Spring 5-29-2024

Welfare Maximization in the Airplane Problem

Alina Chadwick
alina.v.chadwick.24@dartmouth.edu

Welfare Maximization in the Airplane Problem

Alina Chadwick

2023-2024

1 Abstract

Given a set of passengers and a set of airplane seats, the goal of the airplane problem is to sit passengers in seats in a way that maximizes the sum of their total welfare, that is, the total happiness of the passengers in the plane. We aim to maximize their welfare subject to three constraints and how much they care about each constraint being satisfied: a group constraint (where passengers may want to sit together), a constraint on where in a row passengers want to sit (i.e. a window seat, a middle seat, or an aisle seat), and finally a constraint on where in the plane (closer to the front or to the back of the plane) passengers would like to be seated. In this paper, we argue that valuations of passengers on seats are both *superadditive* and, more strongly, *supermodular*, and we are therefore able to employ a greedy framework that maximizes the sum of social welfare for passengers with supermodular valuations on seats.

2 Introduction

The *airplane problem* is the following. There is a set of passengers who are looking to be seated in a simple plane. We can represent this simple plane as two sets of three-seat rows (with seats *ABC* and *DEF* for each row) with a set number of rows. Figure 1 depicts a sample simple plane.

Passengers may be traveling solo or may belong to a *group* of passengers. We limit groups to no more than three passengers; if a group has a size of greater than three, it may be broken down into smaller groups (i.e. a group of four can be broken up into two groups of two, a group of five can be broken up into a group of two and a group of three, and so on). Each group of passengers is considered one *entity* and has the same *valuation function* in our problem that maps all possible sets of seats to how happy the entity would be from receiving that given set of seats. Each entity gains *utility* or *welfare* from an assignment that they value positively (i.e. being assigned that seat or set of seats is better to them than not being assigned anything). Our goal is to maximize the sum of total welfare among the passengers, meaning that this problem is an instance of the *welfare maximization problem*.

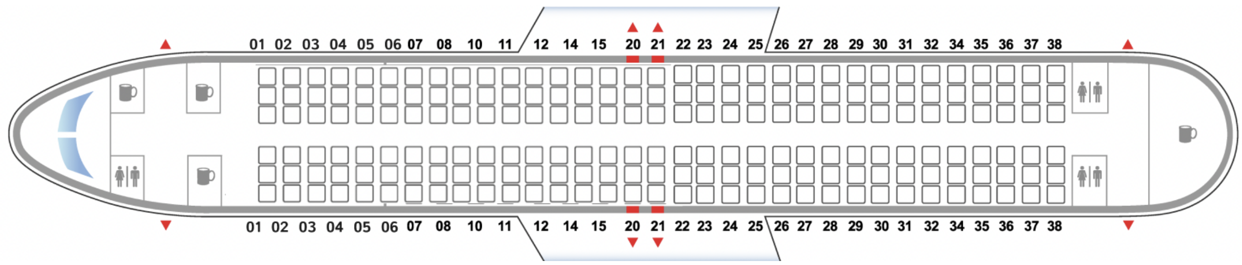

Figure 1. A sample simple plane.

Each entity has a valuation function on the seats of the plane, subject to some constraints. These constraints include the geometric constraint of seating groups together in addition to constraints regarding preferences on whether passengers prefer to sit in window, middle, or aisle seats (the *WMA* constraint) and preferences on whether an entity prefers to sit in the front or the back of the plane (the *F2B* constraint). Each entity also has a *willingness to pay*, which corresponds to the price they are willing to pay on a group of seats that is the size of the number of passengers in their entity. If an entity *e* has a willingness to pay that is twice the willingness to pay of entity *e'*, their valuation function will be scaled accordingly. Specifically, if all valuation functions are on a scale of 0 to ∞ and entities *e* and *e'* value a given bundle of seats the same amount, the valuation function will be scaled such that the valuation of group *e* will be twice that of group *e'*. In other words, we care more about accommodating the welfare of those who are willing to pay more than that of those who are not.

Intuitively, the valuation functions for this problem are *superadditive*, meaning that an entity reaches a high level of utility once it receives the number of seats needed such that each group member has a seat, but receive lower levels of utility for any seat allocated after that. Formally, superadditive valuation are defined as follows:

Definition 2.1. *A valuation function v is superadditive if for any disjoint sets A and B, $v(A \cup B) \geq v(A) + v(B)$.*

The groups constraint of this problem is superadditive for the following reason: when a group is assigned some seats, they obtain very little utility if that set of seats that has fewer seats than the number of people in their entity, or if the seats are not near each other. However, when a group receives a set of seats that is as big as the size of the group and the seats are close together, the utility for the group increases substantially. In other words, if a group of size s receives a set of seats close together that is of size $s-1$, they will receive much less utility from the $(s-1)$th seat than they would from the sth seat. However, for any additional seat, the marginal utility that they would receive would decrease.

To illustrate an example, suppose we have a family of three who want to sit together. We will treat them as one entity. This entity receives very little utility when they receive one seat or two seats (or three seats that are disjoint or located far away from one another), as they want to have all of their members seated in the plane and seated together; if they receive exactly three seats next to each other, the family is happy. However, if they receive four or more seats, they aren't gaining as much utility from each additional seat, since the entity doesn't have passengers to put into those seats. Their valuation of those additional seats is still positive, as they may want those seats to store items (e.g. musical instruments, medical equipment) or may use them if there is a technical difficulty with another one of their seats. Therefore, their valuations are superadditive. In this paper, we will argue that the valuation of entities on seats is not only superadditive, but also *supermodular*.

3 Preliminaries

Definition 3.1. *We first formally define an instance of the airplane problem* $I(E,S,v)$*:*

- E *is a set of* n *entities (passengers or groups of passengers)* $1,\dots,n$.
- S *is a set of* m *seats* $1,\dots,m$.
- v *is a vector of valuation functions* $v_1,\dots,v_n$ *where* v_i *is the valuation function associated with entity* i.

A *feasible solution* is a mapping $SOL : E \to S$ allocating each seat to exactly one entity (passenger or group of passengers). Our goal is to maximize social welfare, $v(SOL) := \sum_{e\in E} v_e(SOL)$.

3.1 Constraints

We will next formally define the constraints of the problem.

Definition 3.2. *The group constraint is defined as follows:*

- *Each passenger belongs to a group of size at least one.*
- *Groups are limited to a size of at most three, given that there are three seats per row. Groups of larger sizes are to be split into groups of size three or smaller, where no split group is of size one or fewer.*
- *Groups of passengers must prioritize sitting with their group over their other preferences (i.e. passengers in a group would rather sit together and not have their WMA or F2B preferences satisfied than sit separately but have their WMA and/or F2B preferences satisfied). In our algorithm, the group constraint will always be satisfied first.*
- *Within the groups, we will maximize the utility of individual passengers subject to their preferences in terms of which seats individuals within the group prefer to sit in over others. This is implicitly incorporated into our framework.*

Definition 3.3. *The WMA (window, middle, aisle) constraint and the F2B (front to back) constraint are defined as follows:*

- *For the WMA constraint, entities select their preference of whether they prefer a window, middle, or an aisle seat and in which order they prefer one over the other.*
- *For the F2B constraint, entities select their preference of whether they prefer seats in the front or the back of the plane as a linear gradient.*
- *Entities rank their preference with respect to their other preferences (i.e. whether their WMA preference is more or less important than their F2B preference, or vice versa).*

3.2 Supermodularity

We will first formally define supermodularity, then argue that valuations of entities on sets of seats are not only superadditive but also supermodular. We will be introducing this modification of the utility function to satisfy the algorithm.

Definition 3.4. *For every subset X, $Y \subseteq S$, with $X \subseteq Y$ and every $s \in S \setminus Y$, we have that*

$$v(X \cup \{s\}) - v(X) \leq v(Y \cup \{s\}) - v(Y)$$

In other words, supermodular valuations are superadditive (i.e. *Supermodular* $\subseteq$ *Superadditive*), since supermodularity requires the same property of complementarity of marginal utility. One of the key distinctions of supermodular valuations from superadditive valuations, however, is that supermodular valuations involve a specific type of complementarity where the marginal value of adding an item to a set is higher when the set is larger. In this problem, that translates to the following idea: losing one seat that is "relevant" to the set of other seats an entity has translates to a larger loss if they have more seats (i.e. a larger group) than if they have fewer.

To illustrate that the valuation functions of entities on seats is not only superadditive but also supermodular, we present a sample valuation over sets of seats. Suppose there are only four seats on the plane: 1A, 1B, 1C, and 5F. Additionally suppose that our entity is of size three (as in the previous examples). Figure 2 displays the entity's utility on all 16 subsets of seats. In this figure, the entity receives more utility from receiving more seats (satisfying superadditivity) and relatively more utility from receiving seats that are closer together than receiving seats that are further apart. Once they receive three seats close together, the entity's marginal utility spikes significantly. For an additional seat after that, the marginal utility becomes smaller, but the entity's valuation on that newer seat is still positive and significant. In our example, the family may want this additional seat to store their things on it or to have another seat to rely on in case there is an issue with any of their three allocated seats. We will now prove that the supermodularity condition holds in our case, given these valuations:

Let $v_e(S)$ be the valuation function of entity e on a subset of seats S. Then, supermodularity implies that, in this example,

Seat or set of seats $\subseteq S$	V_e
$\emptyset$	0
1A	1
1B	1
1C	1
5F	1
1A, 1C	2
1A, 5F	2
1B, 5F	2
1C, 5F	2
1A, 1B	3
1B, 1C	3
1A, 1C, 5F	4
1B, 1C, 5F	5
1A, 1B, 5F	5
1A, 1B, 1C	15
1A, 1B, 1C, 5F	18

Figure 2. An entity's utility on all subsets of seats

.

$$v_e(1A, 1B + 1C) - v_e(1A, 1B) \leq v_e(1A, 1B, 5F + 1C) - v_e(1A, 1B, 5F).$$

According to Figure 2, that translates to

$$15 - 3 \leq 18 - 5$$

Therefore, entities' valuations of seats in the airplane problem are not only superadditive, but are also supermodular based on their defined valuation functions.

4 Algorithm

We must first define a *supermodular dependency set* as defined in Feige and Izsak [1], as it is critical to the algorithm. To define the supermodular dependency set, however, we must first define the *marginal valuation function*.

Definition 4.1. *Let e be an entity in E and let s be a set in S. The marginal valuation function $v_e(s|X)$ is a function mapping each subset $X \subseteq S \setminus \{s\}$ to the marginal value of s given X:*

$$v_e(s|X) := v_e(X \cup \{s\}) - v_e(X)$$

We can now go on to define a supermodular dependency set:

Definition 4.2. *A supermodular dependency set for a given seat s, $D_e(s)$, subject to an entity's valuation function v_e is the set of other seats s' in S such that there exists a subset in $S \setminus \{s\}$ such that $v_e(s|S) > v_e(s|S \setminus \{s'\})$. A supermodular degree, d, is the maximum number of seats in a supermodular dependency set for a given v_e.*

In the example with the family of three, their supermodular degree would be two. If this family were to be assigned one seat s, they would only achieve their maximum marginal utility if they were also assigned to two seats next to s in addition to the original seat s. Specifically, for this entity, their supermodular dependency set on seat $1B$ could consist of seats $1A$ and $1C$.

4.1 Greedy Algorithm

We utilize the framework proposed in Feige and Izsak [1]. In any given iteration, for an entity $e \in E$ and a seat s, let *$D_e(s)$* be the set of seats that are supermodularly dependent on seat s subject to valuation v_e.

Algorithm 1 Greedy algorithm for Approximating Welfare Maximization for Supermodular Valuations in the Airplane Problem

---

Input: An instance of the airplane problem $I(E, S, v)$ and a supermodular dependency graph for each valuation function. Output: A solution with approximation guarantee $\frac{1}{d+2}$, where d is the supermodular degree of $I(E, S, v)$.

$\quad$ $UnoccupiedSeats \leftarrow S$
$\quad$ $Approx \leftarrow \emptyset$
$\quad$ **while** $UnoccupiedSeats \neq \emptyset$ **do**
$\quad\quad$ $MaxMarginalWelfare \leftarrow -1$
$\quad\quad$ **for all** $s \in UnoccupiedSeats, e \in E$ **do**
$\quad\quad\quad$ **if** $v_e(s,\, D_e(s) \cap UnoccupiedSeats \mid \{s' \in S \,|\, (s' \rightarrow e) \in Approx\}) > MaxMarginalWelfare$ **then**
$\quad\quad\quad\quad$ $MaxMarginalWelfare \leftarrow v_e(s,\, D_e(s) \cap UnoccupiedSeats \mid \{s' \in S \,|\, (s' \rightarrow e) \in Approx\})$
$\quad\quad\quad\quad$ $BestAlloc \leftarrow (\{s\} \cup (D_e(s) \cap UnoccupiedSeats) \rightarrow e)$
$\quad\quad\quad\quad$ $HappiestEntity \leftarrow e$
$\quad\quad\quad\quad$ $AllocatedSeat \leftarrow s$
$\quad\quad\quad$ **end if**
$\quad\quad$ **end for**
$\quad\quad$ $Approx \leftarrow Approx \cup BestAlloc$
$\quad\quad$ $UnoccupiedSeats \leftarrow UnoccupiedSeats \setminus (AllocatedSeat \cup D_{HappiestEntity}(AllocatedSeat))$
$\quad$ **end while**

In this algorithm, if the given entity gains a higher marginal welfare than previously seen from a given seat s and it's corresponding dependency set $D_e(s)$, that value of marginal welfare is stored as the maximum, the allocation is stored as the best allocation, and the entity is stored as the happiest entity, until all unoccupied seats and entities have been reached. By treating groups of passengers as entities, we unite their valuation functions on seats and their corresponding dependency sets. This, along with the fact that we have supermodular valuations, allows us to utilize this algorithm.

Treating groups of passengers as entities and assuming their valuation functions are supermodular allows us to tackle the group constraint with this algorithm. The other two constraints (the *WMA* constraint and the *F2B* constraint) are also implicitly addressed in this algorithm, as they are factored into *MaxMarginalWelfare*. On a smaller scale, members of a group will also arrange themselves in a manner that maximizes their total utility after the algorithm is run. Given that groups are limited to a maximum size of three, this step is trivial.

5 Analysis

5.1 Approximation Guarantee

In this subsection, we will provide a high-level overview of why this algorithm (Feige and Izsak [1]) provides a $\frac{1}{d+2}$ approximation guarantee, where d is the maximum supermodular dependency degree.

Let OPT be the optimal solution with value **opt** and let $Approx$ be the approximation we reach in this algorithm with value **appx**. For iteration i of the loop in line 2 of the algorithm, let $Approx_i$ be the allocations made at the first i iterations and OPT_i be the allocations made by OPT for the items that have not been allocated yet. Let $Hybrid_i = Approx_i \cup OPT_i$ be a hybrid solution. Let $Hybrid_i^e$ and OPT_i^e be the items allocated to $Hybrid_i$ and OPT_i to an entity $e \in E$, respectively. Feige and Izsak [1] first prove the following lemma:

Lemma 5.1. *Let i be an iteration of the loop in line 2 of the algorithm and let e^* be the entity to whom items are allocated in that given iteration. Then,*

$$v_{e^*}(Approx_i^{e^*}) - v_{e^*}(Approx_{i-1}^{e^*}) \geq \tfrac{1}{d+2} \textstyle\sum_{e=1}^{n} v_e(OPT_{i-1}^e | Approx_{i-1}^e) - v_e(OPT_i^e | Approx_i^e).$$

Let $\delta v_{e^*} = v_{e^*}(Approx_i^{e^*}) - v_{e^*}(Approx_{i-1}^{e^*})$. They go on to prove that

1. For an item allocated to some $e' \neq e^*$ in OPT, the loss of the value of e' for not getting the item is at most δv_{e^*}.
2. Having received items in iteration i, the loss in marginal value of future items for e^* is at most δv_{e^*}.

Since at most $d+1$ items are allocated at each iteration, the first point only "contributes" to the "damage" up to $(d+1)\cdot\delta v_{e^*}$. The second point "contributes" up to another δv_{e^*}, and for any other player $Hybrid_{i-1} = Hybrid_i$. Let $s_1, ..., s_{d'}$ be the seats allocated at iteration i and let E' be the set of entities $e' \neq e^*$ such that at least one of the seats $s_1, ..., s_{d'}$ is allocated to e' in OPT_{i-1}. Let $e' \in E'$ and let $\hat{s}_1, ..., \hat{s}_{d''} \in s_1, ..., s_{d'}$ be all the items of $s_1, ..., s_{d'}$ allocated to e' in OPT_{i-1}. Then, in the formal calculations of this lemma, the authors find that

$$v_{e'}(OPT_{i-1}^{e'} | Approx_{i-1}^{e'}) - v_{e'}(OPT_i^{e'} | Approx_i^{e'}) \leq d'' \cdot v_{e^*}(Approx_i^{e^*}) - v_{e^*}(Approx_{i-1}^{e^*}).$$

Since there are only $d'' \leq d+1$ items allocated and since for every player $e \notin E' \cup \{e^*\}, Hybrid_i^e = Hybrid_{i-1}^e$, they find that

$$(d+1)\cdot(v_{e^*}(Approx_i^{e^*}) - v_{e^*}(Approx_{i-1}^{e^*})) \geq \textstyle\sum_{e\in E\setminus e^*} v_e(OPT_{i-1}^e | Approx_{i-1}^e) - v_e(OPT_i^e | Approx_i^e).$$

By monotonicity $v_e(Hybrid_i^{e^*}) \geq v_e(Hybrid_{i-1}^{e^*})$, they find that

$$v_{e^*}(OPT_{i-1}^{e^*} | Approx_{i-1}^{e^*}) - v_{e^*}(OPT_i^{e^*} | Approx_i^{e^*}) \leq v_{e^*}(Approx_i^{e^*}) - v_{e^*}(Approx_{i-1}^{e^*}),$$

proving the lemma. They use this lemma to prove the following approximation guarantee:

Theorem 5.1. *The welfare maximization problem with supermodular degree at most* d *admits a polynomial time greedy* $\frac{1}{d+2}$*-approximation algorithm.*

To prove the theorem, let x_i be the profit of Algorithm 1 at iteration i. Then,

$$\begin{aligned} \textbf{opt} &= \textstyle\sum_{e=1}^{n} v_e(OPT_0^e | Approx_0^e) \\ &= \textstyle\sum_{e=1}^{n} \sum_{i=0}^{t-1} (v_e(OPT_i^e | Approx_i^e) - v_e(OPT_{i+1}^e | Approx_{i+1}^e)) \\ &\leq (d+2) \cdot \textstyle\sum_{i=1}^{t} x_i = (d+2) \cdot \textbf{appx}, \end{aligned}$$

proving the approximation guarantee.

5.2 Satisfying Constraints

In this algorithm, if a seat s gets allocated to an entity e, it is allocated along with other seats in it's given *dependency set*, $D_e(s)$. In our example with the family of three, seat $1B$ would likely have a dependency set of $1A$ and $1C$. The allocation of a seat along with its dependency set allows us to tackle the group constraint seamlessly. By treating groups of passengers as individual entities of size $|e|$ and setting up their valuation functions such that they receive the highest marginal utility when they receive the $|e|$th seat, we are able to ensure that when we can maximize welfare subject to this groups constraint, as many groups as possible will receive at least $|e|$ seats. In other words, we will satisfy the maximum number of groups subject to their constraints and their willingness to pay in this algorithm.

In terms of the other two constraints, *WMA* and *F2B*, while we do not explicitly satisfy them in the algorithm, they are implicitly met, since the algorithm is maximizing marginal welfare (which takes into account all preferences).

6 Applications

In the airline industry, seats are currently either first-come-first-serve or there exists system that requires passengers to pay to choose their own seat if they do not want a random seat. Ideally, airline companies would be incentivized by ratings or competition with other airlines for customer satisfaction to implement this new model (as it would likely result in customers experiencing higher levels of happiness than they do in the current system in place), where all passengers are able to provide their preferences and are seated according to the arrangement that maximizes their total welfare subject to the aforementioned constraints. In this system, there could be an "opt-out" option if, for example, a given passenger i had a highly specific preference of sitting in a particular seat (e.g. 17A). In this opt-out option, this passenger i could pay a higher fee to pick their seat first, before the algorithm is run. Then, the algorithm would be run normally, just with that seat excluded from S and that entity excluded from in E.

This system would be incompatible with the current class system (first class, business class, economy plus, economy, etc.), but could work on each class individually. For example, if there was a predetermined range of passengers' willingness to pay, entities above a certain threshold could be automatically put in first class. In that case, we can run this algorithm for each class.

7 Related Work

Maximizing welfare subject to constraints is a popular topic in economics and algorithmic game theory. Often, the question of interest involves maximizing *Nash Social Welfare* (Kaneko and Nakamura [6]), as it has some interesting fairness properties such as envy-freeness, as well as a balanced tradeoff between fairness and efficiency. However, much of the work in that field involves maximizing Nash Social Welfare for agents who have *submodular valuations*, including Garg, Kulkarni, and Kulkarni [9] and Garg et al. [2], or *Rado valuations*, including Garg, Husic, and Vegh [5]. The authors of Garg et al. [2], employ a bidding system, in which agents place bids on items in the first part of their approach, which is then followed by a round of local search to maximize the product of agents' welfare. The authors of Garg, Husic, and Vegh [5], on the other hand, approximate Nash Social Welfare by approximating an integer program and rounding. Unfortunately many findings from Garg et al. [2], Garg, Husic, and Vegh [5], and Garg, Kulkarni, and Kulkarni [9] aren't generalizable to our problem due to inherent differences in valuation functions between submodular or Rado valuations and supermodular valuations. Other work, including Cole and Gkatzelis [3], has approached approximating Nash Social Welfare under additive valuations by using fractional allocations. Finally, another subcategory of papers, such as Caragiannis et al. [4], is focused on discussing fairness properties of maximizing Nash Social Welfare. Originally, our goal was also to maximize Nash Social Welfare in the airplane problem, however, there is unfortunately little existing work in maximizing Nash Social Welfare subject to supermodular or superadditive valuations. As a result, that direction is a great extension of this paper.

The paper most related to our question of social welfare maximization subject to supermodular valuations is Feige and Izsak [1]. We utilize their algorithm in a black-box fashion to compute an allocation that results in maximum social welfare in the airplane seating problem. By incorporating the willingness to pay into the group constraint, we turn the utility function of each entity into a submodular function and then apply the greedy algorithm proposed in Feige and Izsak [1] directly. Many other papers use the concept of weights in their algorithms to allow some agents' preferences to "matter" more than those of other agents. In our case, willingness to pay will be equivalent to giving the agents different weights.

8 Acknowledgements

This work would not be possible without the immense guidance, feedback, and insight from my thesis advisor, Professor Hsien-Chih Chang.

References

[1] Uriel Feige and Rani Izsak (2013). *Welfare Maximization and the Supermodular Degree*.

[2] Jugal Garg, Edin Husić, Wenzheng Li, László A. Végh, and Jan Vondrák (2023) *Approximating Nash Social Welfare by Matching and Local Search*.

[3] Richard Cole, Vasilis Gkatzelis (2015). *Approximating the Nash Social Welfare with Indivisible Items*.

[4] Ioannis Caragiannis, David Kurokawa, Herve Moulin, Ariel Procaccia, Nisarg Shah, Junxing Wang (2016). *The Unreasonable Fairness of Maximum Nash Welfare*.

[5] Jugal Garg, Edin Husic, Laszlo Vegh (2020). *Approximating Nash Social Welfare under Rado Valuations*.

[6] Mamoru Kaneko, Kenjiro Nakamura (1979). *The Nash Social Welfare Function*.

[7] Deeparnab Chakrabarty (2022). *A Crash Course on Linear Programs*.

[8] Aida Rahmattalabi, Shahin Jabbari, Himabindu Lakkaraju, Phebe Vayanos, Max Izenberg, Ryan Brown, Eric Rice, Milind Tambe (2021). *Fair Influence Maximization: A Welfare Optimization Approach*.

[9] Jugal Garg, Pooja Kulkarni, Rucha Kulkarni (2023). *Approximating Nash Social Welfare under Submodular Valuations through (Un)Matchings*.