S/N)\big)$. Moreover, only stations more convenient than a direct trip need to be considered. In a plausible city model (but not in our Gaussian toy model unfortunately), this caps the effect of the S term.

IV. Evaluation by Simulation

Evaluation should ideally be done using the data from the on-site validation periods of Mobility2.0 (in Barcelona and Reggio Emilia). However, they have been neither long enough nor large enough to provide enough data for our statistical model. Creation of synthetic data to feed the MJP with simulated requests would have been possible, but it would have required scalability properties for the server well beyond the scope of the project.

For these reasons, it has been decided that it is better to focus on a preliminary evaluation to get a sense of what global optimization can bring. To this end, we designed a streamlined simulator, distinct from the Multimodal Journey Planner, that is set in a simplified setting but uses the same station selection algorithms as the MJP and that can scale to thousands of stations. One advantage of going towards a synthetic benchmark is that the parameters of our models can vary more freely.

As described before, two main schemes can be used to select a charging station: the user chooses either the fastest possible journey (greedy) or the best journey with respect to the extended cost defined in Section III (global). While users use in general a mix of these two policies, we will evaluate them separately and compare their performance under a large-scale deployment. The case where only a fraction of users are using the global policy is addressed at the end of Section IV-B.

A few shortcuts are taken for the sake of simplicity. They are not expected to modify the validity of our evaluation and could be replaced by something more elaborate without any notable problem:

- All stations have 10 slots;

- All the types of users have the same weight $w = 1/T$;

- The initial range of a FEV is either unlimited or uniform in some range $[R^-, R^+]$;

- The required energy r_{ts} or r_{td} is directly related to travel time; while this is not true in general, the

TABLE I.   Mean travel times (in minutes) for the Gaussian toy model.

	Convenient station	Normal station	Only FEV	Only public transportation
Travel time	20	40	60	80
Energy	10	20	60	0

TABLE II.   Comparison of the different allocation algorithms (unlimited range).

Allocation type	Quadratic mean travel time (min.)	Improvement w.r.t. greedy
Greedy	43.52	–
Global	39.79	8.6%
Off-line	37.60	13.6%

purpose of this evaluation is not to model energy consumption and therefore this first-order approximation is sufficient.

The simulator has been implemented as a command-line C++ utility, completely separate from the actual MJP. However, the greedy and the global allocation algorithms correspond to the real specification. When the range of vehicles is unlimited, it is also possible to run the streamlined transportation simplex algorithm proposed in Chapter 8 of [9] on the whole population of users generated for the simulation, in order to compare the obtained performance[2]. This is interesting because this algorithm (which is expensive and cannot be run on-line) will provide a bound for the attainable performance.

A. Experimental Setup

Since we have made our model abstract enough, the only thing that is necessary is to provide numbers for the travel times. We use in this evaluation a Gaussian toy model, which is defined as follows:

- the travel times are non-negative truncated Gaussian variables with means given in Table I and standard deviations equal to 20% of their means;

- for each type of user, there is one station (picked at random) which is more convenient than the others. The intent is to increase competition between users;

- When using a station, it is assumed that the second part of the trip will be done using public transportation, and therefore the required energy is smaller than the travel time;

Note that the values given in Table I are somewhat arbitrary. They have not been chosen to estimate how much time a FEV user can gain in a real situation: the goal is to observe the behavior of the global algorithm and characterize its performance. In this setting, using a station to park the car and take public transportation is time-efficient, as is supposed to be the case in a peak-hour scenario. Using only public transportation takes however longer than using only the FEV.

Most of the time, we use a setting where the number of charging slots is lower than the number of users:

- $N = 20000$ users;

- $S = 1000$ stations with 10 slots each, that is a total of 10000 available slots;

- $T = 3000$ types of users (equivalent to ≈ 55 origins and 55 destinations).

These numbers have been chosen as follows: 1000 charging stations is an ambitious number for a big city. To see better
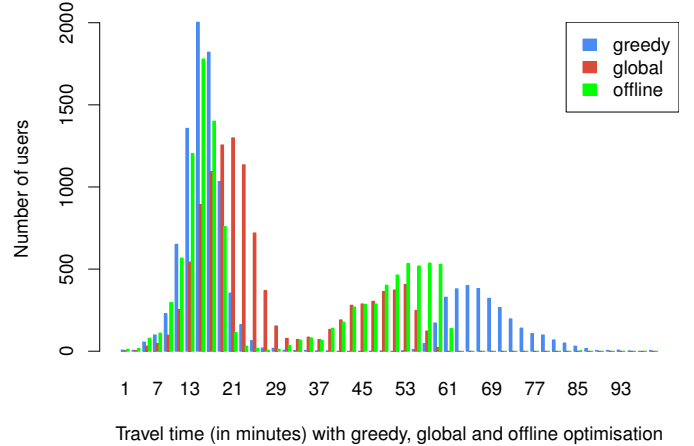

Fig. 1.   Distribution of the travel times according to the different allocation schemes (unlimited range).

the effect of competition for parking space, it is reasonable to pick a number of users that is twice as large as the number of slots. In terms of performance assessment, we will explore the users' travel time in three ways: quadratic mean for all users, distribution and improvement seen with global optimization.

B. Experimental Results with Unlimited Range

It is convenient to start by evaluating the performance of the algorithms with unlimited range. Indeed, this is the only case where it is possible to run the transportation simplex algorithm, which yields the best possible allocation for a given set of users: the run time of the algorithm increases with the number of types of users and, when the range is limited, each user is different because some do not have a good enough energy level to attain some stations.

The most synthetic result is shown in Table II. One sees that global allocation improves the quadratic mean travel time by 8.6%. While this does not seem to be much at first glance, this number is to be compared to the (theoretic) lower bound of the transportation simplex algorithm. This maximal expected gain is 13.6%, with an algorithm that cannot run on-line, contrary to our operational requirements. The gain obtained with the global algorithm is thus very encouraging.

The distribution of travel times, depending on the allocation technique used, is given in Fig. 1. Greedy allocation produces the highest number of quick trips, but the downside is that it also leads to the highest number of long travel times. Global allocation reduces the duration of these long trips, but this is done at the cost of making short trips more expensive. Finally, the off-line allocation strikes a balance and somewhat provides the advantages of both methods at the same time.

---

[2]We used the code distributed at http://darrentmacdonald.com/emd/

TABLE III.   CLASSIFICATION OF USERS IN TERMS OF GAIN G DUE TO GLOBAL ALLOCATION (UNLIMITED RANGE).

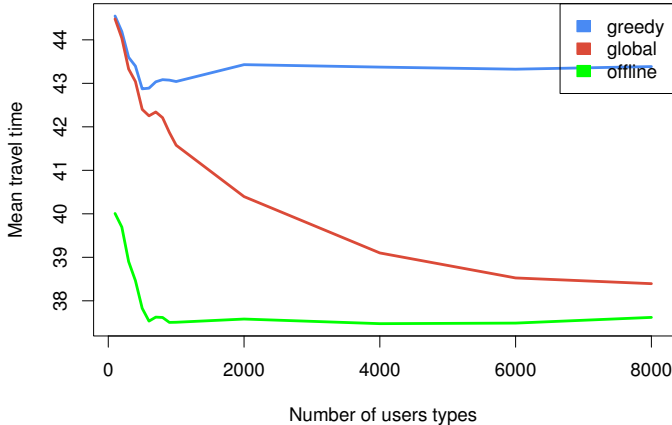|	Weight	Mean gain (min.)
$G < -20$ (loss)	14%	-32
$|G| < 20$	71%	-2
$G > 20$ (gain)	15%	49

Fig. 2.   Influence of the number of user types on the mean quadratic travel time for each of the allocation procedures (unlimited range).

Let us turn now to the time that each user gains (or loses) when everybody follows the global allocation guidelines instead of being greedy. It is possible to separate the users in 3 classes, as shown in Table III. One can see that, while the majority of users will not see the difference, the gain of time experienced by the 15% of lucky users is larger than the time lost by the 14% that do not benefit from the scheme.

The influence of the precision of the model—that is, the number of types of users—on the performance is shown in Fig. 2. Contrary to the greedy and off-line optimization cases, the global optimization algorithm improves notably when the quality of the input data increases. This implies that the efficiency of a Mobility2.0 service will improve with time, but also that the number of user types in the model should be large enough to be efficient. It is important to note that classifying 20000 users into 5000 classes requires several months of historical data. The fact that the off-line algorithm is better than our global optimization algorithm comes mostly from the fact that it has a perfect knowledge of all users in the session; it is not clear that the non-optimality of our algorithm is a problem.

Until now, only the case where all users use the global allocation policy has been considered. In a real situation, part of the Mobility2.0 users will disregard the global allocation proposal and pick shortest journeys. Moreover, a number of users not using the system will compete for charging stations; for simplicity, we assume that these users are identical to "greedy" Mobility2.0 users. Fig. 3 show the influence of the proportion of "good citizens" among the population on the performance of the system. After increasing a little bit at 10%, the mean quadratic travel time becomes better as soon as more than 20% of users use the global allocation choice. Interestingly enough, it is the mean travel time of the whole
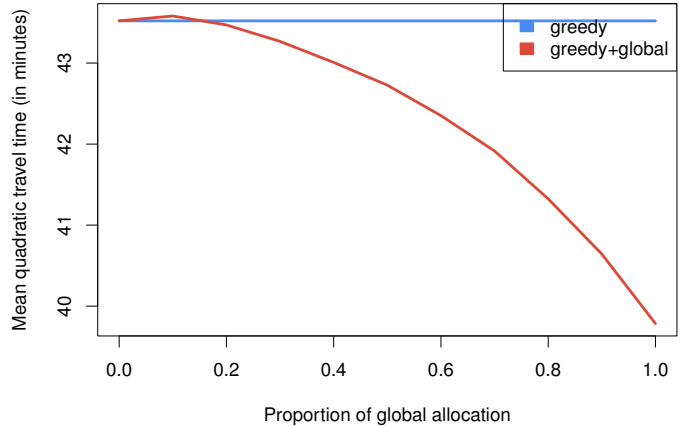

Fig. 3.   Influence of the proportion of users using global allocation on the mean quadratic travel time for each of the allocation procedures (unlimited range).

TABLE IV.   COMPARISON OF THE DIFFERENT ALLOCATION ALGORITHMS (LIMITED RANGE).

Allocation type	Quadratic mean travel time (min.)	Improvement w.r.t. greedy
Greedy	47.97	–
Global	40.54	15.48%

TABLE V.   CLASSIFICATION OF USERS IN TERMS OF GAIN G DUE TO GLOBAL ALLOCATION (LIMITED RANGE).

	Weight	Mean gain (min.)
$G < -20$ (loss)	16%	-34
$|G| < 20$	67%	-2
$G > 20$ (gain)	17%	61

population that improves here, not only the Mobility2.0 users subset.

C. Experimental Results with Limited Range

When the range of FEVs is limited, we expect some vehicles to be prioritized over others because their range is too small. In the following evaluation, the range is taken to be uniform in the range $[45, 90]$, which means that direct trips using only the FEV may not be doable by some users. In this new situation, since all users are different, it is not possible anymore to run the off-line algorithm as a comparison, for complexity reasons. The rest of the analysis of Section IV-B can be repeated.

As can be seen in Table IV and Fig. 4, the results are even better than they were in previous section. Despite the fact that the range of the vehicles is random, which makes the prediction of future demand more difficult, a more substantial gain is obtained by anticipating the arrival of FEVs with low range. Similarly, the distribution of results in classes is still good (Table V).

There is another side effect that is more surprising: for each user, it is possible to compare the travel time knowing his FEV range, to the one that would be experienced if the range was unlimited (Fig. 5). For the greedy allocation, the travel time is
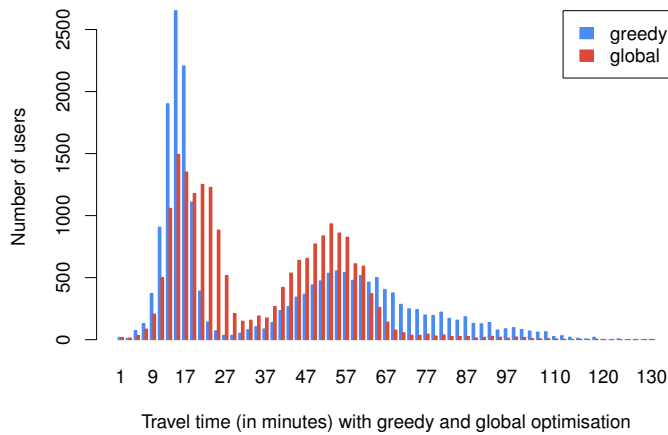
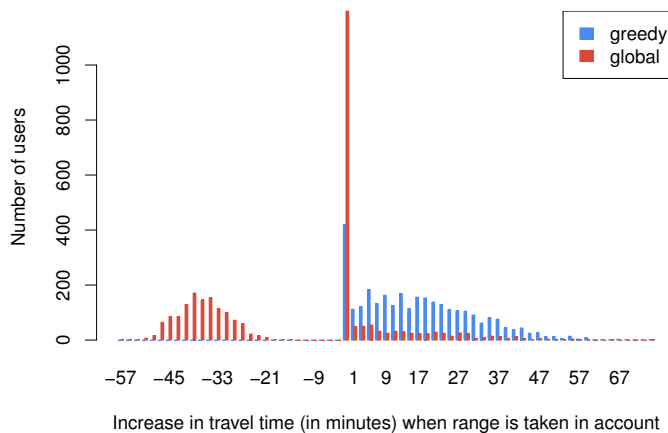Fig. 4. Distribution of the travel time according to the different allocation schemes (limited range).

Fig. 5. Distribution of the increase of travel time with limited FEV range. Only the 3002 users for which the range makes a difference are shown here.

unsurprisingly longer when the range of the FEV is limited. For the global version however, there are a number of users for which the travel time improves when the range is limited, because they get some priority over FEVs with a better range (this is the red heap at the left of the distribution). It can be debated whether this is a good feature, since it will be probably counter-intuitive for many users.

V. CONCLUSION

We have shown in this study how the global allocation algorithm can provide a good improvement to the travel times experienced by the users of the Mobility2.0 system. This scheme allows alleviating the longest users' travel times without sacrificing other users. Moreover, it acts as a prioritizing policy that helps users with a low range. However, these results can only be attained if many users are "good citizens", and especially those who would need to sacrifice some of their time for the benefit of others. Further work is required to understand what incentives can be proposed to users to make them do "the right thing", e.g. financial incentives for drivers to use the recommended re-charging spots. Most importantly, these results show that the global allocation policy is an important tool for achieving the optimal use of an existing public FEV re-charging infrastructure, maximizing its benefit to drivers.

When it comes to being scalable to a million of users, an interesting behavior of the algorithm will only be observed if the number of available parking slots is of the same order of magnitude as the number of users. Otherwise, the effect of any allocation algorithm will be marginal at best, since most users will be constrained to do a FEV-only trip. It is not clear at this point that one can be in a situation where a good allocation strategy can improve significantly the travel times of FEV users. But if this situation can be built, the global algorithm that is proposed here is a very good candidate.

REFERENCES

[1] The Mobility2.0 project. [Online]. Available: http://www.mobility2.eu/

[2] A. Solar, A. Bolovinou, G. Heijenk, J.-M. Lasgouttes, and R. Giménez, "Mobility 2.0: Co-operative ITS systems for enhanced personal electromobility," in *27th International Electrical Vehicle Symposium & Exhibition (EVS27)*, Barcelona, Spain, 2013. [Online]. Available: http://hal.inria.fr/hal-00921558

[3] M. Fernández, A. Bolovinou, G. Karagianinis, J.-M. Lasgouttes, J. Mittag, I. Varea, R. Kurpatov, Z. Kovács, and F. Mastrandea, "Co-operative commuting assistant for enhanced personal electromobility," in *10th ITS European Congress*, Helsinki, Finland, 2014.

[4] M. Gholibeigi, G. J. Heijenk, D. Moltchanov, and Y. Koucheryavy, "Analysis of a receiver-based reliable broadcast approach for vehicular networks," in *2014 IEEE Vehicular Networking Conference, VNC 2014*, 2014, pp. 89–96.

[5] A. Bolovinou, I. Bakas, A. Amditis, F. Mastrandrea, and W. Vinciotti, "Online prediction of an electric vehicle remaining range based on regression analysis," in *Electric Vehicle Conference (IEVC), 2014 IEEE International*. IEEE, 2014.

[6] OpenTripPlanner, multimodal trip planning & analysis. [Online]. Available: http://www.opentripplanner.org/

[7] J. G. Wardrop, "Some theoretical aspects of road traffic research," *Proceedings of the Institute of Civil Engineers, Part II*, vol. 1, pp. 325–378, 1952.

[8] M. G. McNally, "The four-step model," in *Handbook of Transport Modelling*, D. A. Hensher and K. J. Button, Eds. Elsevier, 2000, pp. 35–52.

[9] F. S. Hillier and G. J. Lieberman, *Introduction to Operations Research*, 7th ed. McGraw-Hill, 2001.

[10] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, pp. 83–97, 1955.

[11] J. Munkres, "Algorithms for the assignment and transportation problems," *Journal of the Society for Industrial and Applied Mathematics*, vol. 5, no. 1, pp. 32–38, 1957.

[12] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.