On Fairness in the Carpool Problem

Moni Naor*

Abstract

We discuss the *carpool problem*, where one has to assign drivers to subsets of participants who commute regularly. Fagin and Williams defined a notion of *fair share* for participants (the FW share). We provide an axiomatic characterization of the fair share and show that the FW share is the unique one satisfying these requirements. We define a coalitional game where the Shapley value is the FW share.

1 Introduction

The fair carpool problem is as follows: A group of n people decide to form a carpool. Each day a subset of these people will arrive to a (fixed) gathering point and one of them should drive. A method for determining which person should drive on any given day is required. This immediately raises the issues of fairness. The problem can thus be separated into two: (i) For a given schedule what is the “fair share” of each participant and (ii) how to assign each day a single driver among the people who show up on this day so that the total number of times each participant drives is (approximately) this share. This is especially challenging when done online, without the benefit of knowing the future schedule.

While this sort of problem was undoubtedly faced by many groups throughout human history, whenever individuals had to perform undesired chores, apparently it was first explicitly studied by Fagin and Williams [2] and continued by Ajtai et al. [1]. They proposed that the fair share of a participant should be the sum of the inverses of the number of people who showed up on the days the participant took part in the carpool. We will call it the FW share.

How close can each participant be to their FW share? First, as pointed out in [1], if the schedule is given ahead of time (the offline case), then there exists an assignment where each participant’s load (the number of times driven) is within 1 of their FW share. When one is not given the schedule in advance, but rather the decision has to be made online each day, then there are deterministic scheduling algorithms that guarantee that for every participant the difference between the load (the number of times driven) and the FW share is $O(n)$ (and this is the best possible deterministic guarantee), and probabilistic ones where the difference is $O(\sqrt{n \log n})$, while the lower bound is $\Omega(\sqrt[3]{\log n})$.

We observe the following two appealing properties regarding the online scheduling algorithms of Ajtai et al. [1], specifically the Local Greedy Algorithm and the reduction from the edge orientation to the general carpool problem (Theorem 5.1 there):

- The unfairness of a participant i (the difference between his FW share and the load) is bounded by the number of people with whom i interacts (i.e. shows up on the some day). In other words, even if there is a very large carpool group, but i interacts with only a few of them, (say n_i) then his unfairness will be very small ($O(n_i)$ deterministically and $O(\sqrt{n_i \log n})$ probabilistically).

*Incumbent of the Judith Kleeman Professorial Chair, Dept. of Computer Science and Applied Math, Weizmann Institute of Science, Rehovot 76100, Israel. E-mail: naor@wisdom.weizmann.ac.il. Supported in part by an IST Grant RAND-APX.

- There is no need for maintaining any global information, nor is there any need for trusted bookkeeping, e.g. currency. It is enough that every two participants record their mutual history (actually a single bit).

Fairness: One possible rational for the FW notion of fairness is that if the driver were not a member of the group, but rather a hired driver, then the meaning of fairness would be clear: the professional driver charges a fixed price for every ride, and each day the people that show up split the price of the driver equally among them. When the driver is just one of the set of people that show up, this reasoning leads to the following: If on a certain day, d people show up, each of them owes the driver $1/d$ of a ride. Another explanation is that if the people who show at any day choose the driver at random, then the FW share represents the expected number of times a participant will be chosen.

While these reasons are compelling, doubt might still linger as to whether some other reasonable notions of fairness can be suggested[1]. How does one argue for the Fagin-Williams notion of fairness? In this note we take an axiomatic approach to the problem, listing four very basic requirements from the fair share. We prove that the only notion of fairness that satisfies these requirements is the one suggested by Fagin and Williams. We then show in Section 3 that the FW share coincides with the Shapley value of a certain cooperative game.

2 Requirements

For a group of n participants denoted by $[n]$ a schedule T is a list $D_1, D_2, \ldots D_\ell$ where each $D_i \subseteq [n]$ corresponds to the people who showed on day i (we assume that no subset is empty, that is on each day someone wants to ride). Given such a schedule we want to assign for each participant i a fair share $\varphi_i(T)$. Using this terminology, the FW share for each participant i is:

$$\sum_{\{1 \le j \le \ell | i \in D_j\}} \frac{1}{|D_j|}.$$

The following are our requirements from the fair share function:

Full Coverage: The sum of the fare shares of all participants equals the total number of days in the schedule, i.e. $\sum_{i=1}^{n} \varphi_i(T) = \ell$.

Symmetry: If two participants have *exactly* the same schedule, that is they appear on the same set of days, then their fair shares should be the same.

Dummy: A participant that never shows on any day has share 0.

Concatenation: Given two schedules T_1, T_2 consider the schedule T resulting from the concatenation of the two. Then the fair share of each player i for T is the sum of the fair shares of i for T_1 and T_2, i.e. $\varphi_i(T) = \varphi_i(T_1) + \varphi_i(T_2)$.

Are these requirements self evident? This certainly seems to be the case with *Full Coverage* and *Symmetry*: for the former we definitely want the fair shares to cover all the days and for the latter this seems to be the essence of fairness. *Dummy* also sounds very natural, after all what can we require from a participant that never shows up? As for *Concatenation*, this is essential if one wants to claim that past behavior does not unduly influence the future returns. Note however that our observation is "value neutral" - all it says is that if you wish to satisfy the above requirements, then the FW share is the only one that does so.

Our main claim is the following:

[1] For instance, we have heard of a carpool where each non-driver pays the driver $1/(d-1)$ in each ride where there are d people. Also given the argument of the hired driver, why should she charge as a taxi cab, and not as a shuttle bus?

Theorem 1 *The FW share of a schedule is the unique one satisfying the requirements: Full Coverage, Symmetry, Dummy and Concatenation.*

Proof: The theorem can be proved by induction on the schedule length. It is true for a schedule consisting of a single day: from the Symmetry, Dummy and Full Coverage requirements, each participant that shows on that day get share $1/|D_1|$ and the rest get 0. If it holds for schedules of length less than ℓ, then for a schedule T of length ℓ consider schedules T_1 which consists of the first $\lfloor \ell/2 \rfloor$ days and T_2 which consists of the last $\lceil \ell/2 \rceil$ days of T. Then from the Concatenation requirement and the induction hypothesis:

$$\varphi_i(T) = \varphi_i(T_1) + \varphi_i(T_2) = \sum_{\{1 \le j \le \lfloor \ell/2 \rfloor | i \in D_j\}} \frac{1}{|D_j|} + \sum_{\{\ell - \lceil \ell/2 \rceil + 1 \le j \le \ell | i \in D_j\}} \frac{1}{|D_j|} = \sum_{\{1 \le j \le \ell | i \in D_j\}} \frac{1}{|D_j|}.$$

$\square$

It also holds that each of the properties *Full Coverage, Symmetry, Dummy* and *Concatenation* is essential to obtaining the uniqueness result of Theorem 1. Without Full Coverage, multiplying the FW share by any (fixed) non-zero value will still satisfy the remaining requirements. Without Symmetry we can assign each day D_j the participant with the smallest index in D_j. Without Dummy we can partition the load of D_j between the participants not in D_j. Without concatenation a solution assigning all the load to the participant(s) with the lexicographically largest incidence vector (the vector indicating on what days she showed up) is acceptable.

3 Coalitional Games and the Shapley Value

Cooperative Game Theory deals with games where the outcome of the game is a coalition (see [5] Part IV). One of the simplest forms of games considered are called *coalitional games with transferable payoff*. Such a game is defined by set N of players and function v that assigned every non-empty $S \subseteq N$ a real number $v(S)$. The number $v(S)$ is viewed as the total payoff to be partitioned among the members of S. The question is then which coalition will be formed and what payoff each member should receive. There are many solution concepts for such games, each with different axiomatic characterizations of the payoff vector. The concept that we will be interested in is the Shapley Value, suggested by Lloyd Shapley in 1953 [6], which has been used as a power index [7] and for cost allocation [8].

We now give a *computational* definition of the Shapley value. Let $\Delta_i(S)$, the *marginal contribution* of player i to a coalition S, be $v(S \cup \{i\}) - v(S)$. The Shapley value of player i is defined as

$$\psi_i(v, N) = \frac{1}{|N|!} \sum_{R \in \mathcal{R}} \Delta_i(S_i(R))$$

where $\mathcal{R}$ is the set of all possible $|N|!$ orderings of N and $S_i(R)$ is the set of players appearing before i in order R. This can be viewed as the expected contribution of i to its predecessors given a random order of the players.

We define a coalitional game v_T for each schedule T of the carpool problem. The set of players N is the set of n participants. For a schedule T let the game v_T be as follows: the value $v_T(S)$ of any subset $S \subseteq N$ is the number of days members of S showed up, i.e.

$$v_T(S) = |\{1 \le j \le \ell | \; D_j \cap S \ne \phi\}|.$$

Theorem 2 *For each player i the Shapley value of the game $\psi_i(v_T, N)$ is*

$$\sum_{\{j | i \in D_j\}} \frac{1}{|D_j|}.$$

Proof: The claim follows from viewing the Shapley Value as the expected marginal contribution of i, i.e. $E[\Delta_i(S_i(R))]$ over random order R of the players. For a given order R, we have that $\Delta_i(S_i(R))$ is the number of days D_j such that $i \in D_j$ and i is the first member of D_j to have arrived under the order R of the players. For a random choice of R, the probability that i is the first member of D_j to have arrived is $1/|D_j|$. Therefore $E[\Delta_i(S_i(R))] = \sum_{\{j|i \in D_j\}} \frac{1}{|D_j|}$. $\square$

Therefore we conclude that for any schedule T, the Shapley value of a player in the game v_T is the FW share of the participant in schedule T.

4 Discussion

What have we gained from posing the carpool problem as coalitional game? In his original paper Shapley proved that the so-named value is the unique solution satisfying three axioms known as Symmetry, Dummy Player and Additivity. These correspond (roughly) to our requirements from a fair share. Perhaps more insight is gained by the characterization of the Shapley Value in terms of *objections* and *counterobjections* due to Myerson [4] (see Section 14.4.1 in [5]). Roughly speaking, following [5], for a given payoff vector an objection of player i against player j is an argument "give me more or I will leave the game causing you to lose." A counterobjection is an argument "It is true that if you leave then I will lose, but if I leave you lose at least as much"

What is known about the Shapley value is that it is the unique value where every objection is balanced by a counterobjection of the above type. Therefore if one views these objections and counterobjections as reasonable or expected interactions during the negotiations between the carpool participants (on the system to be used), then it is reasonable to predict that the solution that will emerge is the Shapley value, or equivalently the FW share.

The most outstanding algorithmic question is to close the gap for probabilistic scheduling mentioned in the introduction (between $\Omega(\sqrt[3]{\log n})$ and $O(\sqrt{n \log n})$). It was shown in [1] that the case where each day two participants arrive (i.e. $|D_j| = 2$ for $1 \leq j \leq \ell$) captures the general case, via a scaling type construction. For the two-participants-a-day case Kalai [3] has recently suggested an algorithm where the maximum unfairness of any participant is not more than $O(\log \ell)$ with high probability for any sequence of length ℓ. It is interesting whether it is possible to get rid of the dependency on ℓ.

Acknowledgements

I thank Ron Fagin for interesting discussions regarding the axiomatization of the FW share and Ronen Shalitel and Udi Wieder for critical comments.

References

[1] M. Ajtai, J. Aspnes, M. Naor, Y. Rabani, L. J. Schulman, O. Waarts, *Fairness in Scheduling*, J. Algorithms 29(2), 1998, pp. 306–357.

[2] R. Fagin and J. H. Williams. *A Fair Carpool Scheduling Algorithm*, IBM Journal of Research and Development, 27(2):133–139, March 1983.

[3] A. Kalai, Personal Communication, Dec. 2001.

[4] R. B. Myerson, *Graphs and Cooperations in Games*, Mathematics of Operations Research 2, 1979, pp. 225–229.

[5] M. J. Osborne and A. Rubinstein, **A Course in Game Theory**, MIT Press, 1994.

[6] L. S. Shapley, *A value for n-person games*, in **Contributions to the Theory of Games**, vol 2, eds H. Kuhn and A. W. Tucker, Princeton University Press, 1953, pp. 307–317. Reprinted in **The Shapley Value: Essays in honor of Lloyd S. Shapley**, ed. A. E. Roth, Cambridge University Press, 1988.

[7] L. S. Shapley and M. Shubik, *A Method for Evaluating the Distribution of Power in a Committee System,* American Political Science Review 48, 1954, pp. 787– 792.

[8] M. Shubik, *Incentives, decentralized control, the assignment of joint costs and internal pricing,* Management Science 8(3), 1962, pp. 325–343.