Game-theoretic Analysis of Advance Reservation Services

Eran Simhon
Devision of System Engineering
Boston University
simhon@bu.edu

David Starobinski
Devision of System Engineering
Boston University
staro@bu.edu

Abstract—In many services, such as cloud computing, customers have the option to make reservations in advance. However, little is known about the strategic behavior of customers in such systems. In this paper, we use game theory to analyze several models of time-slotted systems in which customers can choose whether or not making an advance reservation of server resources in future time slots. Since neither the provider nor the customers know in advance how many customers will request service in a given slot, the models are analyzed using Poisson games, with decisions made based on statistical information. The games differ in their payment mechanisms, and the main objective is to find which mechanism yields the highest average profit for the provider. Our analysis shows that the highest profit is achieved when advance reservation fees are charged only from customers that are granted service. Furthermore, informing customers about the availability of free servers prior to their decisions do not affect the provider's profit in that case.

I. INTRODUCTION

In recent years there has been a growing interest in applying advance reservation (AR) in cloud computing [1] and networks [2]. Thus, several software packages for resource management, including Haizea [3], an open source lease management architecture, and IBM Platform Computing Solutions [4], a cluster workload management package, support AR. In both of these packages, an administrator can decide whether or not to enable AR and define the AR pricing scheme.

In most systems supporting AR, customers can choose whether making AR or not. Since the payoff of each customer is affected by decisions of other players, it is natural to analyze the behavior of such systems as strategic games. To our knowledge, such games have not been analyzed before and currently there are no results on customers' behavior in systems that support AR.

The general framework that we consider is a slotted loss system where customers can decide in advance whether reserving future resources for a fee. A customer opting not to make a reservation lowers its chance of finding a server available at the desired time. Customers are not flexible, i.e., they leave the system if they cannot be served at their desired time slots. For each given slot, the number of customers requesting service (referred to as *demand*) is an independent Poisson random variable. The arrival time of a customer prior to the start of a slot, also referred to as its *lead time*, is a positive, continuous i.i.d random variable with general distribution. This model falls under the framework of Poisson games. As shown in [5], only in Poisson games the following properties hold: (1) each player

sees the same probability distribution of the number of players in the game, excluding itself, as the number of players seen by an external observer; (2) if players are randomly ascribed to different types with fixed probabilities, the number of players of each type is also Poisson distributed.

The provider can choose the AR mechanism and fee to maximize its profit. In this paper, we aim to answer two main questions: (1) Should the provider inform the customers prior to their decisions whether there are free resources or not? (2) Should the provider charge the AR fee in advance, even from customers that attempt to make AR but do not get service (which may occur if all servers in the desired slot are already reserved)?

In order to answer those two questions, we analyze three Poisson games that differ in their AR mechanisms. Our main contribution is to show that, at equilibrium, informing customers that free resources are available does not impact the provider's average profit. However, charging an AR fee from all customers attempting AR (i.e., not only those granted service) can only decrease the profit.

Once a model is chosen, another question arises: what is the AR fee that maximizes the provider's profit? The answer to this question seems to be more complicated. As we observe, there is a range of fees, such that choosing a fee within this range leads to multiple equilibria with one of them yielding zero profit.

The rest of the paper is organized as follows. In section II, we briefly review related work. In Section III, we describe in details the three different games. In Section IV, we characterize the different possible equilibria in all of the three games. In Section V, we compare between the profits obtained in the different games and determine which scheme produces the highest profit. In the last section, we conclude the work and suggest several directions for future research.

II. RELATED WORK

Research on queueing systems that support advance reservations started over twenty years ago. Most of the research focuses on performance evaluation and algorithmic aspects of different AR systems. For example, in [6], a scheduling model that supports AR is proposed and several performance metrics are evaluated. in [7], an AR model with flexible time window is suggested and it is shown that this model has a lower blocking probability and a higher utilization than a

model without window. In [8], the authors analyze the effect of AR on the complexity of path selection. In [9], the utilization of models with and without AR are compared. A simulation based comparison between different payment mechanisms is reported in [10]. Algorithms for network routing that support advance channel reservations are proposed in [2]. For a survey on the field, see [11].

The application of game theory to analyze customers' behavior in queues is pioneered in [12]. In this paper the author considers an $M/M/1$ queue where customers observe the queue length and then decide whether to join or bulk. The model introduced in [12] stimulated many follow-up works. An unobservable $M/M/1$ queue, where customers need to decide whether to join or bulk without knowing the queue state, is analyzed in [13]. An observable processor sharing system, where customers decide whether or not to join the system after observing the number of customers within, is analyzed in [14]. In [15], the authors analyze an observable $M/M/1$ queue with priorities, where customers decide on a payment and accordingly priorities are assigned. In [16], the authors analyze an unobservable $M/M/N/N$ system that is initially empty and customers decide whether to join or bulk based on their arrival time. In [17], the authors introduce the concert queueing game, where customers, interested in early service with minimal wait, can choose their arriving time into a system with a specific opening time. For a review of the field see [18].

III. THE MODELS

First we describe the assumptions that are common in all three models:

1) There are N servers.
2) The service time axis is slotted.
3) The *demand*, which represents the number of customers that request service in a specific slot (each customer requests one server), is an independent Poisson distributed random variable D with parameter λ.
4) The customers arrive to the system in a random fashion. Customer i's *lead time* is denoted T_i and it is the time between its arrival and the slot starting time. The random variables T_i are i.i.d and continuous (supported in $[0, \infty)$), with cumulative distribution function (CDF) $F(t)$.
5) Upon arrival, each customer can choose one of two actions: make AR or not make AR, denoted AR and AR' respectively.
6) If the demand for a slot is larger than N, the servers are allocated to the first N customers that made AR. If fewer than N customers made AR, the remaining servers are randomly allocated between the customers that have not made AR.
7) The customers and the provider know the number of servers N and statistical information on the system (i.e., λ and F). However, they do not know the demand and the lead time of the rest of the customers in advance.

-	Make AR		Not make AR	
Model	Served	Not Served	Served	Not served
1 and 3	$1 - C$	0	1	0
2	$1 - C$	$-C$	1	0

TABLE I: Payoff summary.

8) The provider charges a fixed reservation fee denoted by C. All the customers have the same utility U from the service. Without loss of generality, we set $U = 1$.

To illustrate the model, consider the following example: a system with many servers has a slot duration that lasts for one day, starting at 12:00 AM. A customer realizes on Monday 6:00 PM that it will need service on Wednesday. Thus, its lead time is 30 hours. At that point, the customer, knowing the statistical information, the number of servers and the AR fee, needs to decide whether to make AR or not.

The three models differ in their reservation mechanisms as described next:

1) Unobservable model 1: customers do not know if there are free servers at the time of reservation. If a customer makes an AR request, it is then informed whether a server will be allocated at the requested slot. In the first case a reservation fee is charged. In the second case, the customer leaves the system with no gain or cost.
2) Unobservable model 2: customers do not know if there are free servers at the time of reservation. A reservation fee is charged from each customer that makes an AR request.
3) Observable model: customers know at their arrivals (prior to their decision) if there are free servers. A customer that has been informed that there is no free server leaves the system. Otherwise, it has to decide whether to make AR or not.

The possible payoffs are summarized in Table I.

IV. EQUILIBRIA ANALYSIS

A. Classification of the equilibria

We analyze the three models as none-cooperative games where each player (customer) aims to maximize its payoff. In the third model, we can ignore the case where a customer is informed that no server is available, since no decision needs to be made in this case. Since the demand for each slot is an i.i.d random variable, the analysis of a single slot is sufficient for analyzing the game.

Given a lead time t, we set $p = F(t)$ and refer to as the *normalized lead time*. $F(t)$ can be interpreted as the average fraction of customers with lead time smaller than t. For each game, we define a strategy function $\sigma : p \longrightarrow \tau \in [0, 1]$, which defines the probability that a tagged customer with normalized lead time $p \in [0, 1]$ makes AR. Through conditioning, given that there are k other customers with normalized lead times $p_1...p_k$ and that they all follow strategy function σ, the tagged customer can find the probability of getting service (we denote

that event by S) for each action it chooses. Therefore, its probability to get service, when choosing action $\alpha \in \{AR, AR'\}$ and given σ is:

$$\mathbb{P}(S|\sigma, p, \alpha) = \mathbb{P}(D < N) +$$
$$\sum_{k=N}^{\infty} \mathbb{P}(D = k) \int_{p_1=0}^{1} \cdots \int_{p_k=0}^{1} \mathbb{P}(S|p, \alpha, \sigma, k, p_1..p_k) \mathrm{d}p_1 \cdots \mathrm{d}p_k. \quad (1)$$

The first term in (1) is the probability that less than N customers arrive (beside the tagged customer), in this case the probability to get service is one. The second term is the weighted sum of the probabilities of getting service when the demand is greater than N. In this case the probability to get service depends on the actions of other customers and their lead times.

Given a model and strategy function, one can find the expected payoff for each action, denoted $U_\sigma(p, \alpha)$, by multiplying $\mathbb{P}(S|\sigma, p, \alpha)$ and $1 - \mathbb{P}(S|\sigma, p, \alpha)$ with the relevant payoffs, as summarized in Table I. For example, in the second model:

$$U_\sigma(p, AR) = \mathbb{P}(S|\sigma, t, AR) \cdot (1 - C)$$
$$+ (1 - \mathbb{P}(S|\sigma, p, AR) \cdot (-C) \quad (2)$$

and

$$U_\sigma(p, AR') = \mathbb{P}(S|\sigma, p, AR') \cdot 1$$
$$+ (1 - \mathbb{P}(S|\sigma, p, AR') \cdot 0. \quad (3)$$

In equilibrium, each customer choses an action that maximizes its expected payoff. Thus, we define an equilibrium strategy (i.e., a strategy that leads to equilibrium) as follows:

Definition 1. *Strategy σ is an equilibrium strategy if the following holds for any $p \in [0, 1]$:*

1) *If $\sigma(p) = 0$ then $U_\sigma(p, AR) \leq U_\sigma(p, AR')$.*
2) *If $0 < \sigma(p) < 1$ then $U_\sigma(p, AR) = U_\sigma(p, AR')$.*
3) *If $\sigma(p) = 1$ then $U_\sigma(p, AR) \geq U_\sigma(p, AR')$.*

Next we define three type of equilibria followed by a lemma stating that those are the only possible equilibria.

Definition 2. *A none-make-AR equilibrium has the following strategy function:*

$$\sigma(p) = 0, \qquad \forall p \in [0, 1].$$

Definition 3. *Given $p_e \in (0, 1)$, a threshold equilibrium p_e has the following strategy function:*

$$\sigma(p) = \begin{cases} 1 & \text{if } p > p_e \\ 0 & \text{if } p \leq p_e. \end{cases}$$

Definition 4. *An all-make-AR equilibrium has the following strategy function:*

$$\sigma(p) = 1, \qquad \forall p \in [0, 1].$$

Lemma 1. *The only possible equilibria are none-make-AR, threshold and all-make-AR equilibria.*

Proof: In the two unobservable games, consider a strategy function σ and a customer with normalized lead time p. If the customer does not make AR, its payoff does not depend on p and $U_\sigma(p, AR') = a$ where a is some constant. On the other hand, from assumption 6 in Section III, it can be inferred that the payoff $U_\sigma(p, AR)$ of a customer that makes AR is a non-decreasing function of p. Based on this observation, we next show that in any scenario, an equilibrium must be one of the three equilibria stated in the lemma.

If $U_\sigma(p, AR') > U_\sigma(p, AR)$ for all $p \in [0, 1]$, the given strategy is an equilibrium strategy only if it is *none-make-AR*. If $U_\sigma(p, AR') < U_\sigma(p, AR)$ for all $p \in [0, 1]$, the strategy is an equilibrium strategy only if it is *all-make-AR*.

Finally, if the two payoff functions intersect, they can either intersect at a single point p_0 or on an interval $[p_1, p_2]$. In the first case, $U_\sigma(p, AR') > U_\sigma(p, AR)$ for all $p < p_0$ and $U_\sigma(p, AR') < U_\sigma(p, AR)$ for all $p > p_0$ therefore, in this case, σ is an equilibrium strategy only if it is a threshold strategy p_0. In the second case, $U_\sigma(p, AR) = a$ for all $p \in [p_1, p_2]$ which can only happen if $\sigma(p) = 0$ for all $p \in [p_1, p_2]$. Otherwise, $U_\sigma(p, AR)$ would not have a fixed value in that range (we ignore the case of $\sigma(p) \neq 0$ over a measure zero subset of $[p_1, p_2]$, since the probability that a customer will have a normalized lead time within this subset is zero). Since $U_\sigma(p, AR') > U_\sigma(p, AR)$ for all $p < p_1$ and $U_\sigma(p, AR') < U_\sigma(p, AR)$ for all $p > p_2$, the given strategy is an equilibrium strategy only if it is the threshold strategy p_2 (or *none-make-AR* if $p_2 = 1$).

In the third model, consider a customer with normalized lead time p and being informed that there are free servers. The probability that these servers will be allocated to other customers cannot decrease as p increases. Therefore, if the customer does not make AR, its expected payoff is a continuous non-increasing function of the normalized lead time while the payoff of making AR is always $1 - C$. Thus, the same conclusion can be drawn for this model. ■

B. Equilibria structure

Based on the observations we made so far, we next find the equilibria structure for each model separately and we show that different ranges of fees lead to different equilibria. The results are summarized in the following theorem.

Theorem 1. *For each model $i = 1, 2, 3$, there exist quantities $\underline{C}$ and $\overline{C}_i \geq \underline{C}$, such that the following holds:*

- *If $C \leq 0$, all-make-AR is the unique equilibrium.*
- *If $0 < C < \underline{C}$, there is at least one threshold equilibrium.*
- *If $\underline{C} \leq C < \overline{C}_i$, there is a none-make-AR equilibrium and at least two threshold equilibria.*
- *If $C = \overline{C}_i$, there is a none-make-AR equilibrium and at least one threshold equilibrium.*
- *If $C > \overline{C}_i$, none-make-AR is the unique equilibrium.*

1) Unobservable model 1: For each type of equilibrium, we find the range of fees in which this equilibrium may occur. We also show that for any fee at least one equilibrium exists.

All-make-AR equilibrium. In the case $C \leq 0$, making AR is obviously a dominant strategy and therefore *all-make-AR* is the unique equilibrium. However, for any $C > 0$, *all-make-AR* equilibrium can not be formed. We show it by assuming that all customers make AR and there is a customer with lead time ϵ. As ϵ decreases, the probability that it is the last customer to arrive increases. For any $\delta > 0$, there exists $\epsilon > 0$ such that the probability that a customer with lead time ϵ is the last to arrive is greater than $1 - \delta$. When the probability to be the last customer to arrive tends to one, the difference between the probability to get service with and without AR tends to zero. Therefore, for any value of $C > 0$, there is a small enough ϵ such that a customer with ϵ lead time is better off not making AR and *all-make-AR* is not an equilibrium.

Threshold equilibrium. In the threshold equilibrium p_e, the number of customers making AR is Poisson distributed with parameter $\lambda(1 - p_e)$ and the number of customers not making AR is Poisson distributed with parameter λp_e. Furthermore, these two random variables are independent. This follows from the second property of Poisson games mentioned in the introduction. We denote the probability of a customer with normalized lead time $p_e \in [0, 1]$ to get service if making AR (respectively not making AR) as $p_{AR}(p_e)$ (respectively $p_{AR'}(p_e)$). The first expression is:

$$p_{AR}(p_e) = e^{-\lambda(1-p_e)} \sum_{k=0}^{N-1} \frac{(\lambda(1-p_e))^k}{k!}, \tag{4}$$

which is the probability that fewer than N customers (beside the one that arrive at the threshold) make AR. The second expressions is:

$$p_{AR'}(p_e) = e^{-\lambda(1-p_e)} \sum_{k=0}^{N-1} \frac{(\lambda(1-p_e))^k}{k!}$$
$$\cdot \left(e^{-\lambda p_e} \sum_{j=0}^{N-k-1} \frac{(\lambda p_e)^j}{j!} + e^{-\lambda p_e} \sum_{j=N-k}^{\infty} \frac{(\lambda p_e)^j}{j!} \left(\frac{N-k}{j+1} \right) \right), \tag{5}$$

which is the sum of the probabilities that k customers ($k < N$) make AR, each probability being multiplied by two terms. The first one, covers the case of total demand smaller or equal to N, In the second term, the demand exceeds N and the probability to get service depends on the proportion of the number of unreserved servers and the number of customers that did not make AR. After some algebra we get:

$$p_{AR'}(p_e) = e^{-\lambda} \sum_{k=0}^{N-1} \frac{(\lambda(1-p_e))^k}{k!}$$
$$\cdot \left(\sum_{j=0}^{N-k-1} \frac{(\lambda p_e)^j}{j!} + \left(e^{\lambda p_e} - \sum_{j=0}^{N-k} \frac{(\lambda p_e)^j}{j!} \right) \left(\frac{N-k}{j+1} \right) \right). \tag{6}$$

In a threshold equilibrium p_e, a customer with normalized lead time p_e is indifferent between the two actions. Thus, a

threshold strategy p_e is an equilibrium if and only if

$$(1 - C) \cdot p_{AR}(p_e) = p_{AR'}(p_e), \tag{7}$$

where the left hand side of Eq. (7) is the expected payoff of making AR, and the right hand side of Eq. (7) is the expected payoff of not making AR. Using Eq. (7), we can express the fee as a function of the threshold:

$$C_1(p_e) = 1 - \frac{p_{AR'}(p_e)}{p_{AR}(p_e)}, \tag{8}$$

which is defined and continuous in the range $[0, 1]$. We observe that $C_1(0) = 0$ and that $C_1(p_e) > 0$ for any $p_e > 0$. By defining

$$\overline{C}_1 = \max_{0 \leq p_e \leq 1} C_1(p_e) \tag{9}$$

we can conclude that for any $C \in (0, \overline{C}_1]$, there is at least one value of p_e such that Eq. (7) holds and therefore there is at least one threshold equilibrium (Note that if $\overline{C}_1 = C_1(1)$ and $C = \overline{C}_1$, then by definition the equilibrium is a *none-make-AR* equilibrium). Furthermore, for any $C \notin (0, \overline{C}_1]$, there is no threshold equilibrium.

None-make-AR equilibrium. If none of the customers makes AR, they all have the same expected payoff $p_{AR'}(1)$. A customer that deviates gets service with probability 1 and its payoff is $1 - C$. Thus, if the provider chooses a fee such that $1 - C \leq p_{AR'}(1)$, then none of the customers will have an incentive to deviate. On the other hand, if $1 - C > p_{AR'}(1)$, then all the customers will have an incentive to deviate. By defining $\underline{C} = 1 - p_{AR'}(1)$, we can conclude that if and only if $C \geq \underline{C}$, *none-make-AR* is an equilibrium.

By definition $\overline{C}_1 \geq \underline{C}$ and therefore we have shown that for any value of C, at least one equilibrium exists. Furthermore, if $\overline{C}_1 > \underline{C}$, for any $C \in (\underline{C}, \overline{C}_1)$, the equation $C = C_1(p_e)$ must have at least two solutions. Therefore, there are at least two different threshold equilibria in this range.

2) Unobservable model 2: In this section, we show that the second game has the same equilibria structure as the first one but with different ranges.

All-make-AR equilibrium. By the same reasoning as in the first model, if and only if $C \leq 0$, *all-make-AR* is the unique equilibrium.

Threshold equilibrium. Given a threshold, the probability to get service with and without making AR is equal to the probabilities in the first model. Thus, the functions p_{AR} and $p_{AR'}$ can be used in the analysis of this model. As in the first game, when all behave in accordance to a threshold strategy p_e, a customer that arrives at the threshold is indifferent between the two strategies. Thus,

$$p_{AR}(p_e) - C = p_{AR'}(p_e), \tag{10}$$

where the left hand side of (10) is the payoff of a customer that arrives at the threshold and makes AR, while the right hand side of (10) is the payoff of not making AR. In this
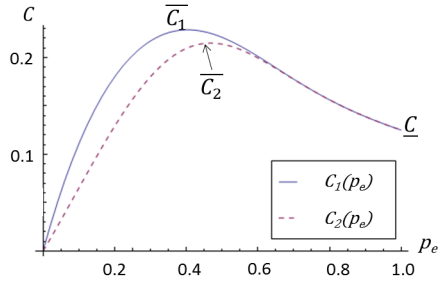
Fig. 1: An example with $N = \lambda = 10$.

model, the fee as a function of the threshold is:

$$C_2(p_e) = p_{AR}(p_e) - p_{AR'}(p_e) \qquad (11)$$

and $\overline{C}_2$ is defined as:

$$\overline{C}_2 = \max_{0 \leq p_e \leq 1} C_2(p_e). \qquad (12)$$

As is the first model, $C_2(0) = 0$ and $C_2(p_e) > 0$ for any $p_e > 0$. Thus, there is a threshold equilibrium if and only if $C \in (0, \overline{C}_2]$.

***None-make-AR* equilibrium** If none of the customers makes AR, the payoffs of making and not making AR are $1 - C$ and $p_{AR'}(1)$ respectively. Which are the same as in the first model. Therefore the range of fees that can lead to a *none-make-AR* is the same as before.

In conclusion, the only difference between the analysis of the two games is that $\overline{C}_1$ may have different value than $\overline{C}_2$. By the definitions of $\overline{C}_1$ and $\overline{C}_2$ it can be observed that $\overline{C}_1 \geq \overline{C}_2$ as illustrated in Fig. 1.

3) Observable model: In this model, customers make their decisions not only based on the statistical information but also based on the knowledge that servers, at the desired slot, are currently available. Next we show that the additional information does not affect their decisions and each fee leads to the same equilibria as in the first unobservable model.

***All-make-AR* equilibrium**. By the same reasoning as in the first model, if and only if $C \leq 0$ *All-make-AR* is the unique equilibrium.

***Threshold* equilibrium**. First we denote the event of being informed that there are free servers as I_1. If all the customers that see this event follow the threshold strategy p_e, then the payoff of a customer arriving at the threshold (and seeing I_1) is $1 - C$ if it makes AR and $\mathbb{P}(S|p_e, I_1, AR')$ otherwise. Next we show that

$$\mathbb{P}(S|p_e, I_1, AR') = \frac{p_{AR'}(p_e)}{p_{AR}(p_e)} \qquad (13)$$

and therefore, Eq. (7) is also valid for the third model. By conditioning on I_1 we get:

$$\mathbb{P}(S|p_e, I_1, AR') = \frac{\mathbb{P}(S \cap I_1|p_e, AR')}{\mathbb{P}(I_1|p_e, AR')}. \qquad (14)$$

Since a customer cannot get service when observing no free server, the nominator $\mathbb{P}(S \cap I_1|p_e, AR')$ is equal to $\mathbb{P}(S|p_e, AR')$. Given a threshold strategy p_e, if a customer

does not know whether there are free servers, its probability to get service is the same as in the two other models. Thus, $\mathbb{P}(S|p_e, AR') = p_{AR'}(p_e)$.

The denominator $\mathbb{P}(I_1|p_e, AR')$ is the probability that a customer that arrives at the threshold and does not make AR, sees the event I_1. This can be rephrased as the probability that less than N servers have been reserved (the fact that the customer does not make AR is irrelevant). This in turn can be rephrased as the probability to get service when making AR exactly at the threshold point without knowing if there are free servers, which is the definition of $p_{AR}(p_e)$. Thus, $\mathbb{P}(I_1|p_e, AR') = p_{AR}(p_e)$.

***None-make-AR* equilibrium**. If none of the customers makes AR, then the customers already know that there are free servers. Thus, being informed by the provider that there are free servers has no impact on customers decisions.

By noticing that the profit of the provider is defined in the same way in both models, that is, the number of customers that make AR and get service multiplied by C, we can infer the following theorem:

Theorem 2. *For any value of C, the first and the third models have the same equilibria and each equilibrium yields the same expected profit for the provider.*

V. COMPARISON OF MODELS

In this section we compare the provider profit in the two unobservable models. Since a fee may lead to more than one equilibrium, we assume that, in this case, the provider can choose an equilibrium. So, instead of choosing the fee, we assume, in this section, that the provider can choose the value of p_e.

We denote the number of customers that make AR requests as M. In the first model, the expected profit as a function of p_e is the expectation of the minimum between N and M, multiplied by the fee:

$$R_1(p_e) = \mathbb{E}[\min(M, N)]C_1(p_e). \qquad (15)$$

While in the second model, the provider expected profit is the expected number of customers that make AR requests multiplied by the fee:

$$R_2(p_e) = \lambda(1 - p_e)C_2(p_e). \qquad (16)$$

Ignoring the trivial cases of $p_e = 0$ and $p_e = 1$ which yield zero profit in both models, we next state the following result:

Theorem 3. *In any threshold equilibrium the first model yields greater expected profit than the second model.*

Proof: We need to show that for every value of $p_e \in (0, 1)$ the following holds true:

$$\frac{R_1(p_e)}{R_2(p_e)} > 1. \qquad (17)$$

From Eq. (8) and (11), we can infer that:

$$\frac{C_1(p_e)}{C_2(p_e)} = \frac{1}{p_{AR}(p_e)}. \qquad (18)$$

Thus, by Eqs. (15), (16) and (18) we obtain

$$\frac{R_1(p_e)}{R_2(p_e)} = \frac{\mathbb{E}[\min(M,N)]}{\lambda(1-p_e)p_{AR}(p_e)}. \tag{19}$$

For convenience we denote $\lambda(1-p_e)$ by $\bar{\lambda}$ and get:

$$\mathbb{E}[\min(M,N)] = \sum_{k=0}^{N} \mathbb{P}(M=k)k + \sum_{k=N+1}^{\infty} \mathbb{P}(M=k)N$$

$$= e^{-\bar{\lambda}} \sum_{k=0}^{N} \frac{\bar{\lambda}^k}{k!}k + e^{-\bar{\lambda}} \sum_{k=N+1}^{\infty} \frac{\bar{\lambda}^k}{k!}N. \tag{20}$$

Substituting Eq. (20) and (4) into Eq. (19), we now need to show that:

$$\frac{e^{-\bar{\lambda}} \sum_{k=0}^{N-1} \frac{\bar{\lambda}^k}{k!}k + e^{-\bar{\lambda}} \sum_{k=N}^{\infty} \frac{\bar{\lambda}^k}{k!}N}{\bar{\lambda}e^{-\bar{\lambda}} \sum_{k=0}^{N-1} \frac{\bar{\lambda}^k}{k!}} > 1 \tag{21}$$

or:

$$\sum_{k=0}^{N-1} \frac{\bar{\lambda}^k}{k!}k + \sum_{k=N}^{\infty} \frac{\bar{\lambda}^k}{k!}N - \bar{\lambda} \sum_{k=0}^{N-1} \frac{\bar{\lambda}^k}{k!} > 0. \tag{22}$$

With some arithmetic manipulations we obtain:

$$\sum_{k=0}^{N-1} \frac{\bar{\lambda}^k k}{k!} + \sum_{k=N}^{\infty} \frac{\bar{\lambda}^k}{k!}N - \bar{\lambda} \sum_{k=0}^{N-1} \frac{\bar{\lambda}^k}{k!}$$

$$= \sum_{k=1}^{N-1} \frac{\lambda^k}{(k-1)!} - \sum_{k=1}^{N} \frac{\lambda^k}{(k-1)!} + \sum_{k=N}^{\infty} \frac{\bar{\lambda}^k}{k!}N$$

$$= -\frac{\bar{\lambda}^N}{N!}N + N\left(e^{\bar{\lambda}} - \sum_{k=0}^{N-1} \frac{\bar{\lambda}^k}{k!}\right)$$

$$= N\left(e^{\bar{\lambda}} - \sum_{k=0}^{N} \frac{\bar{\lambda}^k}{k!}\right) > 0. \tag{23}$$

∎

VI. Conclusions and Future Research

We modeled a slotted system of reusable resources, where customers are allowed to make advance reservations while AR bears an additional fee. We considered three different AR mechanisms and analyzed each one as a Poisson game, where customers differ by their lead time.

We showed that in all three games, for any fee, at least one of the following equilibria exists: all the customers make AR, none of the customers makes AR or a threshold equilibrium where only customers that arrive early enough make AR. We also showed that a profit can be obtained only in the threshold equilibrium and, by choosing the right fee, any threshold equilibrium can be achieved. However, there may be other equilibria for that fee.

As this research is the first to study a stochastic AR system as a game, many directions for future research are possible. To name a few of them: a model where customers have some cancellation probability; generalized versions where customers need more than one server in a slot or more than one slot; and a model with different classes of customers.

References

[1] B. Sotomayor, "Haizea and private clouds,blog.dsa-research.org blog," 2009.

[2] R. Cohen, N. Fazlollahi, and D. Starobinski, "Path switching and grading algorithms for advance channel reservation architectures," *Networking, IEEE/ACM Transactions on*, vol. 17, no. 5, pp. 1684–1695, 2009.

[3] "Haizea," http://haizea.cs.uchicago.edu.

[4] "IBM platform computing solutions," http://www.redbooks.ibm.com/redbooks/pdfs/sg248073.pdf.

[5] R. B. Myerson, "Population uncertainty and poisson games," *International Journal of Game Theory*, vol. 27, no. 3, pp. 375–392, 1998.

[6] W. Smith, I. Foster, and V. Taylor, "Scheduling with advanced reservations," in *Parallel and Distributed Processing Symposium, 2000. IPDPS 2000. Proceedings. 14th International.* IEEE, 2000, pp. 127–132.

[7] N. R. Kaushik, S. M. Figueira, and S. A. Chiappari, "Flexible time-windows for advance reservation scheduling," in *Modeling, Analysis, and Simulation of Computer and Telecommunication Systems, 2006. MASCOTS 2006. 14th IEEE International Symposium on.* IEEE, 2006, pp. 218–225.

[8] R. A. Guérin and A. Orda, "Networks with advance reservations: The routing perspective," in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1. IEEE, 2000, pp. 118–127.

[9] J. T. Virtamo, "A model of reservation systems," *Communications, IEEE Transactions on*, vol. 40, no. 1, pp. 109–118, 1992.

[10] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging it platforms: Vision, hype, and reality for delivering computing as the 5th utility," *Future Generation computer systems*, vol. 25, no. 6, pp. 599–616, 2009.

[11] N. Charbonneau and V. Vokkarane, "A survey of advance reservation routing and wavelength assignment in wavelength-routed wdm networks," 2011.

[12] P. Naor, "The regulation of queue size by levying tolls," *Econometrica: journal of the Econometric Society*, pp. 15–24, 1969.

[13] N. M. Edelson and D. K. Hilderbrand, "Congestion tolls for poisson queuing processes," *Econometrica: Journal of the Econometric Society*, pp. 81–92, 1975.

[14] E. Altman and N. Shimkin, "Individual equilibrium and learning in processor sharing systems," *Operations Research*, vol. 46, no. 6, pp. 776–784, 1998.

[15] K. Balachandran, "Purchasing priorities in queues," *Management Science*, vol. 18, no. 5-Part-1, pp. 319–326, 1972.

[16] M. Haviv, O. Kella, and Y. Kerner, "Equilibrium strategies in queues based on time or index of arrival," *Probability in the Engineering and Informational Sciences*, vol. 24, no. 1, p. 13, 2010.

[17] R. Jain, S. Juneja, and N. Shimkin, "The concert queueing game: to wait or to be late," *Discrete Event Dynamic Systems*, vol. 21, no. 1, pp. 103–138, 2011.

[18] R. J. Hassin and M. Haviv, *To Queue or Not to Queue: Equilibrium Behaviour in Queueing Systems.* Kluwer Academic Pub, 2003, vol. 59.