Approximate Solutions for the Capacity Allocation Problem

Thiam Hui Lee* Woonghee Tim Huh† Retsef Levi‡

1 Introduction

In this paper, we develop a non-parametric data-driven approach to one of the classical revenue management problems — the capacity allocation problem under monotonic fare arrivals. We introduce a new class of solutions called *ε-backward accurate policies*, which are solutions that approximately satisfy the conditions for optimality with an error at most a factor of ε, and prove that these policies are near-optimal. That is, for each accuracy level $\alpha > 0$, for all ε sufficiently small, the expected revenue of each ε-backward accurate monotonic policy is at least $(1 - \alpha)$ times the optimal one. We then describe a simple algorithm that uses uncensored independent demand samples to compute an instance from this class of solutions with high probability. Furthermore, we give polynomial uniform bounds (finite sample bounds) on the number of demand samples needed to ensure the expected revenue of the computed solution is, with a given confidence probability, within some specified ratio of that of an optimal solution, where by uniform bounds we mean bounds that hold for *any* demand distribution.

This canonical problem, of which a prototypical setting is deciding how to allocate seats dynamically to different customers with lower paying customers arriving before higher paying ones, was initially worked on by Littlewood [4], Belobaba [1] and Wollmer [7]. Later, Brumelle and McGill [2] derived closed form optimality conditions that depend on "fill event" probabilities, which Robinson [5] generalized to fare class arrivals that are non-monotonic. In those papers, it is assumed that the demand distributions are known. However, in many practical settings, we neither have exact knowledge of the demand distributions nor the ability to compute them tractably, and must therefore make decisions based on what information we can gleam from historical data. A natural question that arises is how one could work around these limitations.

To address this issue, Van Ryzin and McGill [6] and Huh and Rusmevichientong [3]

*IEOR Department, Columbia University, New York, NY; Email: thl2102@columbia.edu
†IEOR Department, Columbia University, New York, NY; Email: th2113@columbia.edu
‡Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA; Email: retsef@mit.edu

developed stochastic gradient ascent algorithms that update the protection levels from selling season to selling season based on realized demand, converging to an optimal solution at a shown asymptotic rate. In contrast to their papers, our work takes a different approach that allows us to say something much stronger than an asymptotic rate of convergence for our algorithm: we find uniform sampling bounds, in fact polynomial ones, explicitly linking desired performance to the number of samples required to guarantee that. In terms of the methodology, we define a class of solutions that satisfy a relaxed version of the optimality conditions of Brumelle and McGill, and find an upper bound for the optimality gap of any solution from that class of solutions. We then describe a sampling-based algorithm to compute a policy within this class with high probability. Finally, we derive an explicit upper bound on the number of samples needed to guarantee the expected revenue using the solution we generate is at least $(1-\alpha)$ times that of the optimal expected revenue with a confidence probability of $(1-\gamma)$, for some arbitrary choice of α and γ. This bound depends only on the choice of α and γ and *do not* depend on the specific demand distribution and is polynomial in the number of demand classes.

In this paper, we assume the availability of uncensored demand samples, and this assumption is realistic as several modern sales channels for flight reservation systems possess the ability to collect on customers' requests, beyond transacted sales data, which gives information on the underlying demand distribution. For example, travel booking websites such as Travelocity or Orbitz that sell flight tickets on the Internet also collate data on the customers' search and browsing history that allows for uncensoring of sales data.

2 The Capacity Allocation Problem

We suppose that a flight has total capacity x_0 to be sold to $M+1$ different classes of demand indexed by $1, \ldots, M+1$, each willing to pay a different fare $f_1, \ldots, f_{M+1}$ and whose quantities are $D_1, \ldots, D_{M+1}$ respectively. We assume that the demand classes (also termed fare classes) are monotonic in their fares, $f_1 \geq f_2 \geq \ldots \geq f_{M+1}$, and that the demands for the fare classes $\mathbf{D} = (D_1, \ldots, D_{M+1})$ are random variables that are independent, have finite expectations, and whose cumulative distribution functions are $F_1, \ldots, F_{M+1}$ respectively. Furthermore, we assume that demand for fare class $k+1$ arrives completely before demand for fare class k, for $k = 1, \ldots, M$. For ease of presentation, we will assume here that the demand distributions are continuous. However, in the full paper we show the same result for generally distributed demands.

It is well known that there exists an optimal policy that is of the form of protection levels. A protection level p_k is the amount of capacity that is reserved exclusively for sales to customers of fare classes 1 through k inclusive, or equivalently that once the remaining amount of capacity is p_k or less we stop selling to customers of fare classes $k+1,\ldots,M$. We will assume $p_0 = 0$, as in any reasonable policy, we will sell all remaining capacity we can in the last stage. Under this type of policy, the evolution of the remaining capacity from stage to stage is as follows: suppose the current stage is k and we have capacity x remaining. In this stage, all demand that will pay f_k for a seat will arrive, to which we will sell capacity up to the point where there is p_{k-1} or less remaining capacity, and so we will generate $f_k \min\{D_k, (x-p_{k-1})^+\}$ revenue in this stage k (where $p_0 := 0$), and begin the next stage $k-1$ with $(x - \min\{D_k, (x-p_{k-1})^+\})$ remaining capacity.

We can describe in a recursive manner the revenue earned by implementing a policy utilizing protection levels $\mathbf{p}^M = (p_0, p_1, \ldots, p_M)$ as follows: let $W_k(\mathbf{p}^{k-1}, x)$ be the revenue generated by classes 1 to k when x seats are available to satisfy these k highest paying classes and the protection policy for the remaining fare classes is $\mathbf{p}^{k-1} = (p_0, p_1, p_2, \ldots, p_{k-1})$. Then

$$
\begin{aligned}
W_k(\mathbf{p}^{k-1}, x) &= \begin{cases} \int_0^{x-p_{k-1}} r_k f_k + W_{k-1}(\mathbf{p}^{k-2}, x) dF_k(r_k) & \\ \quad + \bar{F}_k(x-p_{k-1}) \left((x-p_{k-1}) f_k + W_{k-1}(\mathbf{p}^{k-2}, p_{k-1}) \right), & \text{if } x > p_k \\ W_{k-1}(\mathbf{p}^{k-2}, x), & \text{if } x \le p_{k-1} \end{cases} \\
W_k(\mathbf{p}^{k-1}, 0) &= 0 \\
W_0(\cdot,\cdot) &= 0
\end{aligned}
$$

In our problem, we wish to select an appropriate protection level vector $\mathbf{p}^M$ in order to maximize $W_{M+1}(\mathbf{p}^M, x_0)$. From Brumelle and McGill [2], we can characterize sufficient conditions for a protection level vector to be optimal as follows: if $\mathbf{p}^{*,M} = (p_0^*, p_1^*, \ldots, p_M^*)$ is a protection level vector satisfying

$$
\frac{f_{k+1}}{f_1} = \Pr\left(D_1 > p_1^*, \ldots, \sum_{i=1}^{k} D_i > p_k^* \right)
$$

for all $k = 1, \ldots, M$, then $\mathbf{p}^{*,M}$ is an optimal protection level vector. This characterization of optimality depends on the probabilities of the events $[D_1 > p_1, \ldots, D_1 + \ldots + D_k > p_k]$, which are the events where when utilizing the protection level policy $(p_1, \ldots, p_{k-1})$, the demands for classes 1 to k are sufficient to sell at least p_k amount of seats. We term these events for each k the *k-th fill event*. Additionally, it can be shown that $\mathbf{p}^{*,M}$ is a monotonic policy, in other words $p_1^* \le p_2^* \le \ldots \le p_M^*$.

3 ε-Backward Accurate Monotonic Policies

We will describe and analyze what we call ε-backward accurate policies, where ε is some positive fraction. These are policies that approximately satisfy the optimality condition described earlier with at most an error of ε. More formally, a protection level policy $\hat{p}^M$ is ε-backward accurate if for all $k \leq M$,

$$\frac{f_{k+1}}{f_1} - \varepsilon \leq \Pr\left(D_1 > \hat{p}_1, \ldots, \sum_{i=1}^{k} D_i > \hat{p}_k\right) \leq \frac{f_{k+1}}{f_1} + \varepsilon.$$

We are able to show a bound on the performance of the ε-backward accurate policy relative to the optimal policy, which is one of the main results of this paper:

Theorem 1. *Let $\varepsilon > 0$. Suppose that $\hat{\mathbf{p}}^M$ is ε-backward accurate, $\mathbf{p}^{*,M}$ is optimal and $\varepsilon < f_{M+1}/f_1$. Then*

$$\frac{W_{M+1}(\mathbf{p}^{*,M}, x_0) - W_{M+1}(\hat{\mathbf{p}}^M, x_0)}{W_{M+1}(\mathbf{p}^{*,M}, x_0)} \leq \frac{2\varepsilon f_1^2 M}{f_M(f_{M+1} - \varepsilon f_1)}.$$

We show this by analyzing differences in the expected revenues of the two policies for different sample paths and proving a uniform upper bound for that difference. That upper bound is then divided by an appropriate lower bound we find for the expected revenue of the x_0-optimal policy, which yields the above result.

4 Sampling-based Algorithm for Constructing an ε-Backward Accurate Policy

We now describe an algorithm that computes from sufficient demand samples an (ε, x_0)-BAM policy with high probability. The algorithm does the following:

Suppose we have N demand path samples $\mathbf{d}^1, \ldots, \mathbf{d}^N$, where $\mathbf{d}_i = (d_1^i, d_2^i, \ldots, d_{M+1}^i)$ is the vector of realized demands $d_1^i, \ldots, d_{M+1}^i$ for fare class $1, \ldots, M+1$ respectively in the i-th sample. Without loss of generality, we can assume that N is an integral multiple of M. We initialize the algorithm by approximating the optimal protection level for class 1, $\hat{p}_1$. First, we form an empirical measure $\hat{\mathbb{P}}$ using N/M demand path samples (let us say $\mathbf{d}^1, \ldots, \mathbf{d}^{N/M}$). For example, $\hat{\mathbb{P}}(D_1 > x) = \sum_{i=1}^{N/M} I(d_1^i > x)$, where $I(\cdot)$ is the indicator function. We find an approximate protection level for class 1, $\hat{p}_1$, by

$$\hat{p}_1 := \inf\left\{p_k \geq 0 : \hat{\mathbb{P}}\left(D_1 > p_k\right) \leq \frac{f_2}{f_1}\right\}$$

Now, we perform the following in an iterative fashion from $k = 2$ to $k = M$. Given that we have computed the first through $(k-1)$-th protection levels $\hat{\mathbf{p}}^{k-1} = (\hat{p}_1, \ldots, \hat{p}_{k-1})$, we compute the k-th protection level, $\hat{p}_k$, in the following manner: use N/M uncensored independent demand samples that have not been used in previous iterations to form an empirical distribution of the demands. Then, using protection levels $\hat{\mathbf{p}}^{k-1}$ for the first through $(k-1)$-th protection level, find the quantity for the k-th protection level such that the optimality conditions are satisfied under the empirical distribution and set $\hat{p}_k$ to be that quantity. This iterative step is summarized as follows:

$$\hat{p}_k := \inf \left\{ p_k \geq 0 : \hat{\mathbb{P}} \left(D_1 > \hat{p}_1, \sum_{i=1}^{2} D_i > \hat{p}_2, \ldots, \sum_{i=1}^{k-1} D_i > \hat{p}_{k-1}, \sum_{i=1}^{k} D_i > p_k \right) \leq \frac{f_{k+1}}{f_1} \right\}$$

where $\hat{\mathbb{P}}$ is the empirical measure defined on the uncensored independent samples of demands used in this iteration using demand samples $\mathbf{d}^{(k-1)(N/M)+1}, \mathbf{d}^{(k-1)(N/M)+2}, \ldots, \mathbf{d}^{k(N/M)}$.

We are able to show that the algorithm gives rise to an ε-backward accurate policy with high probability when given sufficiently many samples:

Theorem 2. *Let some $\varepsilon > 0$ and $\delta \in (0,1)$, and $K(\varepsilon, \delta) := \ln(2/\delta)/(2\varepsilon^2)$. Furthermore, let a given set of protection levels $\hat{\mathbf{p}}^{k-1} := (\hat{p}_1, \ldots, \hat{p}_{k-1})$ be ε-backward accurate. If the number of samples used to compute $\hat{p}_k$ is in the iterative step is greater than or equal to $K(\varepsilon, \delta)$, then $\hat{\mathbf{p}}^k$ is ε-backward accurate with probability at least $1 - \delta$, i.e.*

$$\Pr \left(\bigcap_{i=1}^{k} \left[\frac{f_{i+1}}{f_1} - \varepsilon \leq \Pr \left(D_1 > \hat{p}_1, \sum_{m=1}^{2} D_m > \hat{p}_2, \ldots, \sum_{m=1}^{i} D_m > \hat{p}_i \right) \leq \frac{f_{k+1}}{f_1} + \varepsilon \right] \right) \geq 1 - \delta.$$

By applying some concentration inequalities and properties of ε-backward accurate solutions to this result, we can obtain the second main result of this paper: a bound on the number of demand samples needed by this algorithm to guarantee it outputs a solution of a given quality:

Theorem 3. *Suppose we wish to construct a protection level policy which is $(1-\alpha)$-optimal, i.e. a policy whose expected revenue is at least $(1-\alpha)$ times the expected revenue of an optimal policy, with probability $(1-\gamma)$ where $\alpha > 0$ and $\gamma \in (0,1)$. Then the algorithm outlined in this paper outputs such a policy when supplied with at least $N(\alpha, \gamma)$ uncensored independent demand vector samples, where*

$$N(\alpha, \gamma) = \frac{M \left(2 f_1^2 M + \alpha f_1 f_M\right)^2 \cdot (\ln 2M - \ln \gamma)}{2 \left(\alpha f_M f_{M+1}\right)^2} = \mathcal{O}\left(\left(\frac{f_1}{f_{M+1}} \right)^4 \cdot \frac{M^3 (\log M - \log \gamma)}{\alpha^2} \right)$$

5 Conclusion and Extensions

We defined a new class of policies, the class of ε-backward accurate policies, for the capacity allocation problem. Our approach uses a relaxed version of the classical fill-rate optimization conditions, that is suitable when demand distributions are not known but accessible through samples. We have shown a bound for the performance of any policy drawn from this class of policies relative to an optimal one. And we have proposed an algorithm that, when given demand samples, computes an instance from this class of policies with a high probability. Furthermore, we have given a upper bound for the number of demand path samples needed to guarantee that the algorithm outputs a solution that satisfies a specified performance requirement with at least a given probability.

In this extended abstract, we assumed that the demands follow or can be approximated by continuous distributions. In the full paper, we extend these results to allow for discrete and mixed demand distributions and show the same performance bounds hold. We also show that we can extend these results to when demand classes arrive in a non-monotonic fashion.

References

[1] P. P. Belobaba, *Application of a probabilistic decision model to airline seat inventory control*, Operations Research **37** (1989), 183–197.

[2] S. L. Brumelle and J. I. McGill, *Airline seat allocation with multiple nested fare classes*, Operations Research **41** (1993), no. 1, 127–137.

[3] W. T. Huh and P. Rusmevichientong, *Adaptive capacity allocation with censored demand data: Application of concave umbrella functions*, Working Paper (2006).

[4] K. Littlewood, *Forecasting and control of passenger bookings*, Proceedings of the 12th AGIFORS Symposium, 1972, pp. 95–117.

[5] L. W. Robinson, *Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes*, Operations Research **43** (1995), no. 2, 252–263.

[6] G. Van Ryzin and J. I. McGill, *Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels*, Management Science **46** (2000), no. 6, 760–775.

[7] R. D. Wollmer, *An airline seat management model for a single leg route when lower fare classes book first*, Operations Research **40** (1992), 831–844.