Flexibility can hurt dynamic matching system performance

Arnaud Cadas[*] Josu Doncel[†] Jean-Michel Fourneau[‡] Ana Bušić[*]

ABSTRACT

We study the performance of stochastic matching models with general compatibility graphs. Items of different classes arrive to the system according to independent Poisson processes. Upon arrival, an item is matched with a compatible item according to the First Come First Matched discipline and both items leave the system immediately. If there are no compatible items, the new arrival joins the queue of unmatched items of the same class. Compatibilities between item classes are defined by a connected graph, where nodes represent the classes of items and the edges the compatibilities between item classes. We show that such a model may exhibit a non intuitive behavior: increasing the matching flexibility by adding new edges in the matching graph may lead to a larger average population at the steady state. This performance paradox can be viewed as an analog of the Braess paradox. We show sufficient conditions for the existence or non-existence of this paradox. This performance paradox in matching models appears when specific independent sets are in saturation, i.e., the system is close to the stability condition.

Keywords

Braess paradox, matching models, performance evaluation

1. INTRODUCTION

We consider the stochastic matching model with multiple classes of items. Compatibilities between item classes are defined by a connected graph, where nodes represent the classes of items and the edges the compatibilities between item classes. Items of different classes arrive to the system according to independent Poisson processes. Upon arrival, an item joins a queue if there are no compatible items present in the system. If there are compatible items, it is instantaneously matched with the oldest compatible item, i.e according to the First Come First Matched (FCFM) policy, and both items leave the system.

An example of the compatibility graph is given in Fig. 1. There are four classes of items. Items of class 3 and 4 can be matched with all the other classes, whereas items of class 1 and class 2 can only be matched with items of class 3 and class 4. After uniformization, the dynamics of this matching model can be described by a discrete time Markov chain. The state $w = w_1 \cdots w_q$ of this Markov chain is a word of remaining items in the system, in increasing order of their arrival time. Each letter w_i is the class of the i-th remaining item. Assume that the system is empty at the initial time and that the trajectory of arrivals is the following: 1, 2, 1, 3 and 4. Then, the trajectory of the Markov chain is 1, 12, 121, 21 and 1.

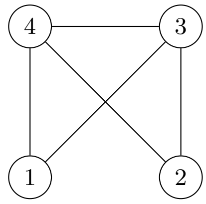

Figure 1: Example of a compatibility graph.

We investigate the existence of a performance paradox in matching models which can be viewed as an analog of the Braess paradox. We study the influence on the mean number of items by adding an edge in a compatibility graph. Adding an edge to the compatibility graph leads to a bigger set of possible matchings. Therefore, it is clear that, when we add an edge to the compatibility graph, the performance of the system will always improve under an optimal policy. However, this is less obvious for other matching policies. We study the conditions under which the mean number of unmatched items under FCFM policy decreases when we add an edge to the compatibility graph. The FCFM assumption allows us to use the result of [15] that shows that the steady-state distribution of remaining items admits a product form solution when the matching discipline is FCFM.

The main contributions of our paper are the following:

- We establish a closed-form expression for the expected total number of items remaining in the system. Specifically, we show that it can be written as a finite sum over all independent sets.
- We give sufficient conditions for the existence or the non-existence of a performance paradox in matching models under an asymptotic assumption similar to the heavy-traffic assumptions in queueing systems [13]. We assume that the sum of the arrival rates of an independent set, which we call saturated, grows to its capacity

[*]INRIA, Paris, France. DI ENS, CNRS, PSL Research University, Paris, France.
[†]University of the Basque Country, UPV-EHU, Leioa, Spain.
[‡]DAVID, UVSQ, Université Paris-Saclay, Versailles, France.

(the sum of the arrival rates of its neighbors), while the other independent sets stay strictly within their capacity. We prove that a performance paradox exists when the saturated independent set has both nodes of the added edge as neighbors. We also prove that a performance paradox does not exist when the saturated independent set contains at least one of the nodes of the added edge.

The rest of the article is organized as follows: we present the related work on matching models, FCFM policy and Braess paradox in Section 2. In Section 3, we describe the model and introduce notation. The closed-form expression for the expected total number of unmatched items is derived in Section 4. Sufficient conditions for the existence or non-existence of the performance paradox are given in Section 5. Finally, we provide conclusions and discuss future work in Section 6.

2. RELATED WORK

The stochastic matching model studied in this paper was first introduced in [14]. It is also known in the literature as the general stochastic matching model [15], where the term general is a reference to the compatibility graph. In fact, this matching model is always used with a non-bipartite compatibility graph as this is a necessary stability condition [14]. Mairesse and Moyal [14] also established that for compatibility graphs that are separable of order $p \geq 3$ the stability region is maximal for all matching policies, and that the policy "Match the Longest" always has a maximal stability region for any non-bipartite compatibility graph. Moyal et al. showed in [15] that the FCFM policy also has maximal stability region and they established the product form of the stationary distribution using a reversibility property related to the one proposed in [1] for a bipartite compatibility graph variant of the stochastic matching model. An extension of FCFM stochastic matching to multigraphs was recently studied in [5].

The stochastic matching model is closely related to the First Come First Served (FCFS) infinite bipartite matching model, proposed by Caldentey et al. [9]. Caldentey et al. use the term FCFS to describe their matching policy but we prefer the term FCFM which emphasizes that there is no service time (matchings are instantaneous). Their matching model has a bipartite compatibility graph: The nodes, representing the classes of items, can be separated in two disjoint sets which we will call the demand nodes and the supply nodes. They consider an infinite sequence of demand items, independent and identically distributed according to a probability distribution on the demand classes, and an independent infinite sequence of supply items, independent and identically distributed according to a probability distribution on the supply classes. The two infinite sequences of items are matched on a FCFM basis. They provide three different Markov chains describing FCFM matchings. The closest to our current paper is the chain that at time n considers the first n demand and first n supply items. The state of the system is described by the word of unmatched items after applying FCFM matching policy. This Markov chain also models the system with stochastic arrivals where, instead of arriving one by one, new items arrive by pairs of one demand item and one supply item. Caldentey et al. [9] provided the necessary stability conditions and analyzed in detail a few particular bipartite compatibility graphs. In [3], Adan and Weiss proved that an alternative Markov description from [9], called server by server matchings, has a product form stationary distribution, and derived the exact matching rates between different compatible classes. FCFM infinite bipartite matching model was further studied by Adan et al. [1], that proved the reversibility of a more detailed Markov chain, as well as the product form of the stationary distribution for different Markov chain descriptions introduced in [9]. Stability properties of stochastic bipartite matching models under different matching policies were studied in [7].

Our paper addresses the performance of the stochastic FCFM matching model with general compatibility graph, when the flexibility increases, i.e when an edge is added to the compatibility graph. We show that increasing the flexibility can decrease the overall performance of the system, which is reminiscent of the Braess paradox. For example, consider a carpooling system where compatibilities between classes represent the geographic proximity of the users. Consider now the case where one class of users declares to be willing to walk or drive longer to be eligible to be matched with a geographical further location. This situation corresponds to a new matching system with a compatibility graph with an additional edge. We show that this can lead to overall longer average queue lengths, and therefore also longer average waiting times.

The existence of a Braess paradox has been explored in several contexts related to queueing networks (see for instance [4, 10]). More closely related to our present work, performance issues due to flexibility were studied in skill based routing models such as queueing systems with redundant requests in [12]. They demonstrate that adding redundancy to a class improves its mean response time but can hurt the mean response time of other classes. Skill based routing models are used in many applications, such as for instance call centers [11], and have several connections to matching models. In [3], the authors show that the FCFS infinite bipartite matching model has the same state description and stationary distribution as the "First Come First Served-Assign the Longest Idle Server" (FCFS-ALIS) parallel queueing model conditioned on all the servers being busy, which implies heavy-traffic assumptions on the arrival rates. The latter also shares the same stationary distribution as the redundancy service model as proven in [2]. Another link between bipartite matching models and redundancy service model was made in [2] through a new model called the FCFS directed infinite bipartite matching model which assumes that arrivals are one by one (instead of pair by pair) and supply items can only be matched with earlier demand items (unmatched supply items are ignored and the state space is only the unmatched demand items). Both models were proved to be path-wise equivalent.

However, these connections between redundancy service model and stochastic matching models rely on specific assumptions, that we do not suppose in our work, such as bipartite compatibility graph and server by server Markov representation, or ignoring some items if they are not immediately matched upon arrival. In addition, the performance study of [12] only covers individual classes. In our paper, we study how flexibility can hurt the overall system by looking at the evolution of the mean total number of remaining items.

3. MODEL DESCRIPTION

We consider a queueing system with n classes of items. Items of different classes arrive to the system according to independent Poisson processes, with rates $\lambda_i,\ i = 1, \dots, n$. The compatibilities between item classes are described by a connected non-bipartite compatibility graph $\mathcal{G} = (\mathcal{V}, \xi)$, where $\mathcal{V} = \{1, \dots, n\}$ is the set of item classes and ξ is the set of allowed matching pairs: items of classes i and j are compatible if and only if $(i, j) \in \xi$. If the incoming item is compatible with at least one unmatched item waiting in the system, we use the FCFM policy which matches the former with the oldest compatible unmatched item. Otherwise, the incoming item is placed at the end of the queue of unmatched items.

After applying standard uniformization technique, with a uniformization constant $\Lambda > \sum_{i=1}^{n} \lambda_i$, we obtain a following matching model in discrete time. In each time slot $t \in \mathbb{N}^*$, one item arrives to the system with probability $1 - \alpha_0$, and there are no arrivals with probability

$$\alpha_0 = \Lambda - \sum_{i=1}^{n} \lambda_i > 0.$$

We assume that this item belongs to a class within the set of item classes $\mathcal{V}$, sampled from a conditional probability distribution $\alpha = (\alpha_1, \cdots, \alpha_n)$ over $\mathcal{V}$ given the event that there is an arrival. Note that this discrete time system has average stationary queue lengths equal to the original continuous time model.

Let us start by defining some notation used throughout the paper. For a class $i \in \mathcal{V}$, let

$$\mathcal{E}(i) = \{j \in \mathcal{V} : (i, j) \in \xi\}$$

denote the set of item classes that are compatible with class i, i.e the set of all the neighbors of the node i in the compatibility graph $\mathcal{G}$. For any subset of item classes $V \subseteq \mathcal{V}$, let

$$\mathcal{E}(V) = \bigcup_{i \in V} \mathcal{E}(i)$$

and

$$|\alpha_V| = \sum_{i \in V} \alpha_i.$$

A subset of nodes $\mathcal{I} \subseteq \mathcal{V}$ is called an *independent set* if there is no edge between any two items in $\mathcal{I}$, i.e. for any $i, j \in \mathcal{I},\ (i, j) \notin \xi$. Let $\mathbb{I}$ be the set of independent sets of $\mathcal{G}$. Furthermore, let $\mathbb{I}_{\mathcal{I}} \subseteq \mathbb{I}$ be the set of independent sets that are subsets of $\mathcal{I} \in \mathbb{I}$, i.e $\mathbb{I}_{\mathcal{I}} = \{\check{\mathcal{I}} \in \mathbb{I} : \check{\mathcal{I}} \subseteq \mathcal{I}\}$.

Let $\mathcal{V}^*$ be the set of finite words over the alphabet $\mathcal{V}$ and denote the empty word by $\emptyset$. Let $\mathbb{W} = \{w = w_1 \cdots w_q \in \mathcal{V}^* : \forall (i, j) \in [\![1, q]\!]^2,\ i \neq j,\ (w_i, w_j) \notin \xi\}$ be the subset of words without a compatible pair of letters, i.e. the set of ordered independent sets of $\mathcal{G}$. For any $w \in \mathbb{W}$ and any $v \in \mathcal{V}$, let $|w|_v$ be the number of occurrences of letter v in word w. Let $|w| = \sum_{v \in \mathcal{V}}$ be the *length* of w.

A state of the system right after the new arrival (if any) has been matched or placed in the queue of unmatched items can be described by a word $w \in \mathbb{W}$. Each letter $w_i \in \mathcal{V}$ represents the class of an unmatched item waiting in the system and the order of the letters represents their order of arrival. The number of items of class v remaining in the system is equal to $|w|_v$.

The FCFM matching process can be modeled as a discrete time Markov chain $W = (W_t)_{t \in \mathbb{N}}$ with values in $\mathbb{W}$, an initial state $w^0 \in \mathbb{W}$ and with the following transitions: assume $W_t = w = w_1 \cdots w_q \in \mathbb{W}$, then with probability $(1 - \alpha_0)\alpha_i$ an item of class $i \in \mathcal{V}$ arrives and

$$W_{t+1} = \begin{cases} w\, i & \text{if } w_l \notin \mathcal{E}(i),\ \forall l \in \{1, \cdots, q\} \\ w_{-i} & \text{otherwise} \end{cases}$$

with w_{-i} being defined as the word w where we remove the first appearance of any letter that belongs to $\mathcal{E}(i)$.

We assume that α satisfies the necessary and sufficient conditions for stability of the model given in [15], i.e

$$|\alpha_{\mathcal{I}}| < |\alpha_{\mathcal{E}(\mathcal{I})}|, \quad \forall \mathcal{I} \in \mathbb{I}. \tag{1}$$

In [15] the authors also prove that the stationary distribution of W has a product form which we recall in the following theorem.

THEOREM 1 ([15]). *Assume that α satisfies* (1)*. Then, the stationary distribution of W, noted π, is equal to*

$$\pi(w) = \pi_{\emptyset} \prod_{i=1}^{q} \frac{\alpha_{w_i}}{|\alpha_{\mathcal{E}(\{w_1, \cdots, w_i\})}|} \quad , \text{ for any } w = w_1 \cdots w_q \in \mathbb{W},$$

where $\pi_{\emptyset}$ is the normalization constant.

We define $Q_t = \sum_{i \in \mathcal{V}} |W_t|_i$ as the total number of items remaining in the system at time t. We denote by $\mathbb{E}[Q]$ the mean total number of items present in the system under the stationary distribution π.

We also consider another compatibility graph $\widetilde{\mathcal{G}} = (\mathcal{V}, \widetilde{\xi})$ where we added the edge (i^*, j^*), i.e $\widetilde{\xi} = \xi \cup \{(i^*, j^*)\}$. We denote by $\widetilde{W}$ the Markov chain defined by $\widetilde{\mathcal{G}}$ on $\widetilde{\mathbb{W}}$ and by $\mathbb{E}[\widetilde{Q}]$ the mean total number of items present in the system with the added edge. Since $\mathcal{E}(\mathcal{I}) \subseteq \widetilde{\mathcal{E}}(\mathcal{I})$ for all $\mathcal{I} \in \mathbb{I}$ and we assume (1) is satisfied, $\widetilde{W}$ is also positive recurrent.

We say that there exists a performance paradox if there exists an edge (i^*, j^*) such that

$$\mathbb{E}[\widetilde{Q}] > \mathbb{E}[Q],$$

that is, if the mean number of items can be increased by adding an edge to the compatibility graph.

For example, we will compare the performance between a matching model with the compatibility graph of Fig. 1 and a matching model with a complete compatibility graph of four nodes (i.e we add an edge between node 1 and node 2 in Fig. 1). We will show that depending on the arrival probability distribution α there exists a performance paradox or not.

4. EXPECTED TOTAL NUMBER OF UNMATCHED ITEMS

We start by computing a closed-form expression for $\mathbb{E}[Q]$, the expected total number of remaining items in the system. This result will be used in the proof of the existence of a performance paradox in the next section.

For an independent set $\mathcal{I} \in \mathbb{I}$, consider an ordered version of $\mathcal{I}$, noted $\mathcal{I}^o = (i_1, \cdots, i_{|\mathcal{I}|})$. We note σ a permutation of its elements, i.e $\mathcal{I}^{\sigma(o)} = (i_{\sigma(1)}, \cdots, i_{\sigma(|\mathcal{I}|)})$ and $\mathfrak{S}_{|\mathcal{I}|}$ the set of all permutations of $[\![1, |\mathcal{I}|]\!]$.

We define

$$T_{\mathcal{I}^o} = \prod_{k=1}^{|\mathcal{I}|} \frac{\alpha_{i_k}}{|\alpha_{\mathcal{E}(\{i_1,\cdots,i_k\})}| - |\alpha_{\{i_1,\cdots,i_k\}}|}$$

and $T_{\mathcal{I}} = \sum_{\sigma \in \mathfrak{S}_{|\mathcal{I}|}} T_{\mathcal{I}^{\sigma(o)}}$. We also define

$$E_{\mathcal{I}^o} = \sum_{l=1}^{|\mathcal{I}|} \frac{|\alpha_{\mathcal{E}(\{i_1,\cdots,i_l\})}|}{|\alpha_{\mathcal{E}(\{i_1,\cdots,i_l\})}| - |\alpha_{\{i_1,\cdots,i_l\}}|} \times \prod_{k=1}^{|\mathcal{I}|} \frac{\alpha_{i_k}}{|\alpha_{\mathcal{E}(\{i_1,\cdots,i_k\})}| - |\alpha_{\{i_1,\cdots,i_k\}}|}$$

and $E_{\mathcal{I}} = \sum_{\sigma \in \mathfrak{S}_{|\mathcal{I}|}} E_{\mathcal{I}^{\sigma(o)}}$.

Moyal et al. [15] proved that the normalizing constant of the stationary distribution can be written as the sum of terms over the independent sets. In the following result, we show that the expected total number of remaining items, defined as an infinite sum over all possible words, can be also written as a finite sum over all independent sets. The proof can be found in [8] the extended version of this paper.

PROPOSITION 1. *Let $\mathbb{E}[Q]$ be the expected value of Q under the stationary distribution π. Then*

$$\mathbb{E}[Q] = \left(1 + \sum_{\mathcal{I} \in \mathbb{I}} T_{\mathcal{I}}\right)^{-1} \left(\sum_{\mathcal{I} \in \mathbb{I}} E_{\mathcal{I}}\right).$$

Let us now consider again the new compatibility graph $\widetilde{\mathcal{G}} = (\mathcal{V}, \widetilde{\xi})$ obtained from $\mathcal{G}$ by adding the edge (i^*, j^*), i.e $\widetilde{\xi} = \xi \cup \{(i^*, j^*)\}$. Let $\widetilde{\mathbb{I}}$ be the set of independent sets of $\widetilde{\mathcal{G}}$. Since $\widetilde{\mathcal{E}}$ only differs from $\mathcal{E}$ in the added edge (i^*, j^*), $\widetilde{\mathbb{I}}$ can be obtained from $\mathbb{I}$ by removing all the independent sets that contain both i^* and j^*.

We are interested in the sign of the difference in the expected values of $\widetilde{Q}$ and Q, i.e the sign of $\mathbb{E}[\widetilde{Q}] - \mathbb{E}[Q]$. Since α satisfies (1), the denominator of each expectation is positive and thus, the sign of $E[\widetilde{Q}] - \mathbb{E}[Q]$ is also the sign of the numerator of $E[\widetilde{Q}]$ times the denominator of $\mathbb{E}[Q]$ minus the numerator of $\mathbb{E}[Q]$ times the denominator of $E[\widetilde{Q}]$. Some terms can also be factorized using a partition of $\mathbb{I}$ based on whether an independent set $\mathcal{I}$ contains the node i^* or the node j^*.

5. PERFORMANCE PARADOX

In this section, we prove sufficient conditions for the existence or non-existence of a performance paradox in matching models under a heavy-traffic assumption. This asymptotic assumption is similar to the assumption (A1) used in [6] to prove the approximate optimality of their matching policy. All the proofs of this section can be found in [8] the extended version of this paper.

For any $\mathcal{I} \in \mathbb{I}$, denote by $|W_t|_{\mathcal{I}} = \sum_{i \in \mathcal{I}} |W_t|_i$, $t \geq 0$ and

$$\Delta_{\mathcal{I}} = |\alpha_{\mathcal{E}(\mathcal{I})}| - |\alpha_{\mathcal{I}}|.$$

Under FCFM policy, for any $t \geq 0$, we have

$$\mathbb{E}[|W_{t+1}|_{\mathcal{I}} - |W_t|_{\mathcal{I}}] \geq -\Delta_{\mathcal{I}},$$

with the equality that is achieved for example when $|W_t|_i > 0$, $\forall i \in \mathcal{I}$ and $|W_t|_i = 0$, $\forall i \in \mathcal{E}(\mathcal{E}(\mathcal{I})) \backslash \mathcal{I}$. A set $\mathcal{I} \in \arg\min_{\mathcal{I} \in \mathbb{I}} \Delta_{\mathcal{I}}$ will be called a *bottleneck set* in the sense that it has the smallest maximal draining speed. Let

$$\bar{\delta} = \min_{\mathcal{I} \in \mathbb{I}} \Delta_{\mathcal{I}} = \min_{\mathcal{I} \in \mathbb{I}} (|\alpha_{\mathcal{E}(\mathcal{I})}| - |\alpha_{\mathcal{I}}|).$$

We have $\bar{\delta} > 0$ as we assume that α satisfies (1). We select a bottleneck set $\hat{\mathcal{I}} \in \arg\min_{\mathcal{I} \in \mathbb{I}} \Delta_{\mathcal{I}}$ with the highest cardinality, i.e.

$$|\hat{\mathcal{I}}| = \max_{\mathcal{I} \in \mathbb{I} \text{ s.t. } \Delta_{\mathcal{I}} = \bar{\delta}} |\mathcal{I}|.$$

Consider a matching model with the compatibility graph of Fig. 1 and let us define α, the conditional probability distribution of the arrivals, as $\alpha_1 = \alpha_2 = 0.22$, $\alpha_3 = 0.45$ and $\alpha_4 = 0.11$. The independent set $\hat{\mathcal{I}} = 3$ is the only bottleneck set that achieves the smallest maximal draining speed $\bar{\delta} = \Delta_{\{3\}} = |\alpha_{\{1,2,4\}}| - |\alpha_3| = 0.1$.

We are interested in how the performance of the system will evolve by adding an edge when $\Delta_{\hat{\mathcal{I}}}$ tends towards 0. First, we define a parameterized family of item class distributions:

$$\alpha_i^{\delta} = \begin{cases} \alpha_i + \frac{\bar{\delta}}{2} \frac{\alpha_i}{|\alpha_{\hat{\mathcal{I}}}|} - \frac{\delta}{2} \frac{\alpha_i}{|\alpha_{\hat{\mathcal{I}}}|} & \text{if } i \in \hat{\mathcal{I}} \\ \alpha_i - \frac{\bar{\delta}}{2} \frac{\alpha_i}{|\alpha_{\mathcal{E}(\hat{\mathcal{I}})}|} + \frac{\delta}{2} \frac{\alpha_i}{|\alpha_{\mathcal{E}(\hat{\mathcal{I}})}|} & \text{if } i \in \mathcal{E}(\hat{\mathcal{I}}) \\ \alpha_i & \text{otherwise} \end{cases}$$

for all $0 < \delta \leq \bar{\delta}$. It is clear that $\alpha^{\bar{\delta}} = \alpha$. By definition of α^{δ}, when δ tends to 0, then $|\alpha^{\delta}_{\mathcal{E}(\hat{\mathcal{I}})}| - |\alpha^{\delta}_{\hat{\mathcal{I}}}| = \delta$ tends to 0. We need to verify that α^{δ} is a distribution.

LEMMA 1. *Let $0 < \delta \leq \bar{\delta}$. We have $\alpha_i^{\delta} > 0$ for all $i \in \mathcal{V}$ and $\sum_{i \in \mathcal{V}} \alpha_i^{\delta} = 1$.*

In order to use the stationary distribution of Theorem 1 and to make sure that $\mathbb{E}[Q]$ is finite, we need to verify that α^{δ} satisfies (1).

PROPOSITION 2. *Let $0 < \delta \leq \bar{\delta}$. Then, α^{δ} satisfies the stability condition* (1).

We have constructed a continuous path of distributions α^{δ}, for $0 < \delta \leq \bar{\delta}$, so we can consider $\mathbb{E}[Q]$ as a function of δ and take its limit when δ tends to 0. We can rewrite α^{δ} as a linear combination of δ, i.e $\alpha_i^{\delta} = a_i + b_i \delta$ with

$$a_i = \begin{cases} \alpha_i + \frac{\bar{\delta}}{2} \frac{\alpha_i}{|\alpha_{\hat{\mathcal{I}}}|} & \text{if } i \in \hat{\mathcal{I}} \\ \alpha_i - \frac{\bar{\delta}}{2} \frac{\alpha_i}{|\alpha_{\mathcal{E}(\hat{\mathcal{I}})}|} & \text{if } i \in \mathcal{E}(\hat{\mathcal{I}}) \\ \alpha_i & \text{otherwise} \end{cases}$$

and

$$b_i = \begin{cases} -\frac{\alpha_i}{2|\alpha_{\hat{\mathcal{I}}}|} & \text{if } i \in \hat{\mathcal{I}} \\ \frac{\alpha_i}{2|\alpha_{\mathcal{E}(\hat{\mathcal{I}})}|} & \text{if } i \in \mathcal{E}(\hat{\mathcal{I}}) \\ 0 & \text{otherwise} \end{cases}.$$

DEFINITION 1. *An independent set $\mathcal{I}$ is called saturated if $\Delta_{\mathcal{I}}^{\delta} = |\alpha^{\delta}_{\mathcal{E}(\mathcal{I})}| - |\alpha^{\delta}_{\mathcal{I}}|$ tends to 0 when δ tends to 0, i.e. if $|a_{\mathcal{E}(\mathcal{I})}| - |a_{\mathcal{I}}| = 0$.*

PROPOSITION 3. *The vector $a = (a_1, \cdots, a_n)$ is stochastic, satisfies the stability condition* (1) *for all $\mathcal{I} \in \mathbb{I} \setminus \{\hat{\mathcal{I}}\}$ and $|a_{\hat{\mathcal{I}}}| = |a_{\mathcal{E}(\hat{\mathcal{I}})}|$, i.e. $\hat{\mathcal{I}}$ is the only saturated independent set for our parametrized family α^{δ}.*

Recall the previous example of a matching model with the compatibility graph of Fig. 1 and with α defined as $\alpha_1 = \alpha_2 = 0.22$, $\alpha_3 = 0.45$ and $\alpha_4 = 0.11$. We define a new collection of conditional probability distributions α^δ based on the bottleneck set $\hat{\mathcal{I}} = \{3\}$, for all $0 < \delta \leq 0.1$, i.e $\alpha_1^\delta = \alpha_2^\delta = 0.2 + \frac{\delta}{5}$, $\alpha_3^\delta = 0.5 - \frac{\delta}{2}$ and $\alpha_4^\delta = 0.1 + \frac{\delta}{10}$. Thus $\Delta_{\hat{\mathcal{I}}}^\delta = |\alpha_{\mathcal{E}(\hat{\mathcal{I}})}^\delta| - |\alpha_{\hat{\mathcal{I}}}^\delta| = \delta$ tends to 0 when δ tends to 0 and $\hat{\mathcal{I}}$ is the only saturated independent set.

We are now ready to present the main result of this paper about the existence or non-existence of a performance paradox for matching models with one saturated independent set.

THEOREM 2. *If $\hat{\mathcal{I}}$ has both i^* and j^* as neighbors, then there exists a performance paradox for δ sufficiently small. If $\hat{\mathcal{I}}$ contains i^* or j^* and $\mathcal{E}(\hat{\mathcal{I}}) \subsetneq \widetilde{\mathcal{E}}(\hat{\mathcal{I}})$, then there does not exist a performance paradox for δ sufficiently small.*

PROOF. The existence or non-existence of the performance paradox depends on the sign of $\mathbb{E}[\widetilde{Q}] - \mathbb{E}[Q]$. We start by rewriting the difference in expected values as in Section 4. Then, we put every term on the same denominator. Using a naive approach, we can multiply the denominator of every term together to get the common denominator. However, by doing so, we introduce a lot of duplicate terms. Removing the duplicate terms results in the common denominator being the multiplication of values like $\Delta_{\mathcal{I}} = |\alpha_{\mathcal{E}(\mathcal{I})}^\delta| - |\alpha_{\mathcal{I}}^\delta|$ for various independent sets $\mathcal{I}$. One of them is the saturated independent set $\hat{\mathcal{I}}$ and it is defined such that $\Delta_{\hat{\mathcal{I}}} = |\alpha_{\mathcal{E}(\hat{\mathcal{I}})}^\delta| - |\alpha_{\hat{\mathcal{I}}}^\delta| = \delta$, which tends to 0^+ when δ tends to 0^+. For all the other independent sets the value tends to a positive constant due to Proposition 3. Thus, the denominator will tend to 0^+ when δ tends to 0^+ and the only concern now is if the polynomial in δ at the numerator has a constant term and of which sign.

In order to get the common denominator, the numerator of each term will also be multiplied by values such as $\Delta_{\mathcal{I}}$ with various independent sets $\mathcal{I}$. The presence of a constant term and its sign will thus depend on whether the satured independent set $\hat{\mathcal{I}}$ is within those independent sets or not.

If $\hat{\mathcal{I}}$ has both i^* and j^* as neighbors, then the numerator will have a positive constant. If $\hat{\mathcal{I}}$ contains i^* or j^* and $\mathcal{E}(\hat{\mathcal{I}}) \subsetneq \widetilde{\mathcal{E}}(\hat{\mathcal{I}})$, then the numerator will have a negative constant. Since the denominator tends to 0^+ when δ tends to 0^+, then for δ sufficiently small, $\mathbb{E}[\widetilde{Q}] - \mathbb{E}[Q]$ is positive in the first case and negative in the second. See [8] for the detailed proof of this theorem. $\square$

Comparing the performance of the matching model with the compatibility graph of Fig. 1 and α defined such as $\hat{\mathcal{I}} = \{3\}$ is the saturated independent set, with the performance of the matching model with the complete compatibility graph lead to a paradox for δ sufficiently small. Indeed, the added edge $(i^*, j^*) = (1, 2)$ has both nodes as neighbors of $\hat{\mathcal{I}} = \{3\}$. In this example, we can even compute exactly how small δ must be for $\mathbb{E}[\widetilde{Q}] - \mathbb{E}[Q]$ to be positive which is for all $0 < \delta \leq 0.0818369$.

The intuition behind the results of Theorem 2 is that the performance paradox occurs when the added edge in the compatibility graph disrupts the draining of a bottleneck. Indeed, as δ shrinks, the maximal draining speed of the saturated independent set $\hat{\mathcal{I}}$ becomes very small compared to the other independent sets. Thus, the set of classes $\hat{\mathcal{I}}$ is a critical part of the system and every item of those classes should be matched anytime a compatible item arrives. However, adding an edge in the compatibility graph between two neighbors of $\hat{\mathcal{I}}$ means that sometimes the FCFM policy will match items of those two neighbors together instead of having them available when items of $\hat{\mathcal{I}}$ arrive. In that case, it is quite intuitive that the load would increase on this already critical part of the system which lead to degrading performances. Assume now that the added edge in the compatibility graph is between a node within $\hat{\mathcal{I}}$ and a node that was not a neighbor of $\hat{\mathcal{I}}$. Then, items of the new neighbor of $\hat{\mathcal{I}}$ can now be sometimes matched with the latter. Thus, reducing the load on the critical part of the system and improving the performance. This other case gives some intuition on the conditions in Theorem 2 where a performance paradox does not exist.

6. CONCLUSIONS

Adding an edge to the compatibility graph of a matching model increases the number of potential matchings, so one might think that the mean number of items in the system decreases. In this work, we have considered a matching model with an arbitrary compatibility graph and we have shown that this is not always the case when the matching policy is FCFM, that is, when an incoming item is matched with the oldest of its compatible items. First, we have shown that the mean number of items in the system can be written as a finite sum over all independent sets. This result has allowed us to show the main contribution of this paper, which gives sufficient conditions for the existence of a performance paradox and for its non-existence.

In future work, we aim to further investigate the existence of a performance paradox when the saturated independent set does not contain any of the nodes of the added edge and also does not have them as neighbors. We also aim to relax the assumption of having exactly one saturated independent set. Another interesting extension of this work is to analyze the existence of a performance paradox for other matching policies. We are also interested in investigating the existence of a performance paradox in bipartite stochastic matching models. Indeed, the stationary distribution of the Markov chain for the bipartite stochastic matching model has a similar product form when the matching policy is FCFM. The proof for the expected number of unmatched items can be extended to bipartite graphs. However, the condition of items arriving in pairs of one demand and one supply item implies that in the bipartite case there are always at least two saturated independent sets. Thus the extension of our main result to bipartite matching models is not straightforward, as it requires the extension to multiple saturated independent sets.

7. REFERENCES

[1] I. Adan, A. Bušić, J. Mairesse, and G. Weiss. Reversibility and further properties of FCFS infinite bipartite matching. *Mathematics of Operations Research*, 43(2):598–621, 2018.

[2] I. Adan, I. Kleiner, R. Righter, and G. Weiss. FCFS parallel service systems and matching models. *Performance Evaluation*, 127-128:253 – 272, 2018.

[3] I. Adan and G. Weiss. Exact FCFS matching rates for two infinite multitype sequences. *Operations Research*, 60(2):475–489, 2012.

[4] N. G. Bean, F. P. Kelly, and P. G. Taylor. Braess's paradox in a loss network. *Journal of Applied Probability*, 34(1):155–159, 1997.

[5] J. Begeot, I. Marcovici, P. Moyal, and Y. Rahme. A general stochastic matching model on multigraphs, 2020. ArXiv preprint: `https://arxiv.org/abs/2011.05169`.

[6] A. Busic and S. Meyn. Approximate optimality with bounded regret in dynamic matching models. *SIGMETRICS Perform. Eval. Rev.*, 43(2):75–77, 2015. ArXiv preprint: `http://arxiv.org/abs/1411.1044`.

[7] A. Bušić, V. Gupta, and J. Mairesse. Stability of the bipartite matching model. *Advances in Applied Probability*, 45(2):351–378, 2013.

[8] A. Cadas, J. Doncel, J.-M. Fourneau, and A. Bušić. Flexibility can hurt dynamic matching system performance, 2020. ArXiv preprint: `https://arxiv.org/abs/2009.10009`.

[9] R. Caldentey, E. H. Kaplan, and G. Weiss. FCFS infinite bipartite matching of servers and customers. *Advances in Applied Probability*, 41(3):695–730, 2009.

[10] J. E. Cohen and F. P. Kelly. A paradox of congestion in a queuing network. *Journal of Applied Probability*, 27(3):730–734, 1990.

[11] N. Gans, G. Koole, and A. Mandelbaum. Telephone call centers: Tutorial, review, and research prospects. *Manufacturing & Service Operations Management*, 5(2):79–141, 2003.

[12] K. Gardner, S. Zbarsky, S. Doroudi, M. Harchol-Balter, E. Hyytiä, and A. Scheller-Wolf. Queueing with redundant requests: exact analysis. *Queueing Systems*, 83(3-4):227–259, 2016.

[13] J. F. C. Kingman. On queues in heavy traffic. *Journal of the Royal Statistical Society. Series B (Methodological).*, 24(2):383–392, 1962.

[14] J. Mairesse and P. Moyal. Stability of the stochastic matching model. *Journal of Applied Probability*, 53(4):1064–1077, 2016.

[15] P. Moyal, A. Bušić, and J. Mairesse. A product form for the general stochastic matching model. *J. Appl. Probab.*, 58(2):449–468, 2021.