



HAL
open science

Stochastic dynamic matching: A mixed graph-theory and linear-algebra approach

Céline Comte, Fabien Mathieu, Ana Bušić

► **To cite this version:**

Céline Comte, Fabien Mathieu, Ana Bušić. Stochastic dynamic matching: A mixed graph-theory and linear-algebra approach. 2023. hal-03502084v4

HAL Id: hal-03502084

<https://hal.science/hal-03502084v4>

Preprint submitted on 25 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Stochastic dynamic matching: A mixed graph-theory and linear-algebra approach

Céline Comte¹ ✉

Eindhoven University of Technology, Eindhoven, The Netherlands
LAAS-CNRS, CNRS, Toulouse, France

Fabien Mathieu ✉

LINCS, Paris, France
Swapcard, Paris, France

Ana Bušić ✉

Inria, Paris, France
DI ENS, CNRS, PSL University, Paris, France

Abstract

Stochastic dynamic matching problems have recently drawn attention in the stochastic-modeling community due to their numerous applications, ranging from supply-chain management to kidney exchange programs. In this paper, we consider a matching problem in which items of different classes arrive according to independent Poisson processes. Unmatched items are stored in a queue, and compatibility constraints are described by a simple graph on the classes, so that two items can be matched if their classes are neighbors in the graph. We analyze the efficiency of matching policies, not only in terms of system stability, but also in terms of matching rates between different classes.

Our results rely on the observation that, under any stable policy, the matching rates satisfy a conservation equation that equates the arrival and departure rates of each item class. Our main contributions are threefold. We first introduce a mapping between the dimension of the solution set of this conservation equation, the structure of the compatibility graph, and the existence of a stable policy. In particular, this allows us to derive a necessary and sufficient stability condition that is verifiable in polynomial time. Secondly, we describe the convex polytope of non-negative solutions of the conservation equation. When this polytope is reduced to a single point, we give a closed-form expression of the solution; in general, we characterize the vertices of this polytope using again the graph structure. Lastly, we study which vectors of the polytope can be achieved by a stable policy. We show that the set of vectors reached by stable greedy policies is included in the interior of the polytope, and that the inclusion is strict in general. In contrast, we conjecture that non-greedy policies can reach any point of the interior of the polytope; whether they can also reach the boundary of the polytope depends on a simple condition on the vertices.

2012 ACM Subject Classification Mathematics of computing → Queueing theory; Mathematics of computing → Markov processes; Mathematics of computing → Matchings and factors

Keywords and phrases stochastic dynamic matching, graph theory, linear algebra, stability, matching rates, conservation equation

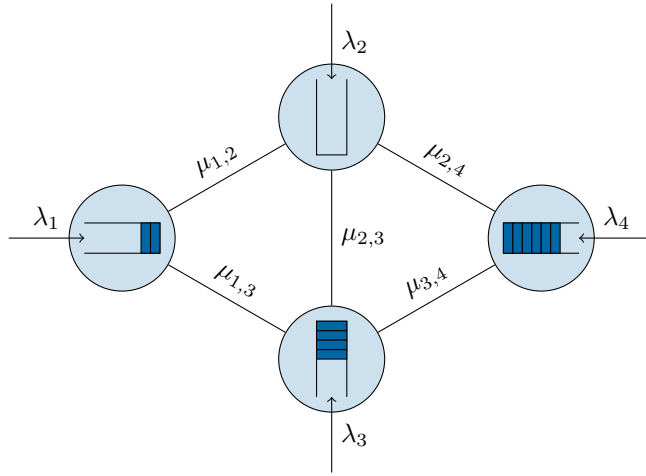
1 Introduction

Stochastic dynamic matching problems, in which items arrive at random instants to be matched with other items, have recently attracted much attention in the stochastic-modeling community. These challenging control problems are highly relevant in many applications, including supply-chain management, pairwise kidney exchange programs, and online marketplaces. In pairwise kidney exchange programs for example, each item represents a donor-receiver pair, and two pairs can be matched if the donor of each pair is compatible

¹ Corresponding author.

with the receiver of the other pair. In online marketplaces, items are typically divided into two categories, called demand and supply, and the goal is to maximize a certain long-term performance criteria by appropriately matching demand items with supply items.

In this paper, we consider the following dynamic matching problem². Items of different classes arrive according to independent Poisson processes. Compatibility constraints between items are described by a simple graph on their classes, such that two items can be matched if their classes are neighbors in the graph. Unmatched items are stored in the queue of their class, and the matching policy decides which matches are performed and when. All in all, a stochastic matching model is described by a triplet (G, λ, Φ) , where $G = (V, E)$ is the compatibility graph, $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ is the vector of per-class arrival rates, and Φ is the matching policy. In Figure 1 for instance, there are four item classes numbered from 1 to 4; classes 2 and 3 are compatible with all classes, while classes 1 and 4 are compatible only with classes 2 and 3.



■ **Figure 1** Illustration of a matching model (G, λ, Φ) on the diamond graph.

We consider two performance indicators, namely the stability and the matching rates along edges. Formally, a matching model (G, λ, Φ) is said to be stable if the associated Markov chain is positive recurrent. Assuming that this matching model is stable, the matching rate μ_k along an edge $k \in E$ with endpoints $i, j \in V$ is the rate at which class- i items and class- j items are matched. Our end goal is to characterize the matching rates that are achievable under stable matching policies and to use this characterization to design matching policies that achieve certain matching rate vectors. Our results show that these two performance indicators are sometimes incompatible, in the sense that optimizing a function of the matching rate vector can lead to instability in some settings.

1.1 Motivation

Besides providing us with criteria to compare the long-term impact of different matching policies, stability conditions and matching rates have practical value in many instances of matching models.

² Detailed definitions of the concepts discussed here will be given in Section 2.

Maximal reward

First assume that a reward w_k is earned every time a match is performed between items of the classes i and j that are endpoints of edge k , for each $k \in E$, and that we look for a policy Φ that maximizes the long-run reward rate $\sum_{k \in E} w_k \mu_k$ while making the system stable. This optimization problem is relevant in applications to organ exchange programs, where the rewards capture the desirability of a given match, accounting for metrics such as the quality of life after transplant and the survival rates of the recipients and donors [7].

Perhaps the most natural policy to consider is a priority policy whereby each incoming item is immediately matched with an unmatched item of the compatible class leading to the highest reward, if any. Unfortunately, it was observed in [26, Section 5] that such a priority policy may lead to instability. Furthermore, the use of a matching policy with a specific priority order for edges does not guarantee that the ordering of the resulting matching rates is consistent with the edge priorities. Depending on the properties of the compatibility graph G and the arrival rate vector λ , we will introduce either a stable matching policy that yields the optimal reward (if such a stable policy exists) or a family of stable policies that we conjecture get arbitrarily close to the optimal.

Chained matching

Matching rates become also crucial when a stable “first-level” matching model (G, λ, Φ) is the source of the (non-Poisson) arrival process of a “second-level” matching model (or, more generally, of another stochastic system). The second-level model consists of a compatibility graph $G' = (E, E')$ whose nodes are the edges of G . In this context, the matching rate vector μ of the first-level model represents the arrival rate vector in the second-level model. The ability to control this vector μ becomes instrumental in stabilizing or optimizing performance in the second-level model. Such two-level matching systems are found in various applications, such as quantum switches [32, 34] for example.

1.2 Contributions

We propose a unified approach to study two closely-related performance criteria, namely the stability and the matching rates along edges. The following observation, powerful despite its simplicity, is fundamental to all our results: under any stable policy, the arrival rate of class- i items is equal to the departure rate of these items, which in turn is equal to the sum of the matching rates along the edges that are incident to node i . In other words, the matching rates satisfy the following *conservation* equation:

$$\sum_{k \in E_i} \mu_k = \lambda_i, \quad i \in V,$$

where $E_i \subseteq E$ is the set of edges that are incident to node i . In matrix form, this equation rewrites $A\mu = \lambda$, where A is the incidence matrix of the compatibility graph G . The solution set of this conservation equation is related to the structure of the graph G via the linear application $y \in \mathbb{R}^m \mapsto Ay \in \mathbb{R}^n$, where n is the number of nodes (or classes), and m is the number of edges. We say that the graph G is surjective (resp. injective, bijective) if the linear application $y \in \mathbb{R}^m \mapsto Ay \in \mathbb{R}^n$ is surjective (resp. injective, bijective), and we give simple equivalent conditions in terms of the graph structure (Definitions 3.1–3.4 and Proposition 3.5). Our main contributions are threefold, and all rely on these definitions.

We first prove that there exists a direct relation between the form of the solution set of the conservation equation, the structure of the compatibility graph, and the existence of a

stable policy. More specifically, we prove that a compatibility graph G is stabilizable (in the sense that there exist a vector λ and a policy Φ such that the matching model (G, λ, Φ) is stable) if and only if the graph G is surjective (Proposition 3.6). We then prove that a matching problem (G, λ) is stabilizable (in the sense that there exists a policy Φ such that the matching model (G, λ, Φ) is stable) if and only if the conservation equation has a solution with positive coordinates (Proposition 3.7). In particular, this allows us to verify stabilizability in a time that is polynomial in the number of classes and edges.

We next describe the affine space of solutions of the conservation equation and the convex polytope of solutions with non-negative coordinates. When this convex polytope is reduced to a single point, we derive a closed-form expression for the solution (Proposition 4.1). When this convex polytope is not reduced to a single point, we characterize its vertices, again using the graph structure. We prove in particular that a non-negative solution of the conservation equation is a vertex of the convex polytope if and only if the subgraph restricted to the support of this vector is injective (Proposition 5.6).

Lastly, we investigate which parts of the polytope can be achieved by a stable policy. We first focus on greedy policies, that is, policies that never postpone a feasible match. We show that greedy policies are limited to the interior of the polytope and that in general, the inclusion is strict (Propositions 6.2, 6.5, 6.7, and 6.8). In contrast, we conjecture that non-greedy policies can potentially reach any point within the interior of this polytope (Conjecture 6.12 and Corollary 6.14). Additionally, depending on the bijectivity of subgraphs of G , these policies can also reach parts or the entirety of the polytope's boundary (Proposition 6.11 and Conjecture 6.12). To support our conjectures, we have conducted simulations that provide empirical evidence.

1.3 State of the art

We now review the relevant work related to (static or dynamic) matching problems.

Non-bipartite or general stochastic matching

Our work is part of a broader research effort on the stochastic matching model that was briefly discussed earlier and will be described in details in Section 2 [5, 6, 12, 15, 22, 26, 28]. Among these works, the following are particularly relevant because directly related to our results on stability. The paper [26] is the earliest work on this matching model. It derives necessary and sufficient stability conditions that are instrumental in several of our results, in particular Propositions 3.6 and 3.7. This work also proves that the *match-the-longest* policy is maximally stable (in the sense that it always leads to stability whenever the matching problem (G, λ) is stabilizable), a result that is also applied in Proposition 3.7. The papers [15, 28] focus on the *first-come-first-matched* policy. In particular, [28] proves that the first-come-first-matched policy is maximally stable, and [15] provides a new sufficient stability condition we prove to be also necessary in Proposition 3.7.

Other variants of the model were studied recently, and an interesting future work would consist of generalizing our results to these variants. In particular, the paper [22] considers item abandonment, the paper [6] considers graphs with self-loops, and the papers [21, 29, 31] allow matches not limited to two items by replacing the graph with a hypergraph.

The recent paper [5] is perhaps the closest to ours, and we provide a detailed discussion to highlight the relation with our paper. The first statement in [5, Theorem 1] is synonym to the equivalence of statements (ii) and (iii) in Proposition 3.7. Our proof is significantly shorter because it relies more heavily on existing results. Our observation at the beginning

of Section 4 that the conservation equation has a unique solution if and only if the graph is bijective (and not surjective-only) summarizes [5, Theorem 3]. Some formulas derived in [5, Section 9] are special cases of the formulas derived in Proposition 4.1. The non-bipartite matching model in [5] is slightly more general because it considers graphs with self-loops, that is, an edge can have identical endpoints, but this paper does not adopt the mixed graph-theory and linear-algebra approach that supports most of our results.

Bipartite stochastic matching

To the best of our knowledge, the first example of a stochastic matching model with an infinite time horizon in the literature, which predated the model that we consider, is the bipartite matching model introduced in [13] and studied in [1, 2, 3, 9, 10, 11, 16]. In this model, the compatibility graph is bipartite, with two parts that correspond to supply and demand items, respectively. This bipartite model differs from ours by its arrival process: time is slotted and, during each time slot, one demand item *and* one supply item arrive. Several works have made contributions about stability [2, 9] and matching rates [2], and they obtained results similar to those derived in the literature on our model. The bipartite nature of the graph simplifies some calculations, for instance by allowing the application of flow-maximization algorithms to calculate optimal matching rates.

Static and fractional matching

The static matching problem, in which the nodes of the graph represent items (rather than classes), has been extensively studied in mathematics, computer science, and economy [25]. Although the questions raised in static and dynamic matching are often different, the conservation equation that we obtain is reminiscent of several results in static matching. For example, finding a maximum-cardinality matching in the graph G (that is, a maximum-cardinality set of edges without common endpoints) is equivalent to finding integers $\mu_k \in \{0, 1\}$ for each edge $k \in E$ that maximize $\sum_{k \in E} \mu_k$ while satisfying the conservation equation with $\lambda_i = 1$ for each $i \in V$. The relaxation of this integer linear program leads to the so-called fractional matching problem, which has been studied in the literature [25, Section 7.2]. Therefore, the fractional matching polytope defined in [25, Section 7.5] is a special case of the convex polytope that we consider in Section 5.2, and our characterization of this convex polytope is a natural generalization of existing characterizations of the fractional polytope³.

1.4 Outline

The remainder of the paper is organized as follows. **Section 2** gives a formal definition of the model. **Section 3** introduces the conservation equation and defines the notions of surjectivity, injectivity, and bijectivity for a graph that will be instrumental throughout the paper. We use these definitions to formulate new necessary and sufficient stability conditions. In particular, we show that stability requires the compatibility graph to be surjective (that is, either bijective or surjective-only). In **Section 4**, we focus on bijective graphs and give a closed-form expression of the unique solution of the conservation equation. **Section 5** characterizes the solution set of the conservation equation for surjective-only graphs. Lastly,

³ The fractional matching polytope is actually defined using non-strict inequalities rather than equalities. However, one can verify that these two convex polytopes have the same non-zero vertices.

in **Section 6**, we focus on the set of matching rates that are effectively achievable under a stable matching policy.

2 Stochastic dynamic matching

We consider a stochastic dynamic matching system in which items arrive at random times to be matched with other items. Each incoming item may be matched with any unmatched item of a compatible class; in this case, both items disappear immediately. Unmatched items are gathered in a waiting queue. In this paper, such a stochastic dynamic matching system will be described by a triplet (G, λ, Φ) , where G is the *compatibility graph*, λ is the vector of *arrival rates*, and Φ is the *matching policy*. We now review each component in details. To facilitate reference and understanding, we provide a summary of our notation in Table 1.

2.1 Compatibility graph

Compatibility constraints between items are described by a graph $G = (V, E)$, called the *compatibility graph* of the model, which is simple (undirected and without self-loop). The number of nodes is represented by n , while m denotes the number of edges.

We use $V = \{v_1, v_2, \dots, v_n\}$ to denote the set of nodes, where each node corresponds to a class in the matching model. In cases where there is no confusion, we may refer to a class v_i simply by its index i . Following the intuition conveyed in Figure 1, we will use the terms “class i ” and “queue i ” interchangeably, and refer for instance to the number of unmatched class- i items as the size of queue i ; this is a convenience of terminology, and this does not prevent matching decisions from being based on additional information, such as the arrival order of items of different classes (see Section 2.3 for more details on matching decisions).

The set of edges is denoted by $E = \{e_1, e_2, \dots, e_m\}$. These edges represent compatibility constraints between item classes, in the sense that a class- i item and a class- j item can be matched with one another if and only if they are adjacent, that is, if there is an edge with endpoints i and j . When there is no ambiguity, we may refer to an edge $e_k \in E$ with endpoints $i, j \in V$ by its index k or by its set of endpoints $\{i, j\}$. In Figure 1 for instance, there are four item classes numbered from 1 to 4. Classes 2 and 3 are compatible with all classes, but classes 1 and 4 are only compatible with classes 2 and 3. The absence of self-loop means that an item of a given class cannot be matched with other items of the same class.

Lastly, we let \mathbb{I} denote the family of independent sets of the compatibility graph G , where an independent set of G is a non-empty set of nodes that are pairwise non-adjacent. The family of independent sets in the compatibility graph of Figure 1 is $\mathbb{I} = \{\{1\}, \{2\}, \{3\}, \{4\}, \{1, 4\}\}$.

2.2 Arrival process

Class- i items arrive according to an independent Poisson process with rate $\lambda_i > 0$, for each $i \in V$. The vector of arrival rates is denoted by $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n) \in \mathbb{R}_{>0}^n$. Scaling all coordinates of λ by the same positive constant is equivalent to changing the time unit, so we can renormalize λ without changing the dynamics. For example, we will sometimes use the *unit* normalization, in which $\sum_{i \in V} \lambda_i = 1$. We also let $I = (I_t, t \in \mathbb{N})$ denote the sequence of independent and identically distributed (i.i.d.) item classes, so that I_t is the class of the $(t + 1)$ -th item, equal to i with probability $\lambda_i / (\sum_{j \in V} \lambda_j)$, for each $t \in \mathbb{N}$. The couple (G, λ) is called a *matching problem* or simply a *problem*. Occasionally, when we need to specify the sequence of incoming items and not merely its distribution, we will also refer to the couple (G, I) as a (matching) problem.

2.3 Policy and matching dynamics

Most of the paper will focus on deterministic size-based policies, that is, policies whereby matching decisions are deterministic functions of the queue size vectors. However, as we will see at the end of this section, our results also apply to a more general definition of a policy. Throughout the paper, we assume that the system is initially empty, meaning that it starts with no unmatched item.

2.3.1 Size-based policies

A (deterministic) size-based *matching policy* is defined formally as a function $\Phi : \mathcal{Q} \times V \rightarrow V \cup \{\perp\}$, where \mathcal{Q} is an infinite subset of \mathbb{N}^n that contains the reachable states of the system. We will delve into further discussion on \mathcal{Q} in Sections 2.3.2 and 2.3.3. For each $q \in \mathcal{Q}$ and $i \in V$, an incoming class- i item that finds the system in state q is matched with an item of class $\Phi(q, i)$ if $\Phi(q, i) \in V$ and is added to the class- i queue if $\Phi(q, i) = \perp$. The matching policy is assumed to be *adapted* to the compatibility graph G in the sense that

$$\Phi(q, i) \in \{j \in V_i : q_j \geq 1\} \cup \{\perp\}, \quad q \in \mathcal{Q}, \quad i \in V. \quad (1)$$

The system dynamics is described by a Markov chain $Q = (Q_t, t \in \mathbb{N})$, called the *queue-size process*, where $Q_t = (Q_{t,1}, Q_{t,2}, \dots, Q_{t,n})$ is an n -dimensional vector giving the number of unmatched items of each class right after the arrival of the t -th item, for each $t \in \mathbb{N}$, with the assumption that the system is initially empty, that is, $Q_0 = 0$. The system dynamics satisfies the recursion

$$Q_{t+1} = \begin{cases} Q_t + \mathbf{1}_{I_t} & \text{if } J_t = \perp, \\ Q_t - \mathbf{1}_{J_t} & \text{if } J_t \neq \perp, \end{cases} \quad (2)$$

where $J_{t+1} = \Phi(Q_t, I_t)$ for each $t \in \mathbb{N}$, and $\mathbf{1}_i$ is the n -dimensional vector with one in coordinate i and zero elsewhere, for each $i \in V$. We assume that the policy Φ is such that the Markov chain Q has state space \mathcal{Q} and is irreducible. By unfolding the recursion (2), we obtain that, for each $t \in \mathbb{N}$,

$$Q_{t,i} = L_{t,i} - \sum_{k \in E_i} M_{t,k}, \quad t \in \mathbb{N}, \quad i \in V, \quad (3)$$

where $E_i \subseteq E$ is the set of edges that are incident to node i in the graph G , for each $i \in V$, $L_{t,i}$ is the number of class- i items among the first t arrivals, for each $t \in \mathbb{N}$ and $i \in V$, and $M_{t,k}$ is the number of times that classes i and j are matched over the first t arrivals, for each $t \in \mathbb{N}$ and $\{i, j\} = e_k \in E$:

$$L_{t,i} = \sum_{s=0}^{t-1} \mathbf{1}_{\{I_s=i\}}, \quad t \in \mathbb{N}, \quad i \in V, \quad (4)$$

$$M_{t,k} = \sum_{s=0}^{t-1} \mathbf{1}_{\{\{I_s, J_s\}=e_k\}}, \quad t \in \mathbb{N}, \quad k \in E, \quad (5)$$

with the convention that the sums are zero if $t = 0$. The triplet (G, λ, Φ) is called a *matching model*, or simply a *model*. Occasionally, when specifying the sequence of incoming item classes is useful, we will also refer to the triplet (G, I, Φ) as a (matching) model.

2.3.2 Greedy policies

A policy Φ is called *greedy* if an incoming item is matched whenever possible, that is, if there is an unmatched item that is compatible. More formally, the policy Φ is greedy if

$$\Phi(q, i) \neq \perp \text{ for each } (q, i) \in \mathcal{Q} \times V \text{ such that } \{j \in V_i : q_j \geq 1\} \neq \emptyset. \quad (6)$$

Equivalently, a policy Φ is greedy if the set of unmatched item classes under this policy is an independent set of the compatibility graph, meaning that the state space \mathcal{Q} of the queue-size process is equal to

$$\mathcal{Q}_G = \{q \in \mathbb{N}^n : q_i q_j = 0 \text{ for each } i, j \in V \text{ such that } \{i, j\} \in E\}. \quad (7)$$

The following greedy policies will be instrumental in Sections 3 and 6:

- **Match-the-longest:** For each $(q, i) \in \mathcal{Q} \times V$ such that $\sum_{j \in V_i} q_j \geq 1$, we choose $\Phi(q, i) \in \arg \max_{j \in V_i} (q_j)$ (ties are broken arbitrarily). This policy was considered in [26, 6, 22, 5].
- **Edge-priority:** This policy selects matches according to a priority order defined on the edges. More specifically, if we let \prec denote a strict total order on the set E of edges, we let $\Phi(q, i) = \arg \min_{j \in V_i : q_j \geq 1} (\{i, j\})$ for each $(q, i) \in \mathcal{Q} \times V$ such that $\sum_{j \in V_i} q_j \geq 1$, where the $\arg \min$ is to be understood over the ordered set (E, \prec) . This definition remains valid if \prec is a partial order compatible with the possible inputs of $\arg \min$, i.e. if all pairs of edges that are incident to the same node and are not part of a triangle are comparable⁴.

2.3.3 Non-greedy policies

The state space \mathcal{Q} under non-greedy policies is a strict superset of the set \mathcal{Q}_G defined in (7). In Section 6, non-greedy policies will be obtained by modifying greedy policies as follows:

- **Filtering:** Given a subset $\underline{E} \subsetneq E$ of edges, replace V_i with $\underline{V}_i = \{j \in V_i : \{i, j\} \in \underline{E}\}$ in the definition (1) of Φ . Intuitively, we eliminate the edges of $E \setminus \underline{E}$ and follow a (greedy or non-greedy) policy on the subgraph $\underline{G} = (V, \underline{E})$.
- **Semi-filtering:** We consider a modified variant of filtering policies in which the subset \underline{E} of edges depends on the system state q at the arrival time. Examples of semi-filtering policies will be given in Section 6.2.

For instance, under the filtering variant of the match-the-longest policy that eliminates edge $\{2, 3\}$ in the matching problem of Figure 1, an incoming class-3 item is either matched with a class-4 item or added to the queue of class 3 (but it is not matched with a class-2 item).

► **Remark 2.1 (Random policies).** Our choice of notation assumes that the policy is deterministic. A random (size-based) policy Φ can be defined as a function $\Phi : \mathcal{Q} \times V \times (V \cup \{\perp\}) \rightarrow [0, 1]$ such that, for each $t \in \mathbb{N}$, J_t is sampled according to the distribution $\Phi(q, i, \cdot)$ given that $Q_t = q$ and $I_t = i$. Saying that the policy is adapted to the compatibility graph G is then equivalent to saying that, for each $q \in \mathcal{Q}$ and $i \in V$, the support of $\Phi(q, i, \cdot)$ is included into $\{j \in V_i : q_j \geq 1\} \cup \{\perp\}$. The policy is greedy if $\Phi(q, i, \perp) = 0$ for each $(q, i) \in \mathcal{Q} \times V$ such that $\{j \in V_i : q_j \geq 1\} \neq \emptyset$ or, equivalently, if $\mathcal{Q} = \mathcal{Q}_G$.

2.3.4 Other policies

Although deterministic size-based policies are notationwise convenient, our results apply to a broader family of policies that are either random or require a more complex state descriptor,

⁴ If three classes form a triangle in the compatibility graph, at most one of them can be non-empty under a greedy policy, so at most one edge of the triangle can be in the input of $\arg \min$.

or both. We now introduce this family of policies, with the goal of being as general as possible. The remainder of Section 2.3 can be skipped by the first-time reader and referred to when necessary; in this case, S can be understood as Q in all definitions and results.

Extended definition

Under this more general definition, the match-maker makes decisions based not only on the vector of queue sizes, but also (possibly) on additional information that is captured by the system state. The state space is a couple $(\mathcal{S}, |\cdot|)$, where \mathcal{S} is a countably infinite set and $|\cdot| : \mathcal{S} \rightarrow \mathbb{N}^n$ is a function that maps any state $s \in \mathcal{S}$ to the vector giving the number of unmatched items of each class in that state, denoted by $|s| = (|s|_1, |s|_2, \dots, |s|_n)$. The existence of the function $|\cdot|$ guarantees that the system state contains enough information to retrieve the number of unmatched items of each class, which is a classical assumption in queueing theory (see for instance [24, Section 3.2]). We assume that there exists a unique state $s \in \mathcal{S}$ such that $|s| = 0$. This state is called the *empty state* and denoted by \emptyset . This assumption guarantees that the intuitive notion of system stability will be captured by the positive recurrence of the Markov chain describing the evolution of the system state⁵.

The policy is now a function $\Phi : \mathcal{S} \times V \times (V \cup \{\perp\}) \times \mathcal{S} \rightarrow [0, 1]$ such that $\Phi(s, i, j, s')$ is the conditional probability that, given an incoming class- i item finds the system in state s , the matching decision is j and the new state is s' . More formally, the dynamics are described by a Markov chain $((S_t, I_t, J_t), t \in \mathbb{N})$, where $I = (I_t, t \in \mathbb{N})$ is the sequence of incoming item classes and, for each $t \in \mathbb{N}$, $i \in V$, $j \in V \cup \{\perp\}$, and $s, s' \in \mathcal{S}$, we have

$$\mathbb{P}(J_t = j, S_{t+1} = s' \mid S_t = s, I_t = i) = \Phi(s, i, j, s').$$

The stochastic process $S = (S_t, t \in \mathbb{N})$ is also a Markov chain, with transition probabilities

$$\mathbb{P}(S_{t+1} = s' \mid S_t = s) = \frac{\sum_{i \in V} \lambda_i \sum_{j \in V \cup \{\perp\}} \Phi(s, i, j, s')}{\sum_{i \in V} \lambda_i}, \quad t \in \mathbb{N}, \quad s, s' \in \mathcal{S}.$$

We assume that the Markov chain S has state space \mathcal{S} and is irreducible, and that $S_0 = \emptyset$. The policy is assumed to be *adapted* to the compatibility graph G and *consistent* in the sense that, for each $(s, i, j, s') \in \mathcal{S} \times V \times (V \cup \{\perp\}) \times \mathcal{S}$, we have $\Phi(s, i, j, s') > 0$ only if

$$j \in \{j' \in V_i : |s|_{j'} \geq 1\} \cup \{\perp\}, \quad \text{and} \quad |s'| = \begin{cases} |s| + \mathbb{1}_i & \text{if } j = \perp, \\ |s| - \mathbb{1}_j & \text{if } j \neq \perp. \end{cases}$$

Using this extended definition, the previously defined policy models can be easily expressed. For example, the matching policy Φ is called *size-based* if $|\cdot|$ is the identity (implying $\mathcal{S} \subset \mathbb{N}^n$) and *deterministic* if, for each $s \in \mathcal{S}$ and $i \in V$, there exists $(j, s') \in (V \cup \{\perp\}) \times \mathcal{S}$ such that $\Phi(s, i, j, s') = 1$. The policy is called *greedy* if $\sum_{s' \in \mathcal{S}} \Phi(s, i, \perp, s') = 0$ for each $(s, i) \in \mathcal{S} \times V$ such that $\{j \in V_i : |s|_j \geq 1\} \neq \emptyset$, and *non-greedy* otherwise. First-come-first-matched [15, 28] is a classical example of a deterministic policy that is not size-based: its state space is a couple $(\mathcal{S}, |\cdot|)$ where \mathcal{S} is a subset of the set of sequences $c = (c_1, c_2, \dots, c_p)$ made of a finite but arbitrarily large number p of elements of \mathcal{I} , and $|c|_i$ is the cardinality of the

⁵ Without this assumption, one may construct two Markov chains associated with the same system, one positive recurrent and the other transient, for example if the state of the latter Markov chain embeds the time t . This assumption is used only in the proofs of Propositions 6.1 and A.1. In both cases, we can verify that the same conclusion holds as long as the set of states $s \in \mathcal{S}$ such that $|s| = 0$ is finite. Assuming that this set is reduced to a singleton is merely a notational convenience.

set $\{q \in \{1, 2, \dots, p\} : c_q = i\}$, for each $i \in \mathcal{I}$. A policy that is neither size-based nor deterministic will appear in the proof of Proposition 6.1.

The stochastic process $Q = (Q_t, t \in \mathbb{N})$ defined by $Q_t = |S_t|$ for each $t \in \mathbb{N}$ is called the *queue-size process*. This process does not satisfy the Markov property in general, but it satisfies the evolution equations (2) and (3), with $L_i = (L_{t,i}, t \in \mathbb{N})$ and $M_k = (M_{t,k}, t \in \mathbb{N})$ defined by (4) and (5) for each $i \in V$ and $k \in E$. The state space of the queue-size process is given by $\mathcal{Q} = \{|s|, s \in \mathcal{S}\}$. The policy is greedy if $\mathcal{Q} = \mathcal{Q}_G$ and non-greedy if $\mathcal{Q} \supsetneq \mathcal{Q}_G$, where \mathcal{Q}_G is still given by (7).

► **Remark 2.2 (Arrival rates vs. arrival sequence).** We will often identify the matching model (G, λ, Φ) with the Markov chain S . This is a slight abuse of language: the triplet (G, λ, Φ) specifies the transition diagram of this Markov chain but, even if Φ is deterministic, characterizing its sample paths requires specifying the sequence I , sampled according to $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$. This slight abuse of language will not cause confusion when discussing stability and matching rates, but the distinction will matter in Section 6.

► **Remark 2.3 (Discrete time vs. continuous time).** The discrete-time Markov chain S gives the sequence of states observed by incoming items, and it was analyzed under various policies in [22, 26]. Yet in queueing theory, it is more common to consider the continuous-time Markov chain describing the system state over time. However, as observed in [15, Section 2.2.2], S is the jump chain of this continuous-time Markov chain, and both Markov chains have the same stationary measures because the departure rate from each state in the continuous-time Markov chain is constant equal to $\sum_{i \in V} \lambda_i$. Therefore, our results are equally relevant to study performance metrics like the mean queue size or the mean waiting time of items.

Equivalent policies

With our extended definition, a decision rule can be associated with an infinite number of policies. For instance, it is always possible to artificially expand the state definition, resulting in an unlimited range of policies. We define here an equivalence relation between policies that captures the intuitive concepts of *yielding identical distributions of matching decisions* (for random policies) and *making the same decisions* (for deterministic policies). This discussion will also prepare the ground for Propositions 6.4 and 6.6.

Consider a policy Φ_1 adapted to a compatibility graph $G = (V, E)$, and let $(\mathcal{S}_1, |\cdot|_1)$ denote its state space. Assume that the function $|\cdot|_1 : \mathcal{S}_1 \rightarrow \mathbb{N}^n$ can be written as a composition of two functions, $\langle \cdot \rangle : \mathcal{S}_1 \rightarrow \mathcal{S}_2$ and $|\cdot|_2 : \mathcal{S}_2 \rightarrow \mathbb{N}^n$, such that \mathcal{S}_2 is the image of \mathcal{S}_1 through $\langle \cdot \rangle$. Moreover, assume that there exists a policy Φ_2 with state space $(\mathcal{S}_2, |\cdot|_2)$, adapted to the graph G , such that for each $s_2, s'_2 \in \mathcal{S}_2$, $i \in V$, and $j \in V \cup \{\perp\}$, we have

$$\sum_{s'_1 \in \mathcal{S}_1 : \langle s'_1 \rangle = s'_2} \Phi_1(s_1, i, j, s'_1) = \Phi_2(s_2, i, j, s'_2) \quad \text{for each } s_1 \in \mathcal{S}_1 \text{ such that } \langle s_1 \rangle = s_2. \quad (8)$$

We say that Φ_1 can be *reduced* to Φ_2 and that $\langle \cdot \rangle$ is a *reduction* function. If we let $((S_{1,t}, I_t, J_{1,t}), t \in \mathbb{N})$ and $((S_{2,t}, I_t, J_{2,t}), t \in \mathbb{N})$ denote the Markov chains associated with policies Φ_1 and Φ_2 , respectively, under the same sequence $(I_t, t \in \mathbb{N})$ of incoming item classes, then for each $t \in \mathbb{N}$, (i) the conditional distribution of $(J_{1,t}, \langle S_{1,t+1} \rangle)$ given that $S_{1,t} = s_1$ and $I_t = i$ is the same for all states $s_1 \in \mathcal{S}_1$ that have the same image $s_2 = \langle s_1 \rangle$, and (ii) $(\langle S_{1,t} \rangle, I_t, J_{1,t})$ and $(S_{2,t}, I_t, J_{2,t})$ have the same distribution. Conclusion (ii) follows from an inductive argument and implies that policies Φ_1 and Φ_2 are stable or unstable under the same conditions and, if stable, yield the same matching rate vector.

The special case where the policy Φ_2 is deterministic will be useful in Propositions 6.4 and 6.6. In this case, (8) says that, for each $s_2 \in \mathcal{S}_2$ and $i \in V$, there exist $j \in V \cup \{\perp\}$ and

$s'_2 \in \mathcal{S}_2$ such that $\sum_{s'_1 \in \mathcal{S}_1: \langle s'_1 \rangle = s'_2} \Phi_1(s_1, i, j, s'_1) = \Phi_2(s_2, i, j, s'_2) = 1$ for each $s_1 \in \mathcal{S}_1$ such that $\langle s_1 \rangle = s_2$. This condition implies that the Markov chains under Φ_1 and Φ_2 are equivalent pathwise (i.e., $\langle S_{1,t} \rangle = \langle S_{2,t} \rangle$ and $J_{1,t} = J_{2,t}$ for each $t \in \mathbb{N}$) and not just in distribution.

In general, we say that two policies Φ_1 and Φ_2 adapted to the graph G are *equivalent* if there exists a policy Φ adapted to the graph G such that both Φ_1 and Φ_2 can be reduced to Φ . This equivalence between Φ_1 and Φ_2 can be interpreted as indicating that the two policies, when they are in equivalent states, i.e. states that have the same reduction, yield identical distributions of matching decisions. It implies that, if we let $(Q_{1,t}, t \in \mathbb{N})$ and $(Q_{2,t}, t \in \mathbb{N})$ denote the queue-size processes under Φ_1 and Φ_2 with the same sequence $(I_t, t \in \mathbb{N})$ of incoming item classes, then we have $\mathbb{P}(Q_{1,t} = q) = \mathbb{P}(Q_{2,t} = q)$ for each $t \in \mathbb{N}$ and $q \in \mathcal{Q}$.

2.4 Performance criteria

We now define the two performance indicators that will be the focus of the paper, namely stability and the matching rate vector. We will see later that the goals of preserving stability and maximizing a certain function of the matching rate vector are sometimes antagonistic.

Stability

We now define the notions of *stability* and *stabilizability*, which will be explored in Section 3.

► **Definition 2.4** (Stability and stabilizability).

- (i) A model (G, λ, Φ) is called *stable* if the Markov chain $(S_t, t \in \mathbb{N})$ is positive recurrent. In this case, we say that the policy Φ stabilizes the problem (G, λ) .
- (ii) A problem (G, λ) is called *stabilizable* if there exists a policy that stabilizes this problem.
- (iii) A compatibility graph G is called *stabilizable* if there exists a vector $\lambda \in \mathbb{R}_{>0}^n$ such that the problem (G, λ) is stabilizable.

Both the *match-the-longest* and *first-come-first-matched* greedy policies, discussed in Section 2.3.2 and Section 2.3.4, respectively, possess the property of stabilizing all stabilizable matching problems [26, 28]. When the matching problem (G, λ) is clear from the given context, we will call a policy (adapted to the graph G) *stable* if the corresponding model (G, λ, Φ) is stable.

Matching rates

Consider a stable matching model (G, λ, Φ) . We define the *matching rate* μ_k along an edge $e_k \in E$ with endpoints i and j as the long-run average number of matches between a class- i item and a class- j item per time unit, that is:

$$\left(\sum_{i \in V} \lambda_i \right) \times \frac{1}{t} M_{t,k} \xrightarrow[t \rightarrow +\infty]{\text{almost sure}} \mu_k, \quad k \in E. \quad (9)$$

This quantity is uniquely defined according the ergodic theorem [30, Theorem 1.10.2]. Note that $M_{t,k}/t$ is the average number of matches *per arrival* (out of the first t arrivals), and the factor $(\sum_{i \in V} \lambda_i)$ only serves to convert this result into the average number of matches *per time unit*. The vector of matching rates associated with the model (G, λ, Φ) is denoted by $\mu = (\mu_1, \dots, \mu_m) \in \mathbb{R}_{\geq 0}^m$. If confusion can arise, we will use the notation $\mu(\Phi)$ or $\mu(G, \lambda, \Phi)$.

► **Remark 2.5 (Zero-rate edges).** The following remark will play a key role in Sections 5.2.1 and 6. The matching rate along an edge $e_k \in E$ with endpoints i and j is given by

$$\mu_k = \lambda_i \sum_{s, s' \in \mathcal{S}} \pi(s) \Phi(s, i, j, s') + \lambda_j \sum_{s, s' \in \mathcal{S}} \pi(s) \Phi(s, j, i, s'),$$

where π is the equilibrium distribution of the Markov chain $(S_t, t \in \mathbb{N})$. Since $\lambda_i > 0$, $\lambda_j > 0$, and the distribution π is positive on its support \mathcal{S} , it follows that the matching rate μ_k is zero if and only if $\Phi(s, i, j, s') = \Phi(s, j, i, s') = 0$ for each $s, s' \in \mathcal{S}$, that is, items of classes i and j are *never* matched with one another under policy Φ . In this case, the policy Φ is also adapted to the subgraph $\underline{G} = (V, \underline{E})$ with $\underline{E} = E \setminus \{e_k\}$.

2.5 Implementation and numerical results

To support the results presented in this paper, we used the Python package *Stochastic Matching* [27]. It provides tools to analyze or simulate the behavior of stochastic matching models. We used it to check the characteristics of the examples presented in Section 5 and to perform simulations in Section 6. Unless otherwise stated, a matching model is evaluated by simulating 10^{10} arrivals.

3 Graph theory and linear algebra

In this section, we present a central result of this paper, the connection between stability and the structure of the compatibility graph. In details, Section 3.1 introduces the conservation equation, a system of linear equations satisfied by all vectors of matching rates. Section 3.2 introduces the related concepts of surjectivity, injectivity, and bijectivity that will play a key role throughout the paper. In Sections 3.3 and 3.4, we combine these concepts to formulate new necessary and sufficient conditions under which a compatibility graph G or a matching problem (G, λ) is *stabilizable* in the sense of Definition 2.4. Section 3.5 illustrates these results with early examples.

3.1 Conservation equation

Computing the matching rate vector achieved by a given stable policy or characterizing the set of matching rate vectors that can be achieved by stable policies is a difficult problem *a priori*. To circumvent this difficulty, we first establish a necessary condition known as the *conservation equation*. This equation is satisfied by the matching rate vectors achieved by all stable policies. It asserts that, in a stable system, the arrival of items and their departure due to matches balance each other in the long run. More formally, given a stable matching model (G, λ, Φ) , the conservation equation (CE) can be derived by dividing (3) by t and taking the limit as t tends to infinity, and it can be written in two equivalent forms, either as a system of linear equations (CE-1) or in matrix form (CE-2):

$$\sum_{k \in E_i} \mu_k = \lambda_i, \quad i \in V, \tag{CE-1}$$

$$A\mu = \lambda, \tag{CE-2}$$

where the $n \times m$ matrix $A = (a_{i,k})_{i \in \{1,2,\dots,n\}, k \in \{1,2,\dots,m\}}$ is the *incidence* matrix of the graph G , defined by $a_{i,k} = 1$ if edge e_k is incident to node v_i and $a_{i,k} = 0$ otherwise. The use of a conservation equation is a conventional approach in queueing theory, but our

| General notation | |
|---|---|
| $\mathbb{N}, \mathbb{N}_{>0}, \mathbb{R}, \mathbb{R}_{\geq 0}, \mathbb{R}_{>0}$ | Sets of non-negative integers, positive integers, real numbers, non-negative real numbers, positive real numbers. |
| $\geq, \leq, >, <$ | Coordinate-wise comparison in \mathbb{R}^n . |
| $ \mathcal{A} $ | Cardinality of the set \mathcal{A} . |
| Graph notation | |
| $G = (V, E)$ | Simple graph G with $ V = n$ vertices and $ E = m$ edges. |
| v_i | Vertex indexed by i (denoted i if there is not ambiguity). |
| e_k, k , or $\{i, j\}$ | Edge indexed by k , with endpoints vertices i and j . |
| \mathbb{I} | Family of independent sets of the graph G . |
| $V_i \subseteq V$ | Set of neighbors of node v_i in the graph G . |
| $V(\mathcal{I}) = \bigcup_{i \in \mathcal{I}} V_i$ | Set of neighbors of the vertices indexed by the independent set $\mathcal{I} \in \mathbb{I}$. |
| $d_{i,j}$ | Distance between nodes i and j . |
| $d_{i,k} = \min(d_{i,j}, d_{i,j'})$ | Distance between node i and edge k with endpoints j and j' . |
| K_ℓ | Complete graph on $\ell \geq 3$ nodes. |
| C_ℓ | Cycle on $\ell \geq 3$ nodes. |
| P_ℓ | Path of length $\ell \geq 0$. |
| $KP_{\ell,r,p}$ | Kayak paddle: two cycles C_ℓ and C_r attached by a path P_p . |
| Linear-algebra notation | |
| $x = (x_1, x_2, \dots, x_n)$ | Arbitrary vector in \mathbb{R}^n . All vectors in \mathbb{R}^n are column vectors. |
| $y = (y_1, y_2, \dots, y_m)$ | Arbitrary vector in \mathbb{R}^m . All vectors in \mathbb{R}^m are column vectors. |
| $A = (a_{i,k})_{i \in V, k \in E}$ | Incidence matrix of the graph G . |
| $A^\top = (a_{k,i})_{k \in E, i \in V}$ | Transpose of the matrix A . |
| $\ker(A) = \{y \in \mathbb{R}^m : Ay = 0\}$ | Right kernel of the matrix A . Its dimension is called the nullity of A . |
| $\ker(A^\top) = \{x \in \mathbb{R}^n : A^\top x = 0\}$ | Left kernel of the matrix A . Its dimension is the nullity of A^\top . |
| $d = m - n$ | Dimension of the right kernel of the matrix A if G is surjective. |
| $\mathcal{B} = (b_1, \dots, b_d)$ | A basis of the right kernel of the matrix A if G is surjective. |
| $\Pi = \{y \in \mathbb{R}^m : Ay = \lambda\}$ | Affine space of the solutions of the conservation equation (CE-2). |
| $\Pi_{\geq 0} = \{y \in \mathbb{R}_{\geq 0}^m : Ay = \lambda\}$ | Polytope of non-negative solutions of (CE-2). |
| $\Pi_{> 0} = \{y \in \mathbb{R}_{> 0}^m : Ay = \lambda\}$ | Set of positive solutions of (CE-2). |
| Matching notation | |
| $\lambda = (\lambda_i)_{1 \leq i \leq n}$ | Vector of arrival rates of the item classes. |
| Φ | A matching policy. |
| $\mu = (\mu_k)_{1 \leq k \leq m} = (\mu_{i,j})_{\{i,j\} \in E}$ | Vector of matching rates along the edges. |
| $\Pi_{\mathcal{P}}$ | Set of matching rates achieved by stable policies. |
| $\Pi_{\mathcal{G}}$ | Set of matching rates achieved by stable greedy policies. |

■ **Table 1** Table of notation.

contributions in this paper primarily stem from the novel mixed approach that combines graph theory and linear algebra, which we will elaborate on in the following sections.

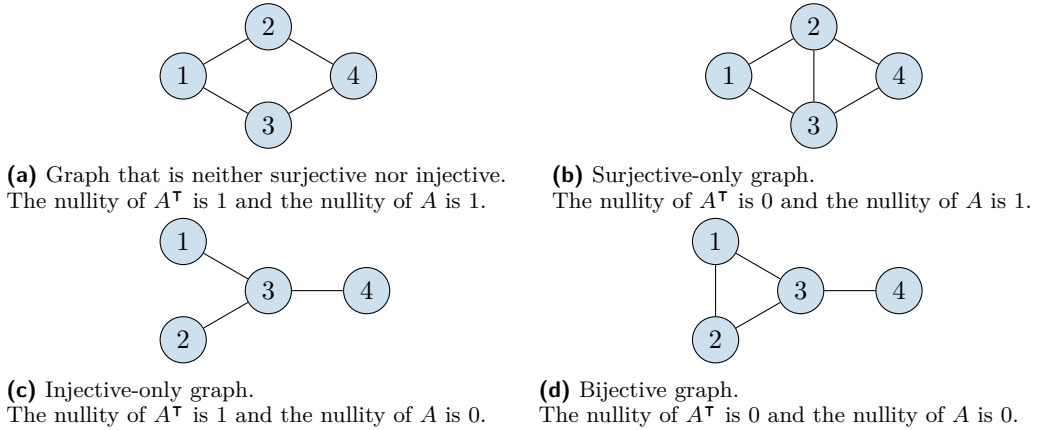
The conservation equation (CE) holds significant importance throughout this paper. While our primary focus is to understand the matching rate vector, we often find it beneficial to temporarily depart from interpreting λ and μ as vectors of arrival and matching rates in a matching model. Instead, we view the conservation equation as a linear equation with λ as a free variable and μ as an unknown. This perspective allows us to explore various aspects of

the equation and its implications. In this context, we will sometimes allow the coordinates of λ and μ to be negative, even if the vectors of arrival and matching rates in a matching model have non-negative coordinates. More specifically, the remainder of the paper focuses on the solutions of (CE) and considers, for a given stabilizable problem (G, λ) , the following sets:

$$\begin{aligned}\Pi &= \{y \in \mathbb{R}^m : Ay = \lambda\} \text{ (studied in Section 5.1),} \\ \Pi_{\geq 0} &= \{y \in \mathbb{R}_{\geq 0}^m : Ay = \lambda\} \text{ (studied in Section 5.2),} \\ \Pi_{> 0} &= \{y \in \mathbb{R}_{> 0}^m : Ay = \lambda\} \text{ (introduced in Section 6),} \\ \Pi_G &= \{\mu(G, \lambda, \Phi) : \Phi \text{ is a stable greedy policy adapted to } G\} \text{ (studied in Section 6.1),} \\ \Pi_{\mathcal{P}} &= \{\mu(G, \lambda, \Phi) : \Phi \text{ is a stable policy adapted to } G\} \text{ (studied in Sections 5.2.1 and 6.2).}\end{aligned}$$

3.2 Surjectivity, injectivity, and bijectivity

Definitions 3.1–3.4 below introduce the notions of *surjectivity*, *injectivity*, and *bijectivity* of a graph. In a nutshell, a compatibility graph G is said to be surjective (resp. injective, bijective) if the linear application $\mu \mapsto A\mu$ defined by its incidence matrix A is surjective (resp. injective, bijective). Interestingly, simple equivalent conditions exist in terms of the graph structure. As we will see later, these notions are fundamental to study the stability of stochastic matching models and the associated matching rate vector. In particular, we will see that a compatibility graph G is stabilizable if and only if G is surjective in the sense of Definition 3.1. Later, we will see that the matching rates in a matching problem (G, λ) are independent of the matching policy Φ (as long as the model is stable) if and only if G is bijective in the sense of Definition 3.3. Examples are shown in Figure 2.



■ **Figure 2** Examples of graphs.

► **Definition 3.1** (Surjective graph). Consider a simple graph $G = (V, E)$ with n nodes and m edges. Let A denote the $n \times m$ incidence matrix of G . The graph G is called surjective if one of the following equivalent conditions is satisfied:

- (i) The function $\mu \in \mathbb{R}^m \mapsto A\mu \in \mathbb{R}^n$ is surjective.
- (ii) For each $\lambda \in \mathbb{R}^n$, the equation $A\mu = \lambda$ of unknown $\mu \in \mathbb{R}^m$ has at least one solution.
- (iii) The left kernel of the matrix A is trivial.
- (iv) Each connected component of the graph G is non-bipartite.

Proof. The equivalence of (i), (ii), and (iii) is a well-known result in linear algebra. We prove that conditions (iii) and (iv) are equivalent. This proof is adapted from [17, Lemma 2.2.3].

The key argument consists of observing that a vector $x = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n$ belongs to the left kernel of the matrix A if and only if

$$\sum_{i=1}^n x_i a_{i,k} = 0, \quad k \in \{1, 2, \dots, m\}.$$

For each $k \in \{1, 2, \dots, m\}$, the k -th equation reads $x_j = -x_i$, where i and j are the endpoints of edge k . An induction argument shows that, for every path i_1, i_2, \dots, i_k in the graph G , we have $x_{i_p} = (-1)^{p-1} x_{i_1}$ for each $p \in \{1, 2, \dots, k\}$.

First assume that condition (iv) is satisfied. Let $x \in \mathbb{R}^n$ be a vector of the left kernel of the matrix A . Since each connected component of G is non-bipartite, for each $i \in V$, there exists a path of length say ℓ that connects node i to a cycle $i_1, i_2, \dots, i_p, i_{p+1} = i_1$ consisting of an odd number p of nodes. We then obtain $x_i = (-1)^{\ell+p+\ell} x_{i_1} = -x_{i_1}$, which implies that $x_i = 0$. Therefore, the left kernel of A is trivial, meaning that condition (iii) is satisfied.

On the contrary, if condition (iv) is not satisfied, then there exists a connected component of G that is bipartite with parts V_+ and V_- . We build a non-zero vector in the left kernel of A by choosing $x_i = 1$ for each $i \in V_+$, $x_i = -1$ for each $i \in V_-$, and $x_i = 0$ for each $i \in V \setminus (V_+ \cup V_-)$. This implies that condition (iii) is not satisfied. \blacktriangleleft

► Definition 3.2 (Injective graph). Consider a simple graph $G = (V, E)$ with n nodes and m edges. Let A denote the $n \times m$ incidence matrix of G . The graph G is called injective if one of the following equivalent conditions is satisfied:

- (i) The function $\mu \in \mathbb{R}^m \mapsto A\mu \in \mathbb{R}^n$ is injective.
- (ii) For each $\lambda \in \mathbb{R}^n$, the equation $A\mu = \lambda$ of unknown $\mu \in \mathbb{R}^m$ has at most one solution.
- (iii) The right kernel of the matrix A is trivial.
- (iv) Each connected component of the graph G contains at most one odd cycle and no even cycle; in other words, each connected component of G is either a tree or a unicyclic graph with an odd cycle.

Proof. The equivalence of conditions (i), (ii), and (iii) is a well-known result in linear algebra. We now prove that conditions (iii) and (iv) are equivalent.

We first assume that the graph G is connected and distinguish the following two cases:

- If G is non-bipartite, according to Definition 3.1, the nullity of A^\top is 0. The rank-nullity theorem implies that the rank of A^\top is n , so that the rank of A is also n . A second application of the rank-nullity theorem implies that the nullity of A is $m - n$. In particular, $\ker(A) = \{0\}$ if and only if $m = n$.
- If G is bipartite, any non-zero vector of the left kernel of A must be parallel (collinear) to the non-zero vector x constructed in the proof of Definition 3.1. This parallelism is due to the constraints $x_i = -x_j$ for all edges $\{i, j\}$. Based on this, the nullity of A^\top is 1, and we conclude from another double application of the rank-nullity theorem that the nullity of A is $m - n + 1$. In particular, $\ker(A) = \{0\}$ if and only if $m = n - 1$.

All in all, we obtain that condition (iii) is true if and only if either the graph G is non-bipartite and contains as many edges as nodes, or the graph G is bipartite and contains one less edge than it contains nodes. This, in turn, is equivalent to condition (iv).

If the graph G is not connected, we can rewrite the matrix A as a bloc matrix in which each bloc corresponds to a connected component, and we can then use the previous argument to prove the equivalence for each connected component. \blacktriangleleft

► **Definition 3.3** (Bijjective graph). *Consider a simple graph $G = (V, E)$ with n nodes and m edges. Let A denote the $n \times m$ incidence matrix of G . The graph G is called bijjective if the following equivalent conditions are satisfied:*

- (i) *The function $\mu \in \mathbb{R}^m \mapsto A\mu \in \mathbb{R}^n$ is bijective.*
- (ii) *For each $\lambda \in \mathbb{R}^n$, the equation $A\mu = \lambda$ of unknown $\mu \in \mathbb{R}^m$ has exactly one solution.*
- (iii) *The matrix A is invertible.*
- (iv) *Each connected component of the graph G contains one cycle and this cycle is odd.*

Proof. The function $\mu \in \mathbb{R}^m \mapsto A\mu \in \mathbb{R}^n$ is bijective if and only if it is both surjective and injective. Hence, the equivalence of conditions (i) to (iv) follows directly from Definitions 3.1 and 3.2. ◀

► **Definition 3.4** (Surjective-only graph and injective-only graph). *A simple graph G is called surjective-only (resp. injective-only) if G is surjective but not injective (resp. injective but not surjective).*

The following proposition gives necessary conditions for surjectivity and injectivity in terms of the number of nodes and edges in the graph.

► **Proposition 3.5.** *Consider an undirected graph $G = (V, E)$ with n nodes and m edges.*

- (i) *If G is surjective, then $n \leq m$.*
- (ii) *If G is injective, then $n \geq m$.*
- (iii) *If G is bijjective, then $n = m$.*
- (iv) *If G is surjective, then G is also injective if and only if $n = m$.*
- (v) *If G is injective, then G is also surjective if and only if $n = m$.*

Proof. These statements (transposed to A) are again well-known in linear algebra. ◀

3.3 Stabilizable compatibility graph

The following proposition gives necessary and sufficient conditions for a graph G to be stabilizable, in terms of either its structure or its incidence matrix.

► **Proposition 3.6.** *Let G be a compatibility graph. The following conditions are equivalent:*

- (i) *The graph G is stabilizable.*
- (ii) *The graph G is surjective.*

Proof. Equivalence between Proposition 3.6 (i) and Definition 3.1(iv) has been proved in [26, Theorem 1]. ◀

In the remainder, we will use the words “stabilizable” and “surjective” interchangeably. Furthermore, unless stated otherwise, we will assume that the graph G is surjective.

Although the equivalence between Proposition 3.6(i) and Definition 3.1(iv) was already proved in [26] in the context of stochastic matching models, no prior literature on stochastic matching models (to the best of our knowledge) has explored the connection between Proposition 3.6(i) and the alternative definitions of surjectivity introduced in Definition 3.1. As we will see later, this new characterization of the stabilizability of a graph G will be useful to analyze the matching rates.

3.4 Stabilizable matching problem

We now turn to the stabilizability of a matching problem (G, λ) . As recalled in Section 2.4, two examples of greedy policies that stabilize the model whenever this matching problem is stabilizable are *match-the-longest* [26] and *first-come-first-matched* [28]. Proposition 3.7 below provides necessary and sufficient conditions for the matching problem (G, λ) to be stabilizable; condition (ii) was already derived in [26], but condition (iii) is new.

► **Proposition 3.7.** *Consider a matching problem (G, λ) with a surjective graph G . The following conditions are equivalent:*

- (i) *The matching problem (G, λ) is stabilizable.*
- (ii) *For each independent set $\mathcal{I} \in \mathbb{I}$, we have $\sum_{i \in \mathcal{I}} \lambda_i < \sum_{i \in V(\mathcal{I})} \lambda_i$.*
- (iii) *The conservation equation (CE) admits a solution $\mu \in \mathbb{R}_{>0}^m$ (i.e., a solution with strictly positive components).*

Proof. Equivalence of (i) and (ii) follows from [26, Proposition 2 and Theorem 2]. For completeness, we observe that [26, Proposition 2] is proved under the assumption that the matching policy Φ is greedy and deterministic and that the state space $(\mathcal{S}, |\cdot|)$ has a particular form, but we can verify that the argument remains valid under the assumptions of Section 2. We now prove that (ii) and (iii) are equivalent. Condition (ii) implies condition (iii) because (a) according to [26], under condition (ii), (G, λ, Φ) is stable if Φ is the *match-the-longest* policy, and (b) the associated vector μ of matching rates satisfies condition (iii) by ergodicity. That condition (iii) implies condition (ii) was proved in [15, Lemma 12]. ◀

One might expect that the time complexity to verify condition (ii) in Proposition 3.7 is exponential in the general case, considering that the number of independent sets grows exponentially with the number n of classes. However, in [26, Proposition 1], it was proven that there exists an $O(n^3)$ -time algorithm for verifying this condition. It is worth noting that this verification process is indirect in the sense that it involves constructing a bipartite double cover of G . In contrast, condition (iii) provides a more direct approach to verify the stabilizability of a matching model (G, μ) . To present this approach more explicitly, we differentiate between two cases based on whether the graph G , assumed to be surjective, is surjective-only or bijective.

► **Remark 3.8.** As observed in [15, Lemma 12], if the graph G is surjective, condition (iii) in Proposition 3.7 gives a simple way of generating vectors $\lambda \in \mathbb{R}_{>0}^n$ such that the problem (G, λ) is stabilizable: it suffices to take $\lambda = A\mu$ for some $\mu \in \mathbb{R}_{>0}^m$. For instance, if $\mu = (\beta, \dots, \beta)$ for some $\beta > 0$, then the coordinates of λ are proportional to the degree of each node.

Verify stabilizability when G is bijective

If the graph G is bijective, then the matrix A is invertible, and (CE) has a unique solution, namely $A^{-1}\lambda$. Proposition 3.7 implies that the matching problem (G, λ) is stabilizable if and only if all coordinates of $A^{-1}\lambda$ are positive. The special case of bijective graphs will be investigated in detail in Section 4, where we will provide a more direct expression for $A^{-1}\lambda$.

Verify stabilizability when G is surjective-only

If the compatibility graph G is surjective-only, (CE) has multiple solutions. To determine if one of these solutions is positive, it suffices to solve a linear optimization problem that

searches for a solution of (CE) whose smallest coordinate is as large as possible:

$$\begin{aligned} & \underset{z=(z_1, z_2, \dots, z_{m+1}) \in \mathbb{R}^{m+1}}{\text{Maximize}} && z_{m+1}, \\ & \text{Subject to} && A(z_1, z_2, \dots, z_m)^\top = \lambda, \\ & && z_i \geq z_{m+1}, \quad i \in \{1, 2, \dots, m\}. \end{aligned} \tag{11}$$

Here, the first m coordinates of the vector z are the coordinates of a vector $\mu \in \mathbb{R}^m$ that satisfies (CE), and the last coordinate of z is a lower bound of the coordinates of this vector μ . Indeed, the equality constraint means that μ satisfies (CE), and the inequality constraint means that the last coordinate of z is less than or equal to its other coordinates. The value to maximize is the last coordinate of the vector z .

If z is a solution of (11), we call the corresponding vector $\mu \in \mathbb{R}^m$ a *maximin* solution of (CE). The optimization problem (11) is a textbook linear optimization problem. It can be solved with a time complexity that is polynomial in the number n of nodes and the number m of edges using many methods, for instance the interior-point-method [23].

The linear optimization problem (11) has a solution with positive coordinates if and only if (CE) has a solution with positive coordinates. According to Proposition 3.7, this is equivalent to saying that the matching problem (G, λ) is stabilizable. Therefore, to verify if a matching problem (G, λ) is stabilizable, it suffices to find a solution of the linear optimization problem (11) and to check if its last coordinate is positive.

► **Remark 3.9.** Observe that the optimization problem (11) always has solutions with finite coordinates. Indeed, the set of vectors that satisfy the constraints of (11) contains at least one valid solution with real-valued coordinates (this is again a consequence of the surjectivity of G). We just need to consider an arbitrary solution μ of (CE) (see Section 5.1.2 for a concrete example using the Moore-Penrose inverse) and to let $z_\mu = (\mu_1, \mu_2, \dots, \mu_m, \min(\mu))$, where $\min(\mu)$ is the smallest coordinate of the vector μ . Any solution better than z_μ has all its coordinates lower-bounded by $\min(\mu)$ and upper-bounded by $\max(\lambda) - \min(0, (n-2)\min(\mu))$. The latter bound is obtained by observing that, if edge k is incident to node i and if $(\mu'_1, \mu'_2, \dots, \mu'_m, x')$ is a solution of (11) such that $x' \geq \min(\mu)$, then $\mu'_k = \lambda_i - \sum_{\ell \in E_i \setminus k} \mu'_\ell$ by (CE). We then use the inequalities $\lambda_i \leq \max(\lambda)$ and $\sum_{\ell \in E_i \setminus k} \mu'_\ell \geq \min(0, (n-2)x') \geq \min(0, (n-2)\min(\mu))$ (this latter inequality is obtained by distinguishing two cases, depending on whether $\min(\mu) \geq 0$ or $\min(\mu) < 0$). Therefore, the solutions better than z_μ belong to a compact set of \mathbb{R}^{m+1} , which ensures the existence of an optimal solution with finite coordinates.

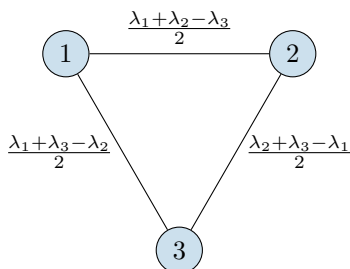
3.5 Early examples

We now provide illustrative examples that demonstrate the concepts introduced in Propositions 3.6 and 3.7 as well as our definitions of surjectivity, injectivity, and bijectivity. These examples will also introduce useful notions that will be further explored in Sections 4–6.

3.5.1 Bijective graphs

We first consider a compatibility graph G that is both surjective and injective. According to Definition 3.3, (CE) has a unique solution for each vector $\lambda \in \mathbb{R}^n$ of arrival rates. Proposition 3.7 implies that the matching problem (G, λ) is stabilizable if and only if the coordinates of this solution are positive. By Remark 3.8, one can easily exhibit a vector $\lambda \in \mathbb{R}_{>0}^n$ of arrival rates that satisfies this condition.

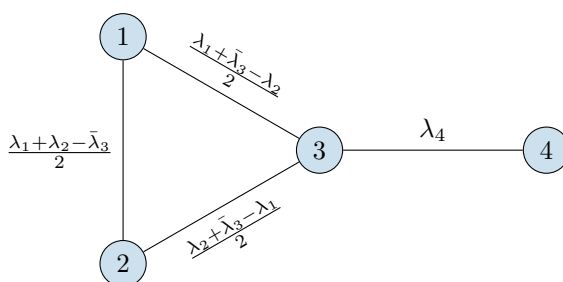
► **Example 3.10** (Triangle). If the graph G is a triangle graph \mathcal{C}_3 , and the vector $\lambda = (\lambda_1, \lambda_2, \lambda_3)$ is given, the solution of (CE) is unique and shown in Figure 3. According to Proposition 3.7(iii), (G, λ) is stabilizable if and only if all coordinates of this solution are positive. This condition is indeed equivalent to Proposition 3.7(ii), which states that $\lambda_1 < \lambda_2 + \lambda_3$, $\lambda_2 < \lambda_1 + \lambda_3$, and $\lambda_3 < \lambda_1 + \lambda_2$. Alternatively, this condition can be expressed as λ_1 , λ_2 , and λ_3 being the lengths of the sides of a non-degenerate triangle (which have to satisfy the triangular inequality). Under these conditions, the model (G, λ, Φ) is stable under the unique greedy policy Φ (as shown in Proposition 6.4).



■ **Figure 3** Matching rates in the triangle graph \mathcal{C}_3 .

► **Example 3.11** (Paw graph). If G is a paw graph, the solution of (CE) is again unique and shown in Figure 4. Here, $\bar{\lambda}_3 = \lambda_3 - \lambda_4$ represents the remaining rate of class 3 after accounting for the needs of class 4. After this subtraction, the matching rates along edges $\{1, 2\}$, $\{1, 3\}$, and $\{2, 3\}$ are defined as in the triangle graph of Figure 3.

It is important to note that, while the existence of a matching rate vector with positive coordinates guarantees the existence of a stable greedy policy like *match-the-longest*, there may still exist greedy policies that lead to instability. For instance, again consider the paw graph of Figure 4 and a *edge-priority* policy where edges $\{1, 3\}$ and $\{2, 3\}$ are preferred over edge $\{3, 4\}$. Intuitively, despite the problem being stabilizable, such a policy may result in a low matching rate $\mu_{3,4}$ due to the priorities: if $\mu_{3,4} < \lambda_4$, this situation would even lead to instability. This observation will be useful in Conjecture 6.10, which requires to develop a greedy policy that does not stabilize a matching problem (G, λ) even when it is stabilizable.



■ **Figure 4** Matching rates in the paw graph. $\bar{\lambda}_3 = \lambda_3 - \lambda_4$ denotes the residual rate that class 3 can provide to classes 1 and 2.

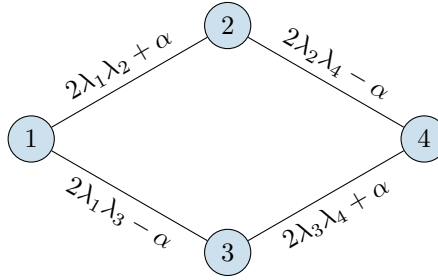
3.5.2 Bipartite graph (that is neither injective nor surjective)

► **Example 3.12** (Square graph). Figure 5 shows a square graph $G = \mathcal{C}_4$. This graph is not surjective because it is bipartite with parts $\{1, 4\}$ (outer part) and $\{2, 3\}$ (inner part).

Therefore, according to Proposition 3.6, this graph is not stabilizable. Yet, given a vector $\lambda = (\lambda_1, \lambda_2, \lambda_3, \lambda_4)$ of arrival rates, (CE) may still have a solution with positive coordinates. This does not contradict Proposition 3.7, as statement (iii) is equivalent to statements (i) and (ii) only if the graph G is surjective. Assuming unit normalization, the conservation equation (CE-2) has a solution if and only if

$$\lambda_1 + \lambda_4 = \lambda_2 + \lambda_3 = \frac{1}{2}. \quad (12)$$

If (12) is satisfied, the solutions of (CE) can be described with a parameter α as shown in Figure 5: starting from the particular solution $\mu = (2\lambda_1\lambda_2, 2\lambda_1\lambda_3, 2\lambda_2\lambda_4, 2\lambda_3\lambda_4)$, all solutions can be generated by alternately adding and subtracting α along the (even) cycle 1–2–4–3–1. The positive solutions correspond to values of α such that $-2 \min(\lambda_1\lambda_2, \lambda_3\lambda_4) < \alpha < 2 \min(\lambda_1\lambda_3, \lambda_2\lambda_4)$.



■ **Figure 5** Matching rates in the square graph C_4 with the normalization $\lambda_1 + \lambda_4 = \lambda_2 + \lambda_3 = \frac{1}{2}$. This graph is not stabilizable.

Focusing on the system dynamics, we can use (3) to prove that the difference in total queue size between the outer part $\{1, 4\}$ and the inner part $\{2, 3\}$ satisfies $Q_{t,1} + Q_{t,4} - (Q_{t,2} + Q_{t,3}) = L_{t,1} + L_{t,4} - (L_{t,2} + L_{t,3})$ for each $t \in \mathbb{N}$. Therefore, the Markov chain $(Q_{t,1} + Q_{t,4} - (Q_{t,2} + Q_{t,3}), t \in \mathbb{N})$ is a random walk on the integer number line $\{\dots, -2, -1, 0, 1, 2, \dots\}$, with transition probability proportional to $\lambda_1 + \lambda_4$ in the increasing direction and to $\lambda_2 + \lambda_3$ in the decreasing direction. If (12) is not satisfied, this random walk is transient, and the difference between the queue sizes of the two parts grows linearly with time. On the other hand, if (12) is satisfied, then the random walk does not have this bias, but the model is still unstable because the corresponding Markov chain is null recurrent⁶.

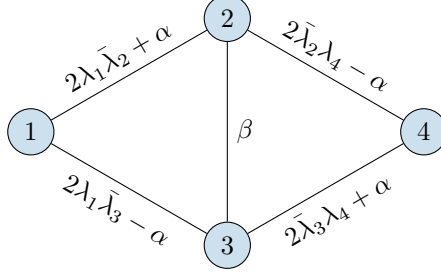
3.5.3 Surjective-only graphs

We finally consider compatibility graphs G that are surjective but not injective. In other words, the graph G is stabilizable and (CE) has an infinite number of solutions. Whether or not these solutions are achievable by a stable matching policy will be discussed in Section 6.

► **Example 3.13** (Diamond (double-fan) graph). Figure 6 shows the diamond graph D , that is, a square graph with an additional edge between nodes 2 and 3. Compared to Example 3.12, this additional edge makes the graph non-bipartite, and therefore surjective, so that the graph is stabilizable according to Proposition 3.6. For ease of computation, we assume that the vector $\lambda = (\lambda_1, \lambda_2, \lambda_3, \lambda_4)$ of arrival rates is normalized so that $\lambda_1 + \lambda_4 = \frac{1}{2}$. Under

⁶ Existing studies of matching in bipartite graphs usually solve this issue by coupling arrivals in both coordinates [2, 9, 13] or by assuming that items have a finite patience time [22].

this assumption, and with $\beta = \frac{1}{2}(\lambda_2 + \lambda_3 - \lambda_1 - \lambda_4) = \frac{1}{2}(\lambda_2 + \lambda_3) - \frac{1}{4}$, $\bar{\lambda}_2 = \lambda_2 - \beta$, and $\bar{\lambda}_3 = \lambda_3 - \beta$, the general solution of (CE) can be described with a parameter α as shown in Figure 6. After subtracting β from λ_2 and λ_3 , the solutions of (CE) for all edges but $\{2, 3\}$ are exactly the same as in the square graph of Example 3.12.



■ **Figure 6** Matching rates in the diamond graph with the normalization $\lambda_1 + \lambda_4 = \frac{1}{2}$. $2\beta = \lambda_2 + \lambda_3 - \lambda_1 - \lambda_4 = \lambda_2 + \lambda_3 - \frac{1}{2}$ is the difference between the arrival rates of the inner part $\{2, 3\}$ and the outer part $\{1, 4\}$. $\bar{\lambda}_2 = \lambda_2 - \beta$ and $\bar{\lambda}_3 = \lambda_3 - \beta$ represent the residual rates that classes 2 and 3 can provide to classes 1 and 4, and they are such that $\lambda_1 + \lambda_4 = \bar{\lambda}_2 + \bar{\lambda}_3 = \frac{1}{2}$, as in Example 3.12.

According to Proposition 3.7(ii), the matching problem (D, λ) is stabilizable if and only if

$$\lambda_2 < \lambda_1 + \lambda_3 + \lambda_4, \quad \lambda_3 < \lambda_1 + \lambda_2 + \lambda_4, \quad \lambda_1 + \lambda_4 < \lambda_2 + \lambda_3, \quad (13)$$

that is, $\bar{\lambda}_3 > 0$, $\bar{\lambda}_2 > 0$, and $\beta > 0$. If these inequalities are satisfied, the positive solutions correspond to values of α such that $-2 \min(\lambda_1 \bar{\lambda}_2, \bar{\lambda}_3 \lambda_4) < \alpha < 2 \min(\lambda_1 \bar{\lambda}_3, \bar{\lambda}_2 \lambda_4)$. Intuitively, compared to the square graph, stabilizable matching problems (D, λ) have a positive difference of 2β between the arrival rates of the inner part $\{2, 3\}$ and the outer part $\{1, 4\}$. This difference is absorbed by the central edge $\{2, 3\}$, which has matching rate β .

Like Example 3.10 and unlike Example 3.11, the matching model (D, λ, Φ) is stable for every greedy policy Φ provided that (13) is satisfied (this will be shown in Corollary A.2).

► **Example 3.14** (Kayak paddle graph). Figure 7 shows a kayak paddle $G = KP_{3,3,1}$, consisting of two triangles linked by an edge. According to Proposition 3.7, the matching problem (G, λ) is stabilizable if and only if there exists $\alpha > 0$ such that $(\lambda_1, \lambda_2, \lambda_3 - \alpha)$ and $(\lambda_4 - \alpha, \lambda_5, \lambda_6)$ are the vectors of arrival rates of two stabilizable triangle graphs \mathcal{C}_3 .

The solutions of (CE) can be described by varying α as shown in Figure 7. Assuming that the matching problem (G, λ) is stabilizable, the solutions of (CE) with positive coordinates correspond to the values of α such that

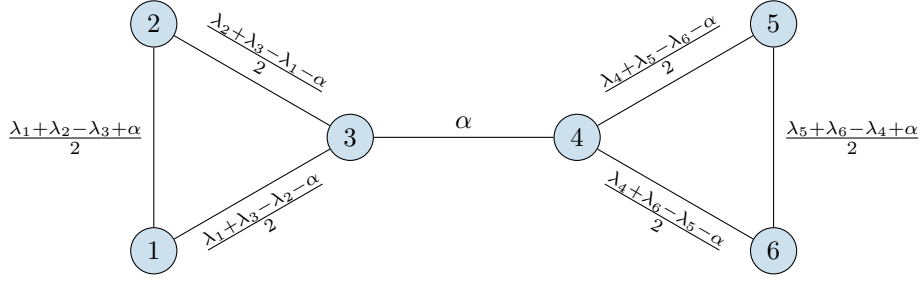
$$0 < \alpha < \min(\lambda_3 - |\lambda_2 - \lambda_1|, \lambda_4 - |\lambda_5 - \lambda_6|).$$

Intuitively, solutions with positive coordinates have a positive matching rate α along edge $\{3, 4\}$. After subtracting this rate from λ_3 and λ_4 , the subgraphs restricted to nodes 1, 2, and 3 and to nodes 4, 5, and 6, respectively, behave like the triangle of Figure 3.

Like Example 3.11 and unlike Examples 3.10 and 3.13, the fact that (G, λ) is stabilizable does not guarantee the stability of all greedy policies.

4 Matching rates in bijective graphs

In the remainder, we will consider exclusively compatibility graphs G that are stabilizable. According to Definition 3.1 and Proposition 3.6, this implies that the graph G is surjective, or equivalently, that each connected component of G is non-bipartite. According to Proposition 3.5(iv) and Proposition 3.7, there are only two possible cases:



■ **Figure 7** Matching rates in the kayak paddle $KP_{3,3,1}$.

1. If $n = m$, the graph G is also bijective: each connected component of the graph G is a unicyclic graph, and its (only) cycle is odd. The conservation equation (CE) has a unique solution given by $\mu = A^{-1}\lambda$. If all components of this solution are positive, then the matching problem (G, λ) is stabilizable, and the solution provides the matching rates achievable by any stable policy.
2. If $n < m$, the graph G is surjective-only: each connected component of the graph G is non-bipartite, and at least one of these connected components contains an even cycle or a pair of odd cycles. The conservation equation (CE) has an infinite number of solutions. If one of them has all-positive components, then the matching problem (G, λ) is stabilizable. Case 1 is studied in this section, while case 2 will be studied in Sections 5 and 6.

In Proposition 4.1 below, we give a simpler expression for the unique solution $\mu = A^{-1}\lambda$ of (CE) in terms of the arrival rate vector λ , under the assumption that the graph G is bijective. We assume without loss of generality that the graph G is connected, as otherwise we can consider each connected component independently. Compared to the expression $\mu = A^{-1}\lambda$, the advantage of Proposition 4.1 is twofold: it does not require calculating a matrix inversion, and it highlights the monotonicity of the matching rates with respect to the arrival rates. This result will be illustrated in Examples 4.2 and 4.3.

► **Proposition 4.1.** *Consider a problem (G, λ) with a compatibility graph $G = (V, E)$ that is connected and bijective, and consider an edge $k \in E$.*

- (i) *If edge k does not belong to the (unique odd) cycle of the graph G , then edge k separates the graph G into two parts, namely a tree and a unicyclic graph. If $V_k \subset V$ denotes the set of nodes that belong to the tree (including one endpoint of edge k), then the matching rate along edge k is given by*

$$\mu_k = \sum_{i \in V_k} (-1)^{d_{i,k}} \lambda_i, \quad (14)$$

where $d_{i,k}$ is the distance between node i and edge k , defined as the minimum distance between node i and an endpoint of edge k .

- (ii) *If edge k belongs to the (unique odd) cycle of the graph G , then the matching rate along edge k is given by*

$$\mu_k = \frac{1}{2} \left(\sum_{i \in V} (-1)^{d_{i,k}} \lambda_i \right). \quad (15)$$

Proof. We first prove (14) for every edge k that does not belong to the cycle. As observed in the proposition, each edge k that does not belong to the cycle separates the graph into two parts, one of which is a tree with node set V_k ; the rooted tree associated with k is obtained

by designating the corresponding endpoint of edge k as the root. We now prove (14) by induction on the height of this rooted tree. Equation (14) is true if the height of this tree is zero. Indeed, in this case, the endpoint of edge k that belongs to the tree, say node i , has no other incident edge, so that applying (CE-1) to node i yields $\mu_k = \lambda_i$, which is consistent with (14). Now assume that the assumption is satisfied for each node whose associated rooted tree has height at most $h - 1$ for some $h \geq 0$, and consider an edge k whose associated rooted tree has height h . By applying (CE-1) to the root i of this associated rooted tree, we obtain

$$\mu_k = \lambda_i - \sum_{\ell \in E_i \setminus \{k\}} \mu_\ell. \quad (16)$$

The induction hypothesis guarantees that (14) is satisfied for every $\ell \in E_i \setminus \{k\}$ (as the height of the associated rooted tree is at most $h - 1$). After injecting this observation to (16), the result for edge k follows by observing that $d_{j,k} = d_{j,\ell} + 1$ for each $\ell \in E_i \setminus \{k\}$ and $j \in V_\ell$, and that $V_k = \{i\} \cup (\bigcup_{\ell \in E_i \setminus \{k\}} V_\ell)$ (all sets being disjoint).

We now prove (15) for each edge k that belongs to the cycle. Since the graph G is unicyclic, deleting edge k from G yields a (connected) tree, which can be seen as a bipartite graph. We let V_+ denote set of nodes in the part that contains both endpoints of edge k (that both endpoints belong to the same part follows from the fact that the cycle is odd) and V_- the set of nodes in the other part. We obtain

$$\sum_{i \in V_+} \lambda_i - \sum_{i \in V_-} \lambda_i = \sum_{i \in V_+} \sum_{\ell \in E_i} \mu_\ell - \sum_{i \in V_-} \sum_{\ell \in E_i} \mu_\ell = 2\mu_k.$$

The first equality follows from (CE-1). The second equality holds because each edge $\ell \in E \setminus \{k\}$ has one endpoint in V_+ and another in V_- , so that μ_ℓ appears exactly once in the first nested sum and once in the second; on the contrary, since both endpoints of edge k belong to V_+ , μ_k appears twice in the first nested sum and zero times in the second. Equation (15) follows by observing that $d_{i,k}$ is even if and only if $i \in V_+$. ◀

We remark that the influence of the arrival rate of a class on the matching rate along an edge only depends on the parity of the distance between the edge and the node. The actual distance does not. In particular, even in a very large (bijective) graph, a node far away from an edge has the same (although possibly reversed) impact as an endpoint of that edge.

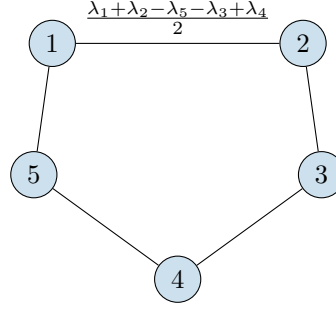
► **Example 4.2** (Cycle graph with 5 nodes). A cycle graph is the simplest bijective graph that we can consider, as it contains an odd cycle and no other edges. In the cycle graph of Figure 8, a direct application of Proposition 4.1(ii) yields

$$\mu_{1,2} = \frac{1}{2}(\lambda_1 + \lambda_2 - \lambda_3 + \lambda_4 - \lambda_5).$$

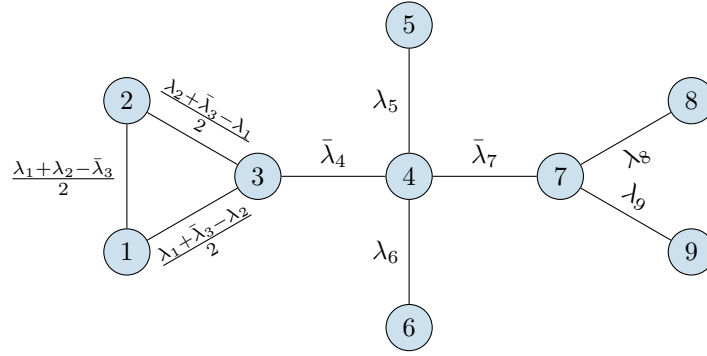
Matching rates along other edges follow by symmetry. From the point of view of edge $\{1, 2\}$, we can partition nodes into two sets, namely $\{1, 2, 4\}$ and $\{3, 5\}$. The former (resp. latter) set contains nodes at an even (resp. odd) distance of edge $\{1, 2\}$, and increasing the arrival rate of these nodes increases (resp. decreases) the matching rate along edge $\{1, 2\}$.

► **Example 4.3** (Lying puppet). We now consider the graph of Figure 9. Edges $\{1, 2\}$, $\{1, 3\}$, and $\{2, 3\}$ belong to the cycle, and the other edges do not. According to Proposition 4.1, we have

$$\mu_{1,2} = \frac{\lambda_1 + \lambda_2 - \bar{\lambda}_3}{2}, \quad \mu_{1,3} = \frac{\lambda_1 - \lambda_2 + \bar{\lambda}_3}{2}, \quad \mu_{2,3} = \frac{-\lambda_1 + \lambda_2 + \bar{\lambda}_3}{2},$$



■ **Figure 8** Matching rates in the pentagon graph C_5 . Only rate $\mu_{1,2}$ is shown for ease of display. The other rates are deduced by permutation.



■ **Figure 9** Matching rates in a “lying puppet” graph with $n = 9$ nodes and $m = 9$ edges. The differences $\bar{\lambda}_7 = \lambda_7 - \lambda_8 - \lambda_9$, $\bar{\lambda}_4 = \lambda_4 - \lambda_5 - \lambda_6 - \bar{\lambda}_7$, and $\bar{\lambda}_3 = \lambda_3 - \bar{\lambda}_4$ are the residual rates that classes 7, 4, and 3 provide to their neighbors of lower index.

where $\bar{\lambda}_3 = \lambda_3 - \mu_{3,4}$, and

$$\begin{aligned} \mu_{4,5} &= \lambda_5, & \mu_{4,6} &= \lambda_6, & \mu_{7,8} &= \lambda_8, \\ \mu_{7,9} &= \lambda_9, & \mu_{4,7} &= \lambda_7 - \mu_{7,8} - \mu_{7,9}, & \mu_{3,4} &= \lambda_4 - \mu_{4,5} - \mu_{4,6} - \mu_{4,7}. \end{aligned}$$

This second set of equations can be obtained either by a direct application of (14) or by applying (CE-1) recursively from the leaves. Indeed, applying (CE-1) to nodes 5, 6, 8, and 9 gives directly the values of $\mu_{4,5}$, $\mu_{4,6}$, $\mu_{7,8}$, and $\mu_{7,9}$, then applying (CE-1) to node 7 gives the value of $\mu_{4,7}$, and finally applying (CE-1) to node 4 gives the value of $\mu_{3,4}$. The values of $\mu_{1,2}$, $\mu_{1,3}$, and $\mu_{2,3}$ are similar to Example 3.11, where the arrival rate λ_3 is again replaced with the effective arrival rate $\bar{\lambda}_3$ from the point of view of classes 1 and 2.

5 Solution of the conservation equation in surjective-only graphs

Consider a surjective-only compatibility graph G . According to Definitions 3.1, 3.2, and 3.4, each connected component of the graph G is non-bipartite, and at least one of these connected components contains an even cycle or a pair of odd cycles. This means that (CE) has an infinite number of solutions. This section describes these solutions. Section 5.1 characterizes the affine space of all real-valued solutions of (CE). Section 5.2 focuses on stabilizable matching problems and describes the convex polytope made of the solutions with non-negative coordinates. Whether or not these solutions are achievable by a matching policy will be discussed in Section 6.

Note that most of the results stated in this section are also applicable to bijective graphs. However, due to the uniqueness of the solution, which is already provided in a closed-form expression in Section 4, bijective graphs hold limited interest in this context. This is why our focus primarily lies on surjective-only graphs.

5.1 Affine space of real-valued solutions

We first consider the set of solutions of (CE) with real-valued (positive, zero, or negative) coordinates:

$$\Pi = \{\mu \in \mathbb{R}^m : A\mu = \lambda\}. \quad (17)$$

In the following, we delve into the properties of Π . In Section 5.1.1, we observe that Π is an affine space of dimension $d = m - n$ that can be described as a translation of the kernel of the incidence matrix A by a particular solution of (CE). Section 5.1.2 gives examples of particular solutions that can be used. Section 5.1.3 gives an algorithm to construct a basis for the kernel of the incidence matrix directly from the graph.

5.1.1 Edge basis, kernel basis

The following proposition characterizes the solution set Π of (CE) using the incidence matrix of the compatibility graph. Equation (18) says that, up to translation, this solution set depends only on the structure of the compatibility graph G , while the arrival rate vector λ impacts only the translation vector.

► **Proposition 5.1.** *Consider a surjective-only compatibility graph G , and let A denote the incidence matrix of G . The solution set Π of (CE) is the affine space obtained by translating the kernel $\ker(A)$ of the matrix A by a particular solution μ° of (CE), that is,*

$$\Pi = \{\mu^\circ + \mu : \mu \in \ker(A)\}. \quad (18)$$

Furthermore, the vector space $\ker(A)$ and the affine space Π have dimension $d = m - n$.

Proof. That the set Π is of the form (18) is a well-known result in linear algebra. Definition 3.1 about surjectivity implies that the rank of A is n , and we conclude from the rank-nullity theorem that the nullity of A is $d = m - n$. The affine space Π has the same dimension according to (18). ◀

Thanks to Proposition 5.1, given a particular solution μ° of (CE) and a basis $\mathcal{B} = (b_1, b_2, \dots, b_d)$ of $\ker(A)$, we can rewrite the affine space Π as

$$\Pi = \{\mu^\circ + \alpha_1 b_1 + \alpha_2 b_2 + \dots + \alpha_d b_d : \alpha_1, \alpha_2, \dots, \alpha_d \in \mathbb{R}\}.$$

In fact, we can define the following affine isomorphism between the coordinate space \mathbb{R}^d and the d -dimensional affine space Π :

$$\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d) \in \mathbb{R}^d \mapsto \mu = \mu^\circ + \alpha_1 b_1 + \alpha_2 b_2 + \dots + \alpha_d b_d \in \Pi. \quad (19)$$

Therefore, given a particular solution μ° of (CE) and a basis $\mathcal{B} = (b_1, b_2, \dots, b_d)$ of $\ker(A)$, there are two natural bases to represent vectors in Π :

- **Edge basis:** A vector of Π is described by its canonical coordinates $\mu = (\mu_1, \mu_2, \dots, \mu_m) \in \mathbb{R}^m$, where μ_k represents a candidate matching rate along edge k , for each $k \in E$.

■ **Kernel basis:** A vector of Π is described by its coordinates $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d) \in \mathbb{R}^d$ in the basis \mathcal{B} , where $d = m - n$ is the dimension of the affine space Π .

Both bases have advantages. The edge basis, by definition, gives directly the candidate matching rates. The kernel basis allows us to work in lower dimension (d instead of m) and to ignore the conservation equation (CE), which is implicitly enforced. In the remainder, we will often use interchangeably the edge coordinates $\mu = (\mu_1, \mu_2, \dots, \mu_m)$ and the kernel coordinates $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d)$ to describe a given vector in Π . Which basis we are actually using will be made clear by our choice of letters (either μ or α).

For graphs that have a low kernel dimension d , it is convenient to mix both approaches and to represent a generic vector of Π , i.e., a generic solution of (CE), in the form $\mu^\circ + \alpha_1 b_1 + \alpha_2 b_2 + \dots + \alpha_d b_d$. For instance, the solutions on Examples 3.13 and 3.14 are actually displayed in Figures 6 and 7 using that convention. This representation, along with the possibility to switch between edge basis and kernel basis, will be used extensively in Sections 5.2 and 6.

► **Remark 5.2 (Change-of-basis formulas).** If B is the $m \times d$ matrix giving the coordinates of the vectors of the basis \mathcal{B} in the edge basis, the change-of-basis formulas are as follows:

- A vector of Π with coordinates α in kernel basis has coordinates $\mu^\circ + B\alpha$ in edge basis;
- A vector of Π with coordinates μ in edge basis has coordinates $B^+(\mu - \mu^\circ)$ in kernel basis, where B^+ is the pseudo-inverse (or Moore-Penrose inverse) of B .

The columns of the matrix B are linearly independent because \mathcal{B} is a basis, so that the pseudo-inverse B^+ has the simple expression $B^+ = (B^\top B)^{-1} B^\top$, where the $d \times d$ matrix $B^\top B$ is invertible because $\ker(B^\top B) = \ker(B) = \{0\}$.

5.1.2 Particular solution

We briefly discuss the derivation of a particular solution μ° of (CE), which can be used as an *origin* for Π . The choice of a basis \mathcal{B} of the kernel of the incidence matrix of the compatibility graph G will be discussed in Section 5.1.3.

Pseudoinverse

A standard way to simultaneously find a particular solution μ° and characterize $\ker(A)$ consists of using the pseudoinverse (or Moore-Penrose inverse) of the matrix A . Definition 3.1 about surjectivity implies that the rows of A are linearly independent, so that the pseudoinverse A^+ of A has the following simple expression:

$$A^+ = A^\top (AA^\top)^{-1},$$

where the $n \times n$ matrix AA^\top is invertible because $\ker(AA^\top) = \ker(A^\top) = \{0\}$. We can then describe a particular solution μ^\perp and the kernel $\ker(A)$ as follows:

$$\mu^\perp = A^+ \lambda, \quad \ker(A) = \{(\text{Id}_{m \times m} - A^+ A) \mu : \mu \in \mathbb{R}^m\}, \quad (20)$$

where $\text{Id}_{m \times m}$ is the m -dimensional identity matrix. The vector μ^\perp is the *least-squares solution* of (CE), and it is orthogonal to $\ker(A)$. In general, some coordinates of this solution can be negative even if non-negative solutions exist. For example, if G is the diamond graph D of Example 3.13, then the matching problem (D, λ) with $\lambda = (4, 5, 2, 1)$ is stabilizable ($\mu = (\frac{7}{2}, \frac{1}{2}, 1, \frac{1}{2}, \frac{1}{2})$ is a positive solution), but the solution given by the pseudoinverse is $\mu^\perp = (\frac{11}{4}, \frac{5}{4}, 1, \frac{5}{4}, -\frac{1}{4})$.

Equation (20) shows that the Moore-Penrose inverse also provides an implicit characterization of $\ker(A)$. However, this characterization is not very practical as it relies on a projection from \mathbb{R}^m to $\ker(A)$. In Section 5.1.3, we will give a more direct characterization by building a basis for $\ker(A)$ based on the structure of the compatibility graph G .

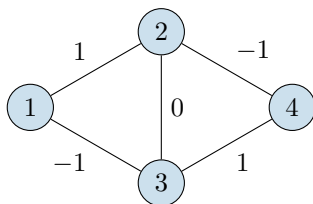
Maximin solution

Alternatively, we can utilize a maximin solution by solving the linear optimization problem (11). Recall that the maximin solution allows us to verify if the matching problem (G, λ) is stabilizable by checking if all edge coordinates of μ° are positive.

5.1.3 Basis of the kernel of the incidence matrix

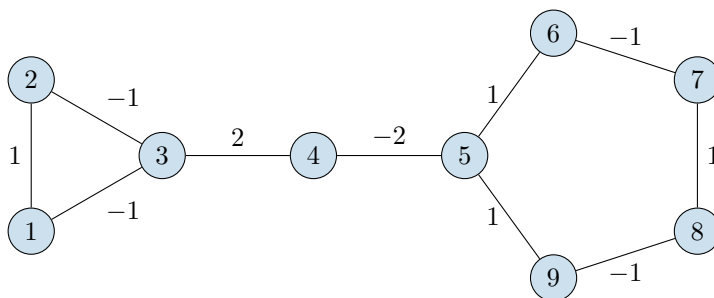
Recall that a vector $y \in \mathbb{R}^m$ belongs to $\ker(A)$ if and only if $Ay = 0$, which reads $\sum_{k \in E_i} y_k = 0$ for each $i \in V$. In other words, a vector $y \in \mathbb{R}^m$ belongs to $\ker(A)$ if and only if, for each $i \in V$, the sum of the coordinates of y associated with the edges that are incident to node i is zero. Using this observation, we first give examples of subgraphs that are supports of vectors that belong to $\ker(A)$, and then we give an algorithm that generates a basis $\mathcal{B} = (b_1, b_2, \dots, b_d)$ of $\ker(A)$.

First observe that an even cycle, if it exists, is the support of a vector in $\ker(A)$: it suffices to assign alternatively the values $+1$ and -1 to the edges of this cycle and the value 0 to all other edges. In the diamond graph of Example 3.13 for instance, if edges are numbered in lexicographical order, then $y = (1, -1, 0, -1, 1)$ is a vector of the unidimensional kernel, associated with the even cycle $1-2-4-3$ (see Figure 10). Intuitively, even cycles can be used to move weights between “odd” and “even” edges without modifying the value of the product Ay . Actually, in this particular example, we can see on Figure 6 that the only way to increase the matching rate along edges $\{1, 2\}$ and $\{3, 4\}$ is if we reduce the matching rate along edges $\{1, 3\}$ and $\{2, 4\}$, and conversely.



■ **Figure 10** Vector of the kernel space of the diamond graph.

Apart from even cycles, other structures of interest are kayak paddles $KP_{\ell,r,p}$ in which the lengths ℓ and r of both cycles are odd. These graphs have a unidimensional kernel, and a base vector can be found by assigning properly the values $+1$ and -1 along the cycles and the values $+2$ and -2 along the path. Figure 11 shows such an assignment for $KP_{3,5,2}$.



■ **Figure 11** Vector of the kernel space of the kayak paddle $KP_{3,5,2}$.

Surprisingly, for any surjective graph G , one can build a basis of $\ker(A)$ using only subgraphs of G that are even cycles and kayak paddles. This is what Algorithm 1 does. It

first identifies an edge set $E \setminus (\mathcal{T} \cup \{a\})$ of cardinality $d = m - n$, where \mathcal{T} is a spanning tree of the graph G and a is an edge in $E \setminus \mathcal{T}$ such that $\mathcal{T} \cup \{a\}$ contains an odd cycle. Then, for each edge $s \in E \setminus (\mathcal{T} \cup \{a\})$, the algorithm builds a base vector b whose support (i) is either an even cycle or a kayak paddle and (ii) contains s and is included into $\mathcal{T} \cup \{a, s\}$. We assume without loss of generality that the graph G is connected (in addition to being surjective-only). If not, we can apply the algorithm to each connected component separately, and then we embed the obtained vectors to \mathbb{R}^n via zero padding.

Data: A connected surjective-only compatibility graph $G = (V, E)$

Result: A basis \mathcal{B} of the kernel of the incidence matrix A of G

```

1  $\mathcal{B} \leftarrow \emptyset$ 
2  $\mathcal{T} \leftarrow$  the set of edges of a spanning tree of  $G$ 
3  $a \leftarrow$  an edge in  $E \setminus \mathcal{T}$  such that  $\mathcal{T} \cup \{a\}$  contains an odd cycle
4 for  $s \in E \setminus (\mathcal{T} \cup \{a\})$  do
5    $b \leftarrow (0, 0, \dots, 0) \in \mathbb{R}^m$ 
6   if  $\mathcal{T} \cup \{a, s\}$  contains an even cycle  $C_\ell$  then
7      $c_1, \dots, c_\ell \leftarrow$  consecutive edges of  $C_\ell$ 
8     for  $d \in \{1, \dots, \ell\}$  do
9        $k \leftarrow$  index of  $c_d$  in  $E$ 
10       $b_k \leftarrow (-1)^d$ 
11   else
12      $\mathcal{T} \cup \{a, s\}$  contains a kayak paddle  $KP_{\ell, r, p}$  with  $\ell$  odd,  $r$  odd, and  $p \geq 0$ 
13      $v_i \leftarrow$  node connecting the kayak's cycle  $C_\ell$  to the kayak central path  $P_p$ 
14      $v_j \leftarrow$  node connecting the kayak's cycle  $C_r$  to the kayak central path  $P_p$ 
15      $c_1, \dots, c_\ell \leftarrow$  consecutive edges of  $C_\ell$ , starting and ending at node  $v_i$ 
16     for  $d \in \{1, \dots, \ell\}$  do
17        $k \leftarrow$  index of  $c_d$  in  $E$ 
18        $b_k \leftarrow (-1)^d$ 
19      $c_1, \dots, c_p \leftarrow$  consecutive edges of  $P_p$ , starting at node  $v_i$  and ending at node  $v_j$ 
20     for  $d \in \{1, \dots, p\}$  do
21        $k \leftarrow$  index of  $c_d$  in  $E$ 
22        $b_k = 2(-1)^{d+1}$ 
23      $c_1, \dots, c_r \leftarrow$  consecutive edges of  $C_r$ , starting and ending at node  $v_j$ 
24     for  $d \in \{1, \dots, r\}$  do
25        $k \leftarrow$  index of  $c_d$  in  $E$ 
26        $b_k \leftarrow (-1)^{d+p+1}$ 
27    $\mathcal{B} \leftarrow \mathcal{B} \cup \{b\}$ 
28 return  $\mathcal{B}$ 

```

■ **Algorithm 1** Construction of a basis of the kernel of the incidence matrix A of the compatibility graph G . This algorithm was initially introduced in [20, Section 3] to build a basis of the eigenspace associated with the eigenvalue -2 of the adjacency matrix of a line graph (i.e., a graph whose nodes and edges represent, respectively, the edges and their incidence relations in another graph).

Figures 12 and 13 show possible runs of Algorithm 1 on the triamond and codomino graphs, which both have a two-dimensional kernel. Note that the basis is not unique and depends on our choice of the spanning tree \mathcal{T} and the augmenting edge a (see lines 2 and 3 in Algorithm 1). We now verify that Algorithm 1 terminates and yields the desired result.

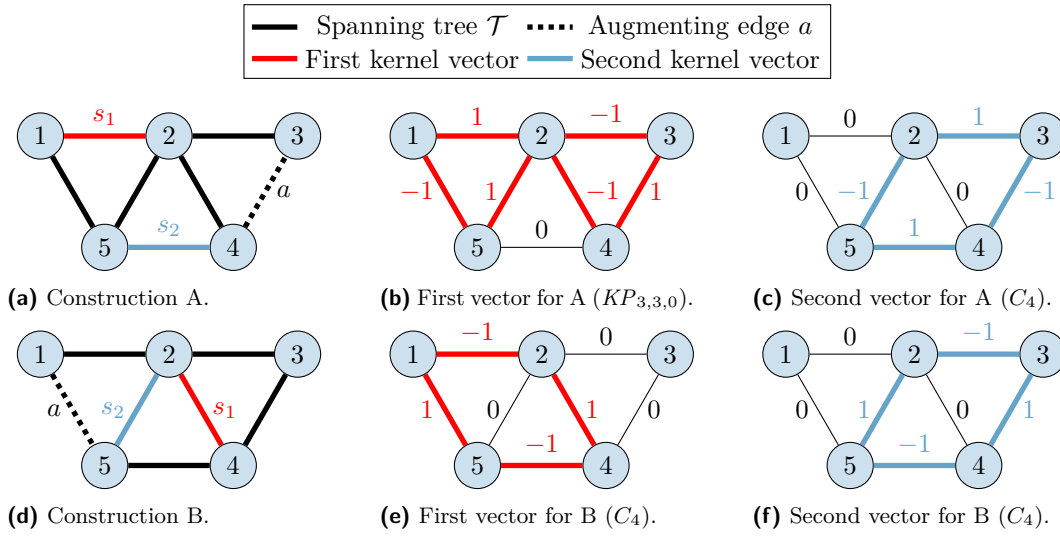


Figure 12 Two possible constructions of a kernel basis for the triamond graph. Construction A yields the basis vectors $b_1 = (1, -1, -1, -1, 1, 1, 0)$ and $b_2 = (0, 0, 1, 0, -1, -1, 1)$. Construction B yields the basis vectors $b_1 = (-1, 1, 0, 1, 0, 0, -1)$ and $b_2 = (0, 0, -1, 0, 1, 1, -1)$.

► **Proposition 5.3.** *Algorithm 1 terminates and returns a basis of the kernel of the incidence matrix A of the compatibility graph G .*

Proof. See Appendix B. ◀

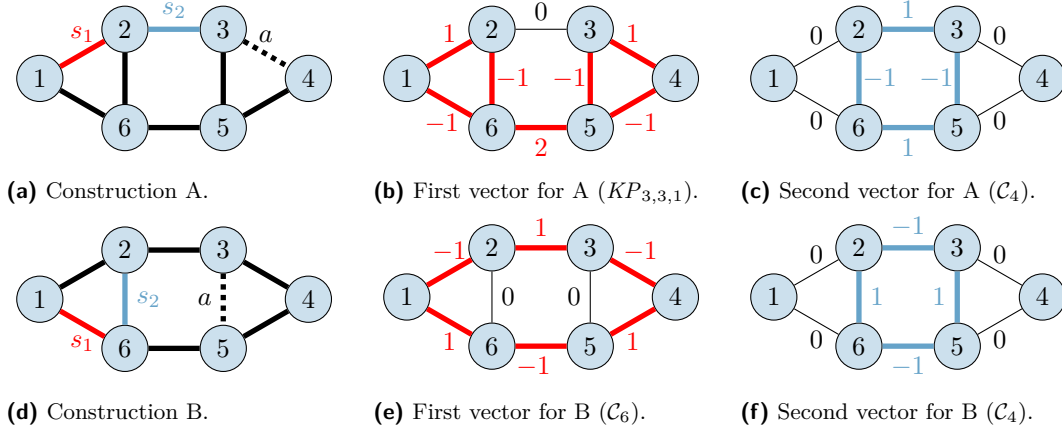
Importantly, equation (18) states that, given an edge $k \in E$, all solutions of (CE) have the same value along edge k if and only if edge k does not belong to the support of any basis vector. According to Algorithm 1, this is equivalent to say that edge k belongs neither to an even cycle nor to a kayak paddle. In the diamond graph of Example 3.13 for instance, the edge $\{2, 3\}$ is the only one that does not belong to the even cycle $1-2-4-3$, and it is also the only one with a fixed value β . In general, if an edge $k \in E$ satisfies this unicity condition, then the matching rate along edge k in a stable matching model (G, λ, Φ) is independent of the policy Φ . Note that there is no general relation between the number of edges with uniquely-defined matching rates and the dimensionality d of the affine space Π .

5.2 Convex polytope of positive solutions

In the remainder, we will focus exclusively on matching problems (G, λ) that are stabilizable. Next, we consider the set $\Pi_{\geq 0}$ of solutions of (CE) that have non-negative coordinates, defined as

$$\Pi_{\geq 0} = \Pi \cap \mathbb{R}_{\geq 0}^m = \{\mu \in \mathbb{R}^m : A\mu = \lambda, \mu \geq 0\}. \quad (21)$$

The set $\Pi_{\geq 0}$ is a d -dimensional convex polytope in \mathbb{R}^m , as it is the intersection of a d -dimensional affine space with the positive orthant $\mathbb{R}_{\geq 0}^m$, both of which are convex. The set $\Pi_{\geq 0}$ is neither empty nor degenerated to a dimension lower than d because the matching problem (G, λ) is assumed to be stabilizable, which means that Π contains a vector with positive coordinates (i.e., in the interior of the positive orthant). It is bounded because each $\mu \in \Pi_{\geq 0}$ satisfies $0 \leq \mu_k \leq \min_{i \in V_k} (\lambda_i)$ for each $k \in E$.



■ **Figure 13** Two possible constructions of a kernel basis for the codomino graph. Construction A yields the vectors $b_1 = (1, -1, 0, -1, 1, -1, -1, 2)$ and $b_2 = (0, 0, 1, -1, 0, -1, 0, 1)$. Construction B yields the vectors $b_1 = (-1, 1, 1, 0, -1, 0, 1, -1)$ and $b_2 = (0, 0, -1, 1, 0, 1, 0, -1)$.

Equation (21) describes $\Pi_{\geq 0}$ in the edge basis. As $\Pi_{\geq 0}$ is a subset of Π , we can also express its elements in the kernel basis introduced in Section 5.1.1. In the kernel basis, $\Pi_{\geq 0}$ is defined by the vectors whose coordinates belong to

$$\tilde{\Pi}_{\geq 0} = \{\alpha \in \mathbb{R}^d : \mu^\circ + \alpha_1 b_1 + \alpha_2 b_2 + \dots + \alpha_d b_d \geq 0\}. \quad (22)$$

As Equations (21) and (22) basically represent the same polytope up to the change-of-basis formulas introduced in Section 5.1.1, in the remainder, we will use the same notation $\Pi_{\geq 0}$ to describe both sets; which basis we are (implicitly) using will be made clear by our choice of letters (μ for the edge basis and α for the kernel basis).

5.2.1 Vertices of the convex polytope

Vertices can be informally defined as the *corners* or extreme points of a convex polytope. Definition 5.4 gives a formal definition of a vertex (as well as those of a face and a facet, which will be useful later).

► **Definition 5.4** (Vertices and facets, adapted from [33]). *Consider a convex polytope Υ of dimension $d \in \mathbb{N}_{>0}$. A (non-empty) face of Υ is a non-empty intersection of Υ with a hyperplane such that Υ is included into one of the two halfspaces defined by the hyperplane. A vertex of Υ is a face of dimension 0, while a facet of Υ is a face of dimension $d - 1$. A vertex can also be defined as follows: a vector $y \in \Upsilon$ is a vertex of Υ if, and only if, it cannot be written as a convex combination of points in $\Upsilon \setminus \{y\}$.*

The vertices of $\Pi_{\geq 0}$ are important if one wants to optimize a linear function of the matching rate vector, as stated in Proposition 5.5 below. Such optimization problems occur in many applications of dynamic matching systems (see the first example in Section 1.1).

► **Proposition 5.5.** *Let $w = (w_1, \dots, w_m) \in \mathbb{R}^m$ be a vector of rewards associated with the edges of the compatibility graph. Consider the problem of finding a non-negative solution of (CE) that maximizes the reward rate $w^\top \mu$:*

$$F = \{\mu \in \Pi_{\geq 0} : w^\top \mu = \max_{z \in \Pi_{\geq 0}} w^\top z\}. \quad (23)$$

F is a non-empty face of $\Pi_{\geq 0}$. In particular, there exists a vertex $\mu \in \Pi_{\geq 0}$ that maximizes the reward (i.e., $\mu \in F$).

Proof. This is a standard result from convex optimization. The fact that $\Pi_{\geq 0}$ is bounded guarantees the existence of a maximum r_{\max} amongst the rewards associated with vectors inside $\Pi_{\geq 0}$. The set F is the intersection of the hyperplane $\{y \in \mathbb{R}^m : w^\top y = r_{\max}\}$ with $\Pi_{\geq 0}$, and it is therefore a non-empty face of $\Pi_{\geq 0}$. The fact that any non-empty face contains a vertex is a consequence of the lattice structure of the faces of polytopes. ◀

The interest of Proposition 5.5 is the following: if for every vertex μ of $\Pi_{\geq 0}$ one can provide a stable policy Φ whose matching rate is μ , or is at least arbitrarily close to μ , then we can maximize any linear reward function (or at least get arbitrarily close to the optimal) with a stable policy. Finding such policies will be the main focus of Section 6.2.

Proposition 5.6 below gives a simple yet powerful characterization of the vertices of $\Pi_{\geq 0}$. The proof of this proposition is borrowed from [14].

► **Proposition 5.6.** *Consider a vector $\mu \in \Pi_{\geq 0}$. Let $\underline{E} = \{k \in E : \mu_k > 0\}$ denote the support of the vector μ and $\underline{G} = (V, \underline{E})$ its support graph. The following statements are equivalent:*

- (i) *The vector μ is a vertex of $\Pi_{\geq 0}$.*
- (ii) *The graph \underline{G} is injective.*

Proof. We actually prove that the negations of (i) and (ii) are equivalent. Let \underline{A} denote the incidence matrix of \underline{G} .

By Definition 5.4, if μ is not a vertex of $\Pi_{\geq 0}$, there exist $z_1, z_2 \in \Pi_{\geq 0} \setminus \{\mu\}$ and $0 < \theta < 1$ such that $\mu = \theta z_1 + (1 - \theta)z_2$. The coordinates of the vectors z_1 and z_2 are non-negative, so this equality implies that their supports are included into the support \underline{E} of the vector μ . In particular, if $\underline{\mu}$ and \underline{z}_1 denote the restrictions of μ and z_1 to coordinates in \underline{E} , respectively, then $\underline{A}\underline{\mu} = \underline{A}\underline{z}_1 = \lambda = \underline{A}z_1 = \underline{A}\underline{z}_1$ with $\underline{\mu} \neq \underline{z}_1$, which means that $\underline{G} = (V, \underline{E})$ is not injective.

Conversely, if \underline{G} is not injective, there exists a non-zero vector \underline{z} in $\mathbb{R}^{|\underline{E}|}$ such that $\underline{A}\underline{z} = 0$. If we embed \underline{z} into $\mathbb{R}^{|E|}$ with zero-padding, we obtain a non-zero vector z such that $Az = 0$, and whose support is included into that of the vector μ . This implies that there exists $\varepsilon > 0$ such that both $\mu - \varepsilon z$ and $\mu + \varepsilon z$ belong to $\Pi_{\geq 0}$. The convex combination $\mu = \frac{1}{2}(\mu - \varepsilon z) + \frac{1}{2}(\mu + \varepsilon z)$ proves that the vector μ is not a vertex of $\Pi_{\geq 0}$. ◀

Examples of vertices are shown in Figures 14–16 and will be discussed in details later. For now, it is sufficient to observe that, since the subgraph restricted to the support of a vertex is injective, this subgraph is either bijective (as in Figures 14d–14h and 16c–16f) or injective-only (as in Figures 15d–15f and 16g).

Proposition 5.7 below characterizes the stabilizability of the matching problem (\underline{G}, λ) for a vertex μ of $\Pi_{\geq 0}$, and relates it to the existence of a policy that stabilizes the matching problem (G, λ) and achieves μ .

► **Proposition 5.7.** *Let $\Pi_{\mathcal{P}}$ be the set of matching rates achieved by stable policies, as defined in Section 3.1:*

$$\Pi_{\mathcal{P}} = \{\mu(G, \lambda, \Phi) : \Phi \text{ is a stable policy adapted to the compatibility graph } G\}.$$

Consider a vertex μ of $\Pi_{\geq 0}$, and define \underline{E} and \underline{G} as in Proposition 5.6. We distinguish two cases, depending on the value of $p = |\underline{E}|$:

- (i) *If $p = n$, i.e., exactly d coordinates of the vertex μ are zero, then the subgraph \underline{G} is bijective and the matching problem (\underline{G}, λ) is stabilizable. Moreover, $\mu \in \Pi_{\mathcal{P}}$.*

- (ii) If $p < n$, i.e., strictly more than d coordinates of the vertex μ are zero, then the subgraph \underline{G} is injective-only and the matching problem (\underline{G}, λ) is not stabilizable. Moreover, $\mu \notin \Pi_{\mathcal{P}}$.

Proof. Let \underline{A} denote the incidence matrix of \underline{G} . We know from Proposition 5.6 that \underline{G} is injective. In particular, the restriction $\underline{\mu}$ of the vector μ to its positive coordinates is the only solution of the conservation equation $\underline{A}\underline{z} = \lambda$, of unknown $\underline{z} \in \mathbb{R}^p$. It also follows from Proposition 3.5 that $p \leq n$, and that the subgraph \underline{G} is bijective if and only if $p = n$. We now consider the two cases separately:

- (i) If $p = n$, then \underline{G} is bijective. Proposition 3.7 implies that the matching problem (\underline{G}, λ) is stabilizable, as \underline{G} is surjective and $\underline{\mu}$ is a solution of the conservation equation $\underline{A}\underline{z} = \lambda$ with positive coordinates. To prove $\mu \in \Pi_{\mathcal{P}}$, we consider the *match-the-longest* policy with a filter on \underline{E} on the matching problem (G, λ) , denoted Φ . Φ behaves exactly like the greedy *match-the-longest* policy on the matching problem (\underline{G}, λ) , which is stable with matching rate $\underline{\mu}$ as we just saw. Hence the matching model (G, λ, Φ) is stable and its matching rate $\mu(\Phi)$ is necessarily equal to μ . This proves that $\mu \in \Pi_{\mathcal{P}}$.
- (ii) If $p < n$, then \underline{G} is injective-only. Proposition 3.6 implies that the matching problem (\underline{G}, λ) is not stabilizable. We prove that $\mu \notin \Pi_{\mathcal{P}}$ by contradiction. Suppose that $\mu \in \Pi_{\mathcal{P}}$, and let Φ be a stable policy on the matching problem (G, λ) such that $\mu(\Phi) = \mu$. Since $\mu_k(\Phi) = 0$ for each $k \in E \setminus \underline{E}$, we know from Remark 2.5 that Φ never performs a match supported by an edge in $E \setminus \underline{E}$. Hence, Φ also defines a stable policy on the matching problem (\underline{G}, λ) , which contradicts the instability of the matching problem (\underline{G}, λ) . Therefore, $\mu \notin \Pi_{\mathcal{P}}$. \blacktriangleleft

In a nutshell, Proposition 5.7 states that the vertices with a bijective support graph can be achieved by stable policies, while those that define an injective-only support graph cannot. This is the first step towards a full characterization of $\Pi_{\mathcal{P}}$, which will be further developed in Section 6.2.

► **Remark 5.8.** We can further explore the relationship between vertices and their support graph. Again consider a vertex μ of $\Pi_{\geq 0}$, and define \underline{G} and \underline{A} as in Proposition 5.6. According to Definitions 3.1, 3.2, and 3.4, the subgraph \underline{G} is injective if and only if each connected component of \underline{G} is either a tree or a unicyclic graph with an odd cycle. For each connected component of \underline{G} that is a tree, and therefore a bipartite graph with parts V_+ and V_- , the existence of vertex μ implies that

$$\sum_{i \in V_+} \lambda_i = \sum_{i \in V_-} \lambda_i. \quad (24)$$

This equation follows by summing (CE-1) over the nodes in V_+ on the one hand, summing (CE-1) over the nodes in V_- on the other hand, and verifying that the left-hand sides of both equations are equal. In fact, one can verify that the vector λ belongs to the image of \underline{A} if and only if λ satisfies (24) for each connected component of \underline{G} that is a tree. Note that this condition is void if \underline{G} is bijective because, in this case, none of the connected components of \underline{G} is a tree.

Conversely, one can wonder which injective subgraph \underline{G} of G defines a vertex of $\Pi_{\geq 0}$. Satisfying (24) for each tree connected component of \underline{G} only guarantees the existence of a (unique) solution $z \in \mathbb{R}^p$ to the conservation equation $\underline{A}z = \lambda$. If each component of z is positive, then we indeed obtain a vertex of $\Pi_{\geq 0}$ by embedding z in \mathbb{R}^m with zero padding; otherwise, \underline{G} does not define a vertex of $\Pi_{\geq 0}$.

5.2.2 Bijective vertices and facets of the convex polytope

With a slight abuse of notation, we will say that a vertex $\mu \in \Pi_{\geq 0}$ is bijective (resp. injective-only) if the subgraph \underline{G} induced by μ is bijective (resp. injective-only). As stated above, the bijectivity of vertices plays an important role in the study of the matching rates that can be achieved by stable policies, which will be the matter of Section 6. We now detail the relationship between the bijectivity of vertices and the inequalities that define $\Pi_{\geq 0}$. While the remainder of this section helps in Section 6.2 to describe the set of matching rate vectors that can be achieved by an arbitrary stable policy, it can be skipped on first reading.

Following Proposition 5.7, the bijectivity of a vertex is determined by the number of its coordinates that are positive in edge coordinates, that is, by the cardinality of the set of edges that form its support. Recall that the d -dimensional polytope $\Pi_{\geq 0}$ is actually characterized by the m inequalities $\mu_k \geq 0$ for each $k \in E$. In particular, this polytope has at most m facets, one for each inequality, but it typically has fewer. Indeed, some inequalities may be redundant and/or not tight, in a sense that will be defined in Definition 5.9 below. For example, by looking more closely at the general solution obtained for the diamond graph in Figure 6, we conclude that:

- The inequality $\mu_{2,3} \geq 0$ is satisfied trivially by every vector $\mu \in \Pi$, as we have $\mu_{2,3} = \beta > 0$. Therefore, this inequality does not define a facet of $\Pi_{\geq 0}$.
- If $\lambda_1 \bar{\lambda}_2 < \bar{\lambda}_3 \lambda_4$, the inequality $\mu_{1,2} \geq 0$ supersedes the inequality $\mu_{3,4} \geq 0$, and conversely. If $\lambda_1 \bar{\lambda}_2 = \bar{\lambda}_3 \lambda_4$, these two inequalities are equivalent. In both cases, the inequalities $\mu_{1,2} \geq 0$ and $\mu_{3,4} \geq 0$ lead to a single facet of $\Pi_{\geq 0}$.
- If $\lambda_1 \bar{\lambda}_3 < \bar{\lambda}_2 \lambda_4$, the inequality $\mu_{1,3} \geq 0$ supersedes the inequality $\mu_{2,4} \geq 0$, and conversely. If $\lambda_1 \bar{\lambda}_3 = \bar{\lambda}_2 \lambda_4$, these two inequalities are equivalent. In both cases, the inequalities $\mu_{1,3} \geq 0$ and $\mu_{2,4} \geq 0$ lead to a single facet of $\Pi_{\geq 0}$.

All in all, the 1-dimensional convex polytope $\Pi_{\geq 0}$ associated with the diamond graph of Example 3.13 has two facets, even if it is defined by five inequalities. Definition 5.9 below will help us relate these notions to the number of zero coordinates of the vertices of the convex polytope $\Pi_{\geq 0}$.

► **Definition 5.9** (Adapted from [4, 33]). *Let $k \in E$.*

- (i) *The inequality $\mu_k \geq 0$ is said to be tight if there exists a vector $\mu \in \Pi_{\geq 0}$ such that $\mu_k = 0$, in which case we also say that this inequality is tight for the vector μ .*
- (ii) *The inequality $\mu_k \geq 0$ is said to be redundant if removing this inequality does not change the polytope $\Pi_{\geq 0}$, in the sense that*

$$\Pi_{\geq 0} = \{\mu \in \mathbb{R}^m : A\mu = \lambda \text{ and } \mu_\ell \geq 0 \text{ for each } \ell \in E \setminus \{k\}\}.$$

Otherwise, this inequality is called irredundant.

- (iii) *The matching problem (G, λ) is called essential if all tight inequalities are irredundant.*
- (iv) *The polytope $\Pi_{\geq 0}$ is said to be simple if every vertex of $\Pi_{\geq 0}$ belongs to exactly d facets, which is the minimal number of facets a vertex belongs to.*

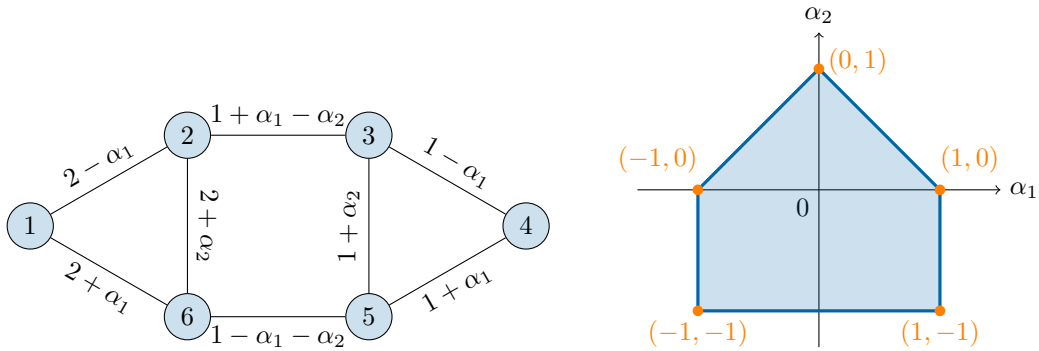
Importantly, the number of positive coordinates of a vertex μ (considered in Proposition 5.7) is the number of inequalities that are not tight for this vertex. More generally, Definition 5.9 has the following intuitive interpretation. An inequality is tight if the convex polytope $\Pi_{\geq 0}$ intersects the hyperplane obtained by transforming this inequality into an equality. Non-tight inequalities are “useless” (and redundant) because they are never satisfied as equalities by any vector in $\Pi_{\geq 0}$. The matching problem (G, λ) is essential if each tight inequality defines a distinct facet of the convex polytope $\Pi_{\geq 0}$. Under this condition, the number of facets

that contain a vertex is equal to the number of inequalities that are tight for this vertex. In particular, as we will see in Proposition 5.13, if the matching problem (G, λ) is essential and the polytope $\Pi_{\geq 0}$ is simple, then every vertex satisfies exactly d (tight) inequalities as equalities, which means that this vertex has d zero coordinates, and therefore $n = m - d$ positive coordinates, so that this vertex is bijective.

All these notions are illustrated in Examples 5.10–5.12 below, which show in particular that a matching problem (G, λ) may be essential even if the polytope $\Pi_{\geq 0}$ is not simple, and conversely. Consistently with Example 3.13 above, these examples use a kernel basis to verify effortlessly whether an inequality is tight and/or irredundant.

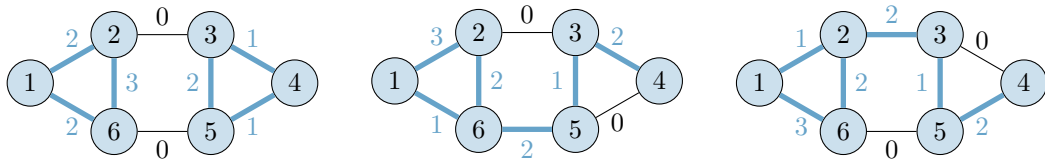
| Edge basis | Kernel basis | Tight? | Irredundant? |
|--------------------|-------------------------------|--------|--------------|
| $\mu_{1,2} \geq 0$ | $\alpha_1 \leq 2$ | ✗ | ✗ |
| $\mu_{1,6} \geq 0$ | $\alpha_1 \geq -2$ | ✗ | ✗ |
| $\mu_{2,3} \geq 0$ | $\alpha_1 - \alpha_2 \geq -1$ | ✓ | ✓ |
| $\mu_{2,6} \geq 0$ | $\alpha_2 \geq -2$ | ✗ | ✗ |
| $\mu_{3,4} \geq 0$ | $\alpha_1 \leq 1$ | ✓ | ✓ |
| $\mu_{3,5} \geq 0$ | $\alpha_2 \geq -1$ | ✓ | ✓ |
| $\mu_{4,5} \geq 0$ | $\alpha_1 \geq -1$ | ✓ | ✓ |
| $\mu_{5,6} \geq 0$ | $\alpha_1 + \alpha_2 \leq 1$ | ✓ | ✓ |

(a) Inequalities.



(b) Generic solution of (CE).

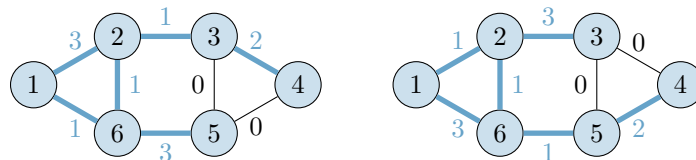
(c) Polytope $\Pi_{\geq 0}$ in kernel basis.



(d) Edge coordinates of $(0, 1)$.

(e) Edge coordinates of $(-1, 0)$.

(f) Edge coordinates of $(1, 0)$.



(g) Edge coordinates of $(-1, -1)$.

(h) Edge coordinates of $(1, -1)$.

■ **Figure 14** An essential matching problem (G, λ) with a simple polytope $\Pi_{\geq 0}$. The vector of arrival rates is $\lambda = (4, 5, 3, 2, 3, 5) \in \mathbb{R}^6$, a particular solution of (CE) is $\mu^\circ = (2, 2, 1, 2, 1, 1, 1, 1) \in \mathbb{R}^8$, and the chosen base vectors for $\ker(A)$ are $b_1 = (-1, 1, 1, 0, -1, 0, 1, -1)$ and $b_2 = (0, 0, -1, 1, 0, 1, 0, -1)$.

► **Example 5.10** (Essential matching problem). Figure 14 considers a codomino graph with the vector of arrival rates $\lambda = (4, 5, 3, 2, 3, 5)$. A particular solution of (CE) is $\mu^\circ = (2, 2, 1, 2, 1, 1, 1, 1) \in \mathbb{R}^8$, and the basis of $\ker(A)$ consists of the vectors $b_1 = (-1, 1, 1, 0, -1, 0, 1, -1)$ and $b_2 = (0, 0, -1, 1, 0, 1, 0, -1)$ obtained in construction B of Figure 13. The generic solution of (CE) is shown in Figure 14b.

The inequalities are listed in Figure 14a. The 2-dimensional polytope $\Pi_{\geq 0}$, shown in Figure 14c in kernel basis, is characterized by five tight inequalities which are also irredundant:

$$-1 \leq \alpha_1 \leq 1, \quad \alpha_2 \geq -1, \quad \alpha_1 - \alpha_2 \geq -1, \quad \alpha_1 + \alpha_2 \leq 1.$$

The matching problem (G, λ) is essential. In kernel basis, the vertices of the convex polytope $\Pi_{\geq 0}$ are $(0, 1)$, $(-1, 0)$, $(1, 0)$, $(-1, -1)$, and $(1, -1)$, and we can verify on Figure 14c that each vertex belongs to exactly 2 facets. Therefore $\Pi_{\geq 0}$ is simple (more generally, all 2-dimensional polytopes are simple). All in all, each vertex of $\Pi_{\geq 0}$ has 2 zero coordinates and 6 positive coordinates in edge coordinates, so that this vertex is bijective. These vertices are represented in edge basis in Figures 14d–14h.

► **Example 5.11** (Non-essential matching problem). Figure 15 shows the same codomino graph as in Example 5.10, with the same basis of $\ker(A)$, but with the vector of arrival rates $\lambda = (2, 4, 4, 2, 2, 2)$. A particular solution of (CE) is $\mu^\circ = (1, 1, 2, 1, 1, 1, 1, 0)$, and the general solution is shown in Figure 15b.

The inequalities are listed in Figure 15a. The 2-dimensional convex polytope $\Pi_{\geq 0}$ is shown in kernel basis in Figure 15c. All inequalities are tight, but only one is irredundant, so we conclude that the matching problem (G, λ) is not essential, even if the polytope $\Pi_{\geq 0}$ is still simple. Correspondingly, even if each vertex belongs to exactly two facets, they all have more than 2 zero coordinates, so none of them is bijective. For example, the vertex $(1, -1)$ in kernel basis has coordinates $(0, 2, 4, 0, 0, 0, 2, 0)$ in edge basis (Figure 15f). This vertex has 5 zero coordinates in edge coordinates (and only 3 positive coordinates) even if it belongs to only 2 facets.

► **Example 5.12** (Non-simple polytope). We finally exhibit an essential matching problem with a non-simple associated polytope. As 2-dimensional polytopes are simple, we need to consider a more complex example. We consider the matching problem of Figure 16a. The arrival rate is $\lambda = (3, 3, 6, 3, 4, 4, 6, 3, 4, 4) \in \mathbb{R}^{10}$. The particular solution and kernel basis are shown on the edges. The set $\Pi_{\geq 0}$, shown in Figure 16b in kernel basis, is an Egyptian pyramid characterized by the following tight inequalities:

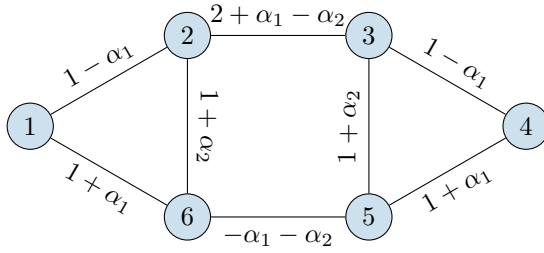
$$\alpha_3 \geq 0, \quad 1 + \alpha_1 - \alpha_3 \geq 0, \quad 1 - \alpha_1 - \alpha_3 \geq 0, \quad 1 + \alpha_2 - \alpha_3 \geq 0, \quad 1 - \alpha_2 - \alpha_3 \geq 0.$$

These five inequalities are irredundant (each one corresponds to exactly one of the five facets), so we conclude that the matching problem (G, λ) is essential. In kernel basis, the vertices of this convex polytope are $(-1, -1, 0)$, $(1, -1, 0)$, $(1, 1, 0)$, $(-1, 1, 0)$, and $(0, 0, 1)$. These vertices are shown in edge basis in Figures 16c–16g. The polytope $\Pi_{\geq 0}$ is not simple because the vertex $(0, 0, 1)$ (the “top” of the pyramid) belongs to 4 facets, while the polytope has dimension 3. Consistently, we can see in Figure 16g that this vertex has 4 zero coordinates and only 9 positive coordinates in edge basis; the subgraph defined by the support of this vertex is injective-only.

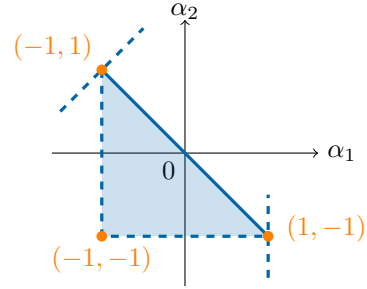
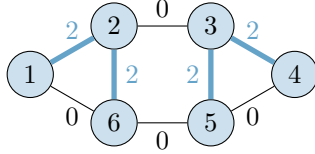
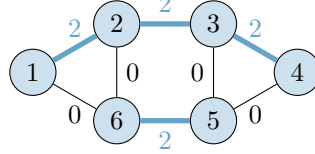
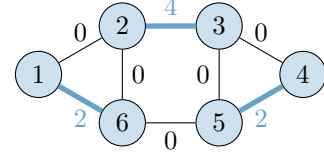
In light of the examples above, we can give the following characterization of the bijective vertices of $\Pi_{\geq 0}$.

| Edge basis | Kernel basis | Tight? | Irredundant? |
|--------------------|-------------------------------|--------|--------------|
| $\mu_{1,2} \geq 0$ | $\alpha_1 \leq 1$ | ✓ | ✗ |
| $\mu_{1,6} \geq 0$ | $\alpha_1 \geq -1$ | ✓ | ✗ |
| $\mu_{2,3} \geq 0$ | $\alpha_1 - \alpha_2 \geq -2$ | ✓ | ✗ |
| $\mu_{2,6} \geq 0$ | $\alpha_2 \geq -1$ | ✓ | ✗ |
| $\mu_{3,4} \geq 0$ | $\alpha_1 \leq 1$ | ✓ | ✗ |
| $\mu_{3,5} \geq 0$ | $\alpha_2 \geq -1$ | ✓ | ✗ |
| $\mu_{4,5} \geq 0$ | $\alpha_1 \geq -1$ | ✓ | ✗ |
| $\mu_{5,6} \geq 0$ | $\alpha_1 + \alpha_2 \leq 0$ | ✓ | ✓ |

(a) Inequalities



(b) Generic solution of (CE).

(c) Polytope $\Pi_{\geq 0}$ in kernel basis. Dashed lines show tight redundant inequalities.(d) Edge coordinates of $(-1, 1)$.(e) Edge coordinates of $(-1, -1)$.(f) Edge coordinates of $(1, -1)$.

■ **Figure 15** Non-essential matching problem (G, λ) with a simple polytope $\Pi_{\geq 0}$. The vector of arrival rates is $\lambda = (2, 4, 4, 2, 2, 2) \in \mathbb{R}^6$, a particular solution of (CE) is $\mu^\circ = (1, 1, 2, 1, 1, 1, 1, 0) \in \mathbb{R}^8$, and the chosen base vectors for $\ker(A)$ are $b_1 = (-1, 1, 1, 0, -1, 0, 1, -1)$ and $b_2 = (0, 0, -1, 1, 0, 1, 0, -1)$.

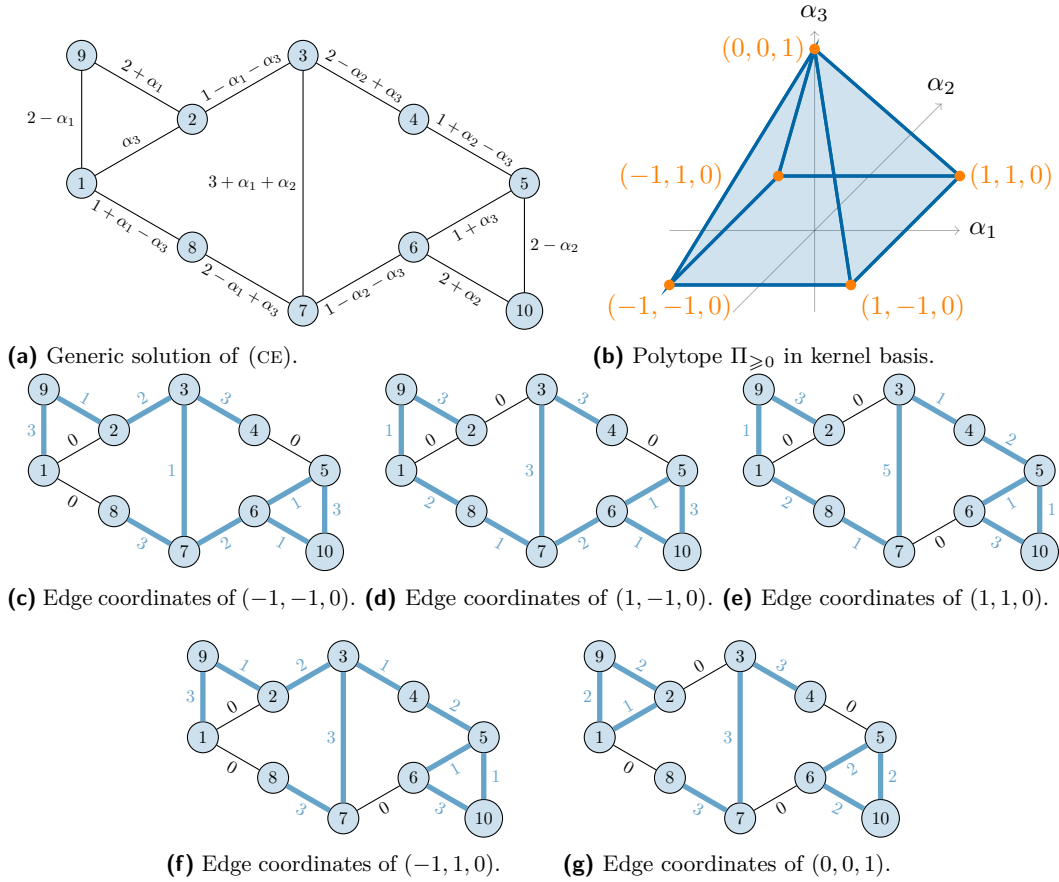
► **Proposition 5.13.** *Let μ be a vertex of $\Pi_{\geq 0}$. The following statements are equivalent:*

- (i) μ is bijective.
- (ii) μ belongs to exactly d facets of $\Pi_{\geq 0}$ and none of the inequalities tight for μ is redundant. In particular, all vertices of $\Pi_{\geq 0}$ are bijective if, and only if, the matching problem (G, λ) is essential and the polytope $\Pi_{\geq 0}$ is simple.

Proof. To prove the equivalence of (i) and (ii), we first remark that the number of zero edge coordinates of a vector $\mu \in \Pi_{\geq 0}$ is by definition the number of inequalities that are tight for μ . It is in particular at least the number of facets that intersect μ , with equality if, and only if, none of the inequalities tight for μ is redundant.

If a vector μ is a vertex of $\Pi_{\geq 0}$, then μ belongs to at least d facets of $\Pi_{\geq 0}$. Now, the vector μ is bijective if, and only if, d of its coordinates are zero, which in view of the remark above is equivalent to say that μ belongs to exactly d facets of $\Pi_{\geq 0}$ and none of the inequalities tight for μ is redundant.

As for the last statement, it follows directly from Definition 5.9. ◀



■ **Figure 16** Essential matching problem (G, λ) with a non-simple polytope $\Pi_{\geq 0}$. The arrival rate is $\lambda = (3, 3, 6, 3, 4, 4, 6, 3, 4, 4) \in \mathbb{R}^{10}$. A particular solution and the chosen base vectors for $\ker(A)$ are implicitly shown on the edges of Figure 16a.

6 Matching rates in surjective-only graphs

Section 5 described the polytope $\Pi_{\geq 0}$, defined as the set of non-negative solutions of (CE). In this section, we investigate which of these solutions may, or may not, be achieved by a (stable) policy. Unless stated otherwise, in this section, we consider a stabilizable matching problem (G, λ) with a surjective-only compatibility graph G , so that the set $\Pi_{\geq 0}$ of non-negative solutions of (CE) is non-trivial.

Given a policy Φ that stabilizes the matching problem (G, λ) , we let $\mu(\Phi)$ (resp. $\alpha(\Phi)$) denote the edge (resp. kernel) coordinates of the vector of matching rates in the model (G, λ, Φ) . We consider the set of matching rates achieved by stable policies (resp. by stable greedy policies), defined as follows:

$$\begin{aligned} \Pi_G &= \{\mu(\Phi) : \Phi \text{ is a greedy policy adapted to the compatibility graph } G\}, \\ \Pi_P &= \{\mu(\Phi) : \Phi \text{ is a stable policy adapted to the compatibility graph } G\}. \end{aligned}$$

Section 6.1 focuses on the set Π_G . We show in particular that $\Pi_G \subseteq \Pi_{>0}$, where $\Pi_{>0}$ is the (non-empty) set of positive solutions of (CE) (that is, solutions with positive coordinates). We also show that, in many cases, the inclusion is strict. In contrast, we observe in Section 6.2 that non-greedy policies can achieve any vector in $\Pi_{>0}$, and that filtering policies can even

approach some faces of $\Pi_{\geq 0}$. More precisely, we conjecture that $\Pi_{>0} \subseteq \Pi_{\mathcal{P}}$ in general, and we prove that $\Pi_{\mathcal{P}} = \Pi_{\geq 0}$ if the support graph of each vertex of $\Pi_{\geq 0}$ is bijective. Informally, these results show that non-greedy policies are more suitable than greedy policies to navigate the set $\Pi_{\geq 0}$.

The following result will be useful throughout this section.

► **Proposition 6.1.** *The sets $\Pi_{\mathcal{G}}$ and $\Pi_{\mathcal{P}}$ are convex.*

Proof. See Appendix C.1. ◀

6.1 Greedy policies

Greedy policies are appealing candidates to control matching problems. Some of them, like *match-the-longest* or *first-come-first-matched*, are easy to implement and stabilize all stabilizable matching problems. Greedy policies may also be more easily accepted socially, especially in situations involving human beings. We show in this section that, unfortunately, in general, greedy policies cannot achieve all achievable matching rates. Indeed, although we do not provide a universal tight description of $\Pi_{\mathcal{G}}$ (which we believe to be a difficult task), we provide two simple examples where $\Pi_{\mathcal{G}}$ is a strict subset of $\Pi_{>0}$. In contrast, Section 6.2 will introduce semi-filtering (non-greedy) policies and show that these policies can be used to achieve every vector in the set $\Pi_{>0}$ and some of the faces of the polytope $\Pi_{\geq 0}$.

Firstly, in Proposition 6.2 below, we show that greedy policies cannot reach the boundary of the convex polytope $\Pi_{\geq 0}$.

► **Proposition 6.2.** *If the compatibility graph G is surjective-only and the matching problem (G, λ) is stabilizable, then the set $\Pi_{\mathcal{G}}$ is non-empty and $\Pi_{\mathcal{G}} \subseteq \Pi_{>0}$.*

Proof. The set $\Pi_{\mathcal{G}}$ is non-empty because, as recalled in Section 2.4, the greedy policies *match-the-longest* and *first-come-first-matched* are stable. We now prove that $\Pi_{\mathcal{G}} \subseteq \Pi_{>0}$. Consider a stable greedy policy Φ and let $\mu = (\mu_1, \mu_2, \dots, \mu_m)$ denote the matching rate vector in the model (G, λ, Φ) . Consider an edge $e_k = \{i, j\}$. Since the policy Φ is greedy, two items of classes i and j are always matched if the following sequence of events occurs: the system is in the empty state \emptyset , then a class- i item arrives, and then a class- j item arrives. Let p_{\emptyset} denote the stationary probability that the model (G, λ, Φ) is in the empty state \emptyset . We know that $p_{\emptyset} > 0$ because the model is stable, and the previous remark implies that $\mu_k \geq p_{\emptyset} \lambda_i \lambda_j / \sum_{\ell \in V} \lambda_{\ell} > 0$. Since this is true for each edge $e_k \in E$ and each $\mu \in \Pi_{\mathcal{G}}$, we conclude that $\Pi_{\mathcal{G}} \subseteq \Pi_{>0}$. ◀

Proposition 6.2 has the following consequence regarding the ability of greedy policies to maximize linear reward functions.

► **Corollary 6.3.** *Consider the problem of maximizing an edge-dependent reward defined by a vector w as stated in Proposition 5.5. Let $r_{\max} = \max_{\mu \in \Pi_{\geq 0}} w^{\top} \mu$ be the optimal reward. One of the two exclusive statements below is true:*

- (i) *All vectors of $\Pi_{\geq 0}$ are optimal, i.e., $w^{\top} \mu = r_{\max}$ for each $\mu \in \Pi_{\geq 0}$.*
- (ii) *All stable greedy policies are suboptimal, i.e., $w^{\top} \mu < r_{\max}$ for each $\mu \in \Pi_{\mathcal{G}}$.*

Proof. We know from Proposition 5.5 that the set of vectors $\mu \in \Pi_{\geq 0}$ that maximize the reward is a non-empty face F of $\Pi_{\geq 0}$. If $F = \Pi_{\geq 0}$, we are obviously in case (i) (i.e., all $\mu \in \Pi_{\geq 0}$ are optimal). One can verify that this corresponds to situations where the vector w

is orthogonal to $\ker(A)$ (for example, if all coordinates of w are identical⁷). Otherwise, because of the lattice structure of the faces of a polytope, F is included into a facet of $\Pi_{\geq 0}$, which means that there is at least one edge $k \in E$ such that the k -th coordinate of all vectors in F is zero in the edge basis. In particular, since all coordinates of the matching rate vector under a stable greedy policy are positive in the edge basis, no greedy policy is optimal. ◀

Proposition 6.2 demonstrates that stable greedy policies can only achieve vectors of matching rates with positive coordinates in the edge basis, i.e., $\Pi_G \subseteq \Pi_{>0}$. In the remainder of this section, we show that whether $\Pi_G \subsetneq \Pi_{>0}$ or $\Pi_G = \Pi_{>0}$ depends on the matching problem (G, λ) under consideration. Specifically, we investigate two examples in Sections 6.1.1 and 6.1.2 where Π_G is a strict subset of $\Pi_{>0}$. In contrast, in Section 6.1.3, we present a well-chosen example where $\Pi_G = \Pi_{>0}$, illustrating a case where greedy policies cover all positive solutions of (CE).

6.1.1 Complete graph

We first consider a matching problem (K_n, λ) , where K_n is the complete graph with $n \geq 3$ nodes. According to Proposition 3.7, this matching problem is stabilizable if and only if $\lambda_i < \frac{1}{2} \sum_{j \in V} \lambda_j$ for each $i \in V$. Propositions 6.4 and 6.5 below state that greedy policies do not allow any degree of freedom, in the sense that the queue-size process and the vector of matching rates are the same under all stable greedy policies. Indeed, the only independent sets of a complete graph are the singletons, hence by (7) the state space of the queue-size process under greedy policies is

$$\mathcal{Q}_G(K_n) = \{0\} \cup \left(\bigcup_{i \in V} \{\ell \mathbb{1}_i : \ell \in \mathbb{N}_{>0}\} \right), \quad (25)$$

where $\mathbb{1}_i$ is the n -dimensional vector with one in coordinate i and zero elsewhere, for each $i \in V$. This observation implies that greediness restricts the freedom in matching decisions: all unmatched items belong to the same class, and an incoming item must be matched with one of them if its class differs. This observation is formalized in Proposition 6.4. Proposition 6.5 further builds on it to conclude that the set Π_G collapses to a single point and directly express the unique vector of matching rates achieved by (stable) greedy policies.

In contrast, the polytope $\Pi_{\geq 0}$ exhibits a dimension $d = m - n = \frac{n(n-3)}{2}$, where $n = |V|$ is the number of classes and $m = |E| = \binom{n}{2} = \frac{n(n-1)}{2}$ is the number of possible matches. In particular, for $n \geq 4$, the dimension $d = m - n$ of Π_G is at least 2. Therefore, limiting ourselves to greedy policies dramatically narrows the range of achievable solutions on complete graphs larger than a triangle.

► **Proposition 6.4** (Equivalence of greedy policies).

Consider the complete graph K_n with $n \geq 3$ nodes.

- (i) There is a unique size-based greedy policy that is adapted to the compatibility graph K_n . This policy, called the natural greedy policy and denoted by $\Phi_G(K_n)$, is the deterministic

⁷ Assume that all coordinates of w are identical. Then a vector $y \in \mathbb{R}^m$ is orthogonal to w if and only if $\sum_{k \in E} y_k = 0$. For each vector $y \in \ker(A)$, we have $Ay = 0$, so that in particular $\sum_{k \in E} (Ay)_k = 0$. But by definition of the matrix A , we also know that $\sum_{k \in E} (Ay)_k = 2 \sum_{k \in E} y_k$.

size-based policy defined on $\mathcal{Q}_{\mathcal{G}}(K_n) \times V$ by

$$\Phi_{\mathcal{G}}(K_n)(q, i) = \begin{cases} j & \text{if } q_j \geq 1 \text{ with } j \neq i, \\ \perp & \text{otherwise.} \end{cases} \quad (26)$$

- (ii) Consider a greedy policy Φ adapted to the graph K_n , and let $(\mathcal{S}, |\cdot|)$ denote its state space. The policy Φ makes the same decisions as the natural greedy policy $\Phi_{\mathcal{G}}$ in the sense that

$$\sum_{s' \in \mathcal{S}} \Phi(s, i, \Phi_{\mathcal{G}}(K_n)(|s|, i), s') = 1, \quad s \in \mathcal{S}, \quad i \in V.$$

In other words, all greedy policies are equivalent in the sense that they can all be reduced to $\Phi_{\mathcal{G}}(K_n)$ (see “Equivalent policies” in Section 2.3.4).

- (iii) For each greedy policy Φ adapted to the graph K_n and each sequence I of item classes, we have $(Q_{\mathcal{G}})_t = Q_t$ for each $t \in \mathbb{N}$, where $Q_{\mathcal{G}}$ and Q are the queue-size processes of the models $(K_n, I, \Phi_{\mathcal{G}}(K_n))$ and (K_n, I, Φ) , respectively.

Proof. See Appendix C.2. ◀

► **Proposition 6.5** (Matching rates under greedy policies).

Consider a stabilizable matching problem (K_n, λ) , where K_n is the complete graph with $n \geq 3$ nodes, and let Φ denote a greedy policy adapted to K_n .

- (i) The model (K_n, λ, Φ) is stable.
(ii) The matching rate vector $\mu_{\mathcal{G}}$ in this model satisfies

$$(\mu_{\mathcal{G}})_k = \lambda_i p_j + \lambda_j p_i, \quad e_k = \{i, j\} \in E, \quad (27)$$

where, for each $i \in V$, p_i is the stationary probability that queue i is non-empty, given by

$$p_i = \frac{\lambda_i}{(\sum_{j \neq i} \lambda_j) - \lambda_i} p_{\emptyset}, \quad (28)$$

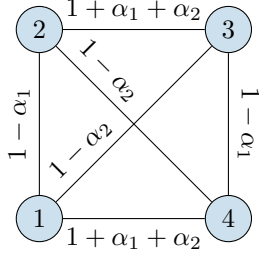
and p_{\emptyset} is the stationary probability that the system is empty, given by

$$p_{\emptyset} = \left(1 + \sum_{i \in V} \frac{\lambda_i}{(\sum_{j \neq i} \lambda_j) - \lambda_i} \right)^{-1}. \quad (29)$$

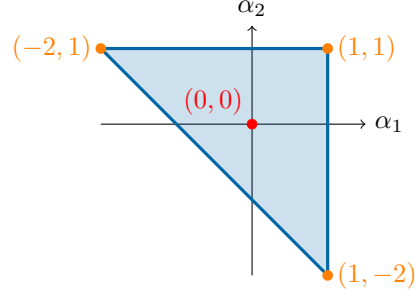
In particular, we have $\Pi_{\mathcal{G}} = \{\mu_{\mathcal{G}}\} \subsetneq \Pi_{>0}$ whenever $n \geq 4$.

Proof. See Appendix C.2. ◀

Figure 17 illustrates this result on a complete graph K_4 in which all arrival rates are equal to 3. The particular solution μ° and the basis $\{b_1, b_2\}$ of $\ker(A)$ that we use are shown on Figure 17a. In the corresponding kernel basis, the polytope $\Pi_{\geq 0}$ is defined by the inequalities $\alpha_1 \leq 1$, $\alpha_2 \leq 1$, and $\alpha_1 + \alpha_2 \geq -1$, that is, it is the triangle with vertices $(-2, 1)$, $(1, -2)$, and $(1, 1)$. Yet, Proposition 6.5 shows that only the solution $\alpha = (0, 0)$ can be achieved by a greedy policy.



(a) Generic solution of (CE).

(b) The polytope $\Pi_{\geq 0}$ in kernel basis. The set Π_G is reduced to the origin vector $\alpha = (0, 0)$.

■ **Figure 17** Matching problem (K_4, λ) with $\lambda = (3, 3, 3, 3)$.

6.1.2 Diamond graph

We now consider the diamond (double fan) graph introduced in Example 3.13, for which we will show that the set Π_G of matching rate vectors achieved by stable greedy policies is again a strict subset of the set $\Pi_{>0}$ of positive solution of (CE), without being restricted to a singleton as in Section 6.1.1 either. More specifically, we consider the diamond matching problem (D, λ) defined as follows:

$$\begin{aligned}
 D &= (V, E) \text{ with } V = \{1, 2, 3, 4\} \text{ and } E = \{\{1, 2\}, \{1, 3\}, \{2, 3\}, \{2, 4\}, \{3, 4\}\}, \\
 \lambda &= (\lambda_1, \bar{\lambda}_2 + \beta, \bar{\lambda}_3 + \beta, \lambda_4) \\
 &\quad \text{with } \lambda_1 > 0, \bar{\lambda}_2 > 0, \bar{\lambda}_3 > 0, \lambda_4 > 0, \beta > 0, \lambda_1 + \lambda_4 = \bar{\lambda}_2 + \bar{\lambda}_3 = \frac{1}{2}, \\
 \mathcal{B} &= \{(1, -1, 0, -1, 1)\}, \\
 \mu^\circ &= (2\lambda_1\bar{\lambda}_2, 2\lambda_1\bar{\lambda}_3, \beta, 2\bar{\lambda}_2\lambda_4, 2\bar{\lambda}_3\lambda_4), \\
 \mu &= (2\lambda_1\bar{\lambda}_2 + \alpha, 2\lambda_1\bar{\lambda}_3 - \alpha, \beta, 2\bar{\lambda}_2\lambda_4 - \alpha, 2\bar{\lambda}_3\lambda_4 + \alpha), \quad \alpha \in \mathbb{R}.
 \end{aligned} \tag{30}$$

Parameterizing λ by β ($= \mu_{2,4}$ for each $\mu \in \Pi$) is a notational convenience that does not lead to any loss of generality: as already observed in Example 3.13, the stabilizability condition (ii) writes $\bar{\lambda}_2 > 0$, $\bar{\lambda}_3 > 0$, and $\beta > 0$. We leave it to the reader to verify that our choices for \mathcal{B} and μ° are correct. Equation (30) implies in particular that the sets $\Pi_{\geq 0}$, $\Pi_{>0}$, $\Pi_{\mathcal{P}}$, and Π_G are real intervals. We will now prove that the interval Π_G , without being reduced to a singleton as in a complete graph, is still a strict subset of the interval $\Pi_{>0}$.

Equivalence of stable greedy policies

Following the same approach as in Section 6.1.1, we first identify a common behavior shared by all greedy policies (Proposition 6.6), and then we exploit this behavior to characterize the matching rates achievable by stable greedy policies (Propositions 6.7 and 6.8). Since $\{1, 4\}$ is the only independent set of the diamond graph D that is not a singleton, the possible queue states under any greedy policy can be partitioned as follows: either all queues are empty, or exactly one class has a non-empty queue, or (only) classes 1 and 4 have non-empty queues. In other words, the state space of the queue-size process under greedy policies adapted to the graph D is

$$\mathcal{Q}_G(D) = \{0\} \cup \left(\bigcup_{i \in V} \{\ell \mathbf{1}_i, \ell \in \mathbb{N}_{>0}\} \right) \cup \{\ell_1 \mathbf{1}_1 + \ell_4 \mathbf{1}_4 : \ell_1, \ell_4 \in \mathbb{N}_{>0}\}. \tag{31}$$

Therefore, greediness entirely determines the decisions made by greedy policies, except if an item of class 2 or 3 enters while there are unmatched items of classes 1 and 4. Using the fact that classes 2 and 3 are both compatible with classes 1 and 4, the next proposition shows that all greedy policies adapted to the diamond graph make the same matching decisions as the natural greedy policy adapted to the complete graph K_3 obtained by “merging” classes 1 and 4 in the diamond graph.

► **Proposition 6.6** (Equivalence of greedy policies).

Given the diamond graph D , we introduce the following notation:

■ *Queue-size projection:* For each $q = (q_1, q_2, q_3, q_4) \in \mathcal{Q}_G(D)$, we let $\langle q \rangle = (q_1 + q_4, q_2, q_3)$.

■ *Class projection:* We let $\langle i \rangle = i$ for each $i \in \{1, 2, 3, \perp\}$ and $\langle 4 \rangle = 1$.

Every greedy policy Φ adapted to the diamond graph D satisfies the following properties:

(i) If Φ is a deterministic size-based policy, then

$$\langle \Phi(q, i) \rangle = \Phi_G(K_3)(\langle q \rangle, \langle i \rangle), \quad (q, i) \in \mathcal{Q}_G(D) \times V.$$

where $\Phi_G(K_3)$ is the natural greedy policy (26) adapted to the complete graph K_3 .

(ii) In general, if Φ is a policy with state space $(\mathcal{S}, |\cdot|)$, then

$$\sum_{s' \in \mathcal{S}} \sum_{\substack{j \in V \cup \{\perp\}: \\ \langle j \rangle = \Phi_G(K_3)(\langle |s| \rangle, \langle i \rangle)}} \Phi(s, i, j, s') = 1, \quad s \in \mathcal{S}, \quad i \in V.$$

(iii) For each sequence I of item classes, we have $\langle Q_t \rangle = (Q_G)_t$ for each $t \in \mathbb{N}$, where Q and Q_G are the queue-size processes of the models (D, I, Φ) and $(K_3, \langle I \rangle, \Phi_G)$, respectively, with $\langle I \rangle = (\langle I_t \rangle, t \in \mathbb{N})$.

Proof. See Appendix C.3. ◀

Loosely speaking, the main take-away of Proposition 6.6(iii) is that, in the matching model (D, I, Φ) , the process $((Q_{t,1} + Q_{t,4}, Q_{t,2}, Q_{t,3}), t \in \mathbb{N})$ does not depend on the specific greedy policy Φ that is applied, and it is in fact equal to the queue size process built by applying the natural greedy policy $\Phi_G(K_3)$ in the complete graph K_3 obtained by merging classes 1 and 4 in the diamond graph D .

Matching rates under stable greedy policies

Propositions 6.7 and 6.8 use the equivalence result of Proposition 6.6 to characterize the matching rate vector under stable greedy policies. These propositions illustrate the complementarity of edge and kernel coordinates: the results of Proposition 6.7 are easier to state using edge coordinates, while those of Proposition 6.8 are easier to state using kernel coordinates. The lower bound (32) on the matching rates in the diamond matching problem are obtained by following a similar approach as in Proposition 6.5(ii).

► **Proposition 6.7** (Matching rates under greedy policies – Edge coordinates).

Consider the matching model (D, λ, Φ) , where (D, λ) is the diamond problem (30) and Φ is a greedy policy adapted to the graph D .

(i) This matching model is stable.

(ii) The matching rate vector $\mu = \mu(D, \lambda, \Phi)$ satisfies $\mu_{2,3} = \beta = \frac{1}{2}(\lambda_2 + \lambda_3 - \lambda_1 - \lambda_4)$, and

$$\begin{aligned} \mu_{1,2} &\geq \lambda_1 p_2 + \lambda_2 p_1, & \mu_{1,3} &\geq \lambda_1 p_3 + \lambda_3 p_1, \\ \mu_{2,4} &\geq \lambda_2 p_4 + \lambda_4 p_2, & \mu_{3,4} &\geq \lambda_3 p_4 + \lambda_4 p_3, \end{aligned} \tag{32}$$

where p_i is the stationary probability that the system contains unmatched items that all belong to class i , for each $i \in V$.

(iii) Let $p_{1,4}$ denote the stationary probability that the system contains unmatched items that belong to class 1 or 4 (or both). We have

$$\begin{aligned} p_2 &= \frac{\lambda_2}{\lambda_1 + \lambda_3 + \lambda_4 - \lambda_2} p_\emptyset, & p_3 &= \frac{\lambda_3}{\lambda_1 + \lambda_2 + \lambda_4 - \lambda_3} p_\emptyset, \\ p_{1,4} &= \frac{\lambda_1 + \lambda_4}{\lambda_2 + \lambda_3 - \lambda_1 - \lambda_4} p_\emptyset, \end{aligned} \quad (33)$$

where p_\emptyset is the stationary probability that the system is empty, whose value follows from the normalization equation $p_\emptyset + p_{1,4} + p_2 + p_3 = 1$. We also have

$$p_1 > \frac{\lambda_1}{\lambda_2 + \lambda_3 + \lambda_4} p_\emptyset, \quad p_4 > \frac{\lambda_4}{\lambda_1 + \lambda_2 + \lambda_3} p_\emptyset. \quad (34)$$

Proof. See Appendix C.3. ◀

Intuitively, the kernel coordinate α given in (30) (also visible in Figure 6) acts like a slider that determines how much edges $\{1, 2\}$ and $\{3, 4\}$ are (dis)avored compared to edges $\{1, 3\}$ and $\{2, 4\}$ in the long run. Remarkably, Proposition 6.8 below shows that the greedy policy that favors edges $\{1, 2\}$ and $\{3, 4\}$ (resp. $\{1, 3\}$ and $\{2, 4\}$) the most *in the long run* is also the policy that favors these two edges the most *in the short run*. This result is proved by a coupling argument. In addition, Proposition 6.8 uses the lower bound (32) in Proposition 6.7 to prove that, in the limit as $\beta \rightarrow +\infty$, all greedy policies yield the same matching rate vector.

► **Proposition 6.8** (Matching rates under greedy policies – Kernel coordinates).

Consider the diamond matching problem (D, λ) described in (30).

(i) The intervals $\Pi_{\geq 0}$, $\Pi_{> 0}$, and Π_G are defined as follows in the kernel basis:

$$\begin{aligned} \Pi_{\geq 0} &= [-2 \min(\lambda_1 \bar{\lambda}_2, \bar{\lambda}_3 \lambda_4), 2 \min(\lambda_1 \bar{\lambda}_3, \bar{\lambda}_2 \lambda_4)], \\ \Pi_{> 0} &= (-2 \min(\lambda_1 \bar{\lambda}_2, \bar{\lambda}_3 \lambda_4), 2 \min(\lambda_1 \bar{\lambda}_3, \bar{\lambda}_2 \lambda_4)), \\ \Pi_G &= [\alpha_-, \alpha_+], \end{aligned}$$

with $-2 \min(\lambda_1 \bar{\lambda}_2, \bar{\lambda}_3 \lambda_4) < \alpha_- \leq \alpha_+ < 2 \min(\lambda_1 \bar{\lambda}_3, \bar{\lambda}_2 \lambda_4)$. Hence, $\Pi_G \subsetneq \Pi_{> 0} \subsetneq \Pi_{\geq 0}$.

(ii) The coordinates α_+ and α_- satisfy the following properties:

- a. $\alpha_+ = \alpha(\Phi_+)$, where Φ_+ is the edge-priority policy adapted to the graph D whereby edges $\{1, 2\}$ and $\{3, 4\}$ have the highest priority.
- b. $\alpha_- = \alpha(\Phi_-)$, where Φ_- is the edges-priority policy adapted to the graph D whereby edges $\{1, 3\}$ and $\{2, 4\}$ have the highest priority.
- c. If $\beta \rightarrow +\infty$ while $\lambda_1, \bar{\lambda}_2, \bar{\lambda}_3$, and λ_4 remain fixed, we have $\alpha_+ \rightarrow 0$ and $\alpha_- \rightarrow 0$.

Proof. See Appendix C.3. ◀

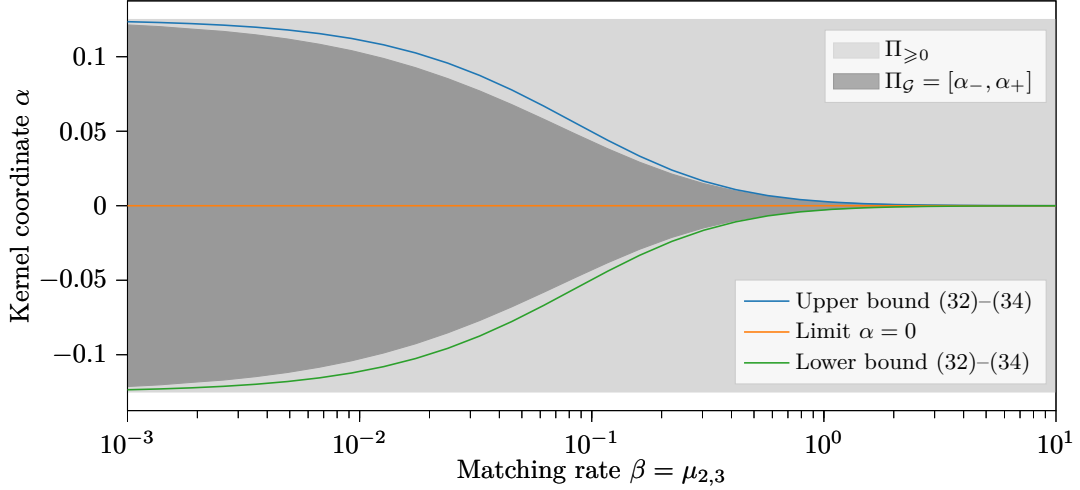
Taken together, Statements (i) and (iic) in Proposition 6.8 show that, as $\beta \rightarrow +\infty$, the interval Π_G becomes reduced to a single point $\alpha = 0$, meaning that all greedy policies yield the same vector of matching rates, with edge coordinates $\mu = (2\lambda_1 \bar{\lambda}_2, 2\lambda_1 \bar{\lambda}_3, \beta, 2\bar{\lambda}_2 \lambda_4, 2\bar{\lambda}_3 \lambda_4)$. The rationale behind this result is the following. In the regime where $\beta = \mu_{2,3} \rightarrow +\infty$, we see in (30) that the arrival rates of classes 2 and 3 become large compared to those of classes 1 and 4. As a result, items of classes 1 and 4 are matched (almost) always immediately, and unmatched items belong to either class 2 or class 3 (but not both at the same time). In the proof of Proposition 6.8(iic), this intuition is formalized by taking the limit of (33) as $\beta \rightarrow +\infty$, which yields

$$p_\emptyset \rightarrow 0, \quad p_2 \rightarrow 2\bar{\lambda}_2, \quad p_3 \rightarrow 2\bar{\lambda}_3, \quad p_{1,4} \rightarrow 0.$$

In this regime, the greediness of the policy allows no degree of flexibility in choosing the class of the item to which an incoming item is matched, so the matching rates are unique.

► **Remark 6.9.** Some quantities in Proposition 6.8 are functions of the matching rate $\beta = \mu_{2,3}$, but this dependency is kept implicit to simplify notation. In particular, the intervals $\Pi_{\geq 0}$ and $\Pi_{> 0}$ do *not* depend on β , but the interval $\Pi_{\mathcal{G}}$ and the coordinates α_+ and α_- do.

Numerical results



■ **Figure 18** Evolution of $\Pi_{\mathcal{G}}$ as a function of β in the diamond problem (30) with $\lambda_1 = \bar{\lambda}_2 = \bar{\lambda}_3 = \lambda_4 = \frac{1}{4}$. For each β , $\Pi_{\mathcal{G}}$ is estimated by simulating the model (D, λ, Φ_+) and leveraging the symmetry of the problem. $\Pi_{\geq 0}$ and other bounds are displayed for comparison. All results are expressed in kernel coordinates.

To illustrate Proposition 6.8, Figure 18 shows a symmetric example with $\lambda_1 = \bar{\lambda}_2 = \bar{\lambda}_3 = \lambda_4 = \frac{1}{4}$. The figure compares $\Pi_{\geq 0}$ and $\Pi_{\mathcal{G}}$ with the bounds (32)–(34) (converted in the kernel coordinates) and the limit $\alpha = 0$. Each point forming the shape of $\Pi_{\mathcal{G}} = [\alpha_-, \alpha_+]$ is obtained by running a simulation consisting of 10^{10} steps (as specified in Section 2.5). As announced by Proposition 6.8, $\Pi_{\mathcal{G}}$ becomes reduced to a single point $\alpha = 0$ when $\beta \rightarrow +\infty$. We also notice that the bounds provided by (32)–(34) are not tight when β is small (in the sense that the difference between α_+ and the upper-bound is no longer negligible compared to the difference between α_+ and the boundary $2 \min(\lambda_1 \bar{\lambda}_3, \bar{\lambda}_2, \lambda_4)$ of $\Pi_{\geq 0}$), while at the same time α_+ and α_- become arbitrarily close to the borders of $\Pi_{\geq 0}$ when β tends to zero (which does not contradict the fact that $\Pi_{\mathcal{G}}$ is a strict subset of $\Pi_{\geq 0}$ as long as $\beta > 0$). The gap between $\Pi_{\mathcal{G}}$ and the bounds (32)–(34) comes from the fact that, to obtain the lower bounds for p_1 and p_4 in (34), we neglected the case where there are both class-1 and class-4 unmatched items, and this case is not negligible when β is small.

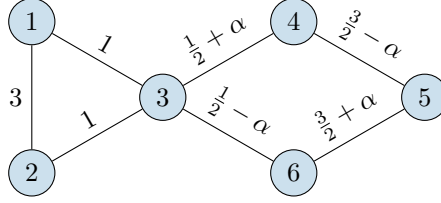
6.1.3 Fish graph

Sections 6.1.1 and 6.1.2 exhibited two stabilizable matching problems in which the matching rates achievable by stable greedy policies were bounded away from the faces of the polytope $\Pi_{\geq 0}$. These two examples may suggest that, in most stabilisable matching problems, the restrictions imposed by greediness are so strong that we will also obtain $\Pi_{\mathcal{G}} \subsetneq \Pi_{> 0}$. We

now consider another matching problem, called the Fish problem, where we conjecture on the contrary that $\Pi_G = \Pi_{>0}$. The Fish matching problem is defined as follows:

$$\begin{aligned}
V &= \{1, 2, 3, 4, 5, 6\}, \\
E &= \{\{1, 2\}, \{1, 3\}, \{2, 3\}, \{3, 4\}, \{4, 5\}, \{5, 6\}\}, \\
\lambda &= (4, 4, 3, 2, 3, 2), \\
\mathcal{B} &= \{(0, 0, 0, 1, -1, -1, 1)\}, \\
\mu^\circ &= (3, 1, 1, \frac{1}{2}, \frac{1}{2}, \frac{3}{2}, \frac{3}{2}), \\
\mu &= (3, 1, 1, \frac{1}{2} + \alpha, \frac{1}{2} - \alpha, \frac{3}{2} - \alpha, \frac{3}{2} + \alpha), \quad \alpha \in \mathbb{R}.
\end{aligned} \tag{35}$$

The general solution μ given in (35) is shown in Figure 19. We can verify by a direct inspection that this matching problem is stabilizable and that $\Pi_{\geq 0} = [-\frac{1}{2}, \frac{1}{2}]$ in kernel coordinates. As in Section 6.1.2, the kernel coordinate α acts like a slider that is positive (resp. negative) if matches along edges $\{3, 4\}$ and $\{5, 6\}$ are more (resp. less) frequent than matches along edges $\{3, 6\}$ and $\{4, 5\}$. Conjecture 6.10 below is supported by an intuition of proof and by extensive simulations discussed later in this section.



■ **Figure 19** Generic solution of (CE) in the Fish matching problem of (35) with $\lambda = (4, 4, 3, 2, 3, 2)$.

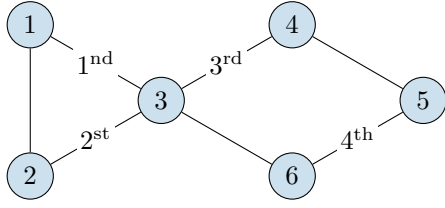
► **Conjecture 6.10.** *In the Fish matching problem of (35), we have $\Pi_G =]-\frac{1}{2}, \frac{1}{2}[= \Pi_{>0}$.*

Intuition of the proof

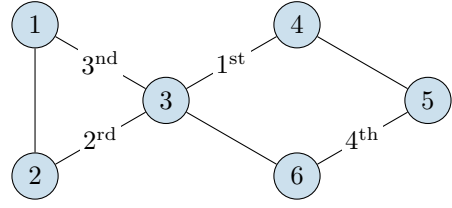
We build two families of stable greedy policies, denoted by $(\Phi_k^+)_{k \in \mathbb{N}}$ and $(\Phi_k^-)_{k \in \mathbb{N}}$, such that $\lim_{k \rightarrow +\infty} \alpha(\Phi_k^+) = \frac{1}{2}$ and $\lim_{k \rightarrow +\infty} \alpha(\Phi_k^-) = -\frac{1}{2}$. The conclusion then follows from the convexity of the set Π_G (Proposition 6.1). We focus on the family $(\Phi_k^+)_{k \in \mathbb{N}}$, as the family $(\Phi_k^-)_{k \in \mathbb{N}}$ is symmetrical (in the sense that it suffices to exchange classes 4 and 6).

The family $(\Phi_k^+)_{k \in \mathbb{N}}$ is defined as follows. Let Φ_∞^+ denote the edge-priority policy where edges have the following descending priority order: $\{1, 3\}, \{2, 3\}, \{3, 4\}, \{5, 6\}$, followed by all other edges in an arbitrary order. Let Φ_0^+ denote the edge-priority policy where edges have the following descending priority order: $\{3, 4\}, \{2, 3\}, \{1, 3\}, \{5, 6\}$, followed by all other edges in an arbitrary order. The important point is that both policies prioritize edges $\{3, 4\}$ and $\{5, 6\}$ over edges $\{3, 6\}$ and $\{4, 5\}$ (with the hope that this lead to a high α), but that Φ_∞^+ gives higher priority to the “tail” while Φ_0^+ gives higher priority to the “trunk”. Now, for each $k \in \mathbb{N}$, Φ_k^+ is the deterministic size-based policy that follows Φ_∞^+ when the queue size of class 4 is at most $k - 1$ and Φ_0^+ when the queue size of class 4 is at least k (that is, $\Phi_k^+(q, i) = \Phi_\infty^+(q, i)$ if $q_4 \leq k - 1$ and $\Phi_k^+(q, i) = \Phi_0^+(q, i)$ if $q_4 \geq k$).

The rationale behind this definition is as follows. According to (35), α is maximal (i.e., equal to $\frac{1}{2}$) when $\mu_{3,4}$ is equal to 1 and $\mu_{3,6}$ is equal to 0. If we allow *non-greedy* policies, $\alpha = \frac{1}{2}$ can be achieved by merely applying the edge-filtering variant of match-the-longest (or any maximally-stable policy) on the bijective subgraph of G obtained by eliminating edge $\{3, 6\}$, that is, by *never* performing a match between classes 3 and 6. As we will see, the



(a) Policy Φ_∞^+ , followed by policy Φ_k^+ when the size of queue 4 is at most $k - 1$



(b) Policy Φ_0^+ , followed by policy Φ_k^+ when the size of queue 4 is at least k

■ **Figure 20** Edge-priority (greedy) policies in the Fish matching problem of (35).

family $(\Phi_k^+)_{k \in \mathbb{N}}$ of (stable greedy) policies emulates this (non-greedy) edge-filtering policy by favoring edge $\{3, 4\}$ over edge $\{3, 6\}$ while making the probability that $q_4 = 0$ arbitrary small as k increases. Roughly speaking, mixing the policies Φ_∞^+ and Φ_0^+ allows us to control q_4 : Φ_∞^+ is not stable and makes q_4 go to infinity, while Φ_0^+ is stable and reduces q_4 when it is large. All in all, Φ_k^+ keeps the value of q_4 around k , so that the probability that $q_4 = 0$ is low when k is large. In the limit, q_4 is always positive, so that the class-3 items that are not matched with classes 1 or 2 are drained by class 4 (while class 6 is matched only with class 5), so that $\mu_{3,4}(\Phi_k^+)$ tends to 1 and $\mu_{3,6}(\Phi_k^+)$ tends to 0 as $k \rightarrow \infty$.

We expect that a rigorous proof that $\mu_{3,6}(\Phi_k^+)$ tends to 0 as $k \rightarrow +\infty$ will involve the following steps, some of which may require further investigation or analysis:

1. For each $k \in \mathbb{N}$, the matching model (G, λ, Φ_k^+) is stable. We believe that stability can be proved by applying a fluid-limit argument that generalizes the framework of [18]⁸. Therefore, for each $k \in \mathbb{N}$, we can consider a random vector $Q_k = (Q_{k,1}, Q_{k,2}, \dots, Q_{k,n})$ distributed like the vector of queue sizes in the matching model (G, λ, Φ_k^+) in stationary regime, and we know that the vector $\mu(\Phi_k^+)$ is well-defined.
2. For each $k \in \mathbb{N}$, by definition of the policy Φ_k^+ , we have

$$\begin{aligned} \mu(\Phi_k^+)_{3,6} &= \lambda_3 P(Q_{k,4} = 0, Q_{k,6} > 0) + \lambda_6 P(Q_{k,3} > 0, Q_{k,5} = 0), \\ &\leq (\lambda_3 + \lambda_6) P(Q_{k,4} = 0), \end{aligned}$$

where the inequality arises because the events $(Q_{k,4} = 0, Q_{k,6} > 0)$ and $(Q_{k,3} > 0, Q_{k,5} = 0)$ are both included into the event $(Q_{k,4} = 0)$ (in the latter case, this is because the greediness of the policy prevents the system from containing unmatched items of classes 3 and 4 at the same time). Consequently, to prove that $\mu(\Phi_k^+)_{3,6} \rightarrow 0$, it suffices to prove that $P(Q_{k,4} = 0) \rightarrow 0$ as $k \rightarrow +\infty$. This argument rigorously supports the intuition that, if class 4 becomes “unstable”, then this class “drains” all class-3 items that are not matched with classes 1 and 2.

3. To prove that $P(Q_{k,4} = 0) \rightarrow 0$ as $k \rightarrow +\infty$, we believe that we can reason as follows:
 - We can prove that, for each $k \in \mathbb{N}$, the conditional average rate at which class 3 is matched with either class 4 or class 6, *given that* $Q_{k,4} \leq k - 1$, is upperbounded by $\lambda_3 \left(1 + \frac{\lambda_1}{\lambda_2 + \lambda_3 - \lambda_1} + \frac{\lambda_2}{\lambda_1 + \lambda_3 - \lambda_2}\right)^{-1} = \frac{9}{11}$. This upperbound follows by coupling the fish problem with a “worst-case” matching problem in which there is an infinite backlog of items of classes 4 and 6 to be matched with class 3. Roughly speaking, this worst-case matching problem can be obtained by taking $\lambda_4 = \lambda_6 = +\infty$ in the fish problem; then, focusing on nodes 1, 2, and 3, the system behaves like a paw graph with an unstable pendant node obtained by merging classes 4 and 6.

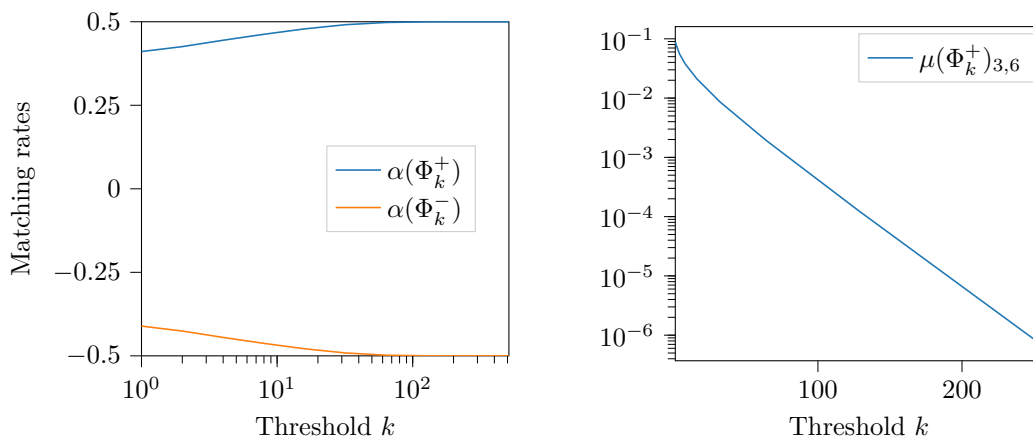
⁸ Formalizing this generalization is outside the scope of this paper and could be the topic of a future work.

- Using this upperbound, we prove that at least one class among classes 4 and 6 becomes unstable, in the sense that $\lim_{k \rightarrow +\infty} P(Q_{k,4} = 0) = 0$ or $\lim_{k \rightarrow +\infty} P(Q_{k,6} = 0) = 0$. Since we can verify otherwise that class 6 does not become unstable, it follows that $\lim_{k \rightarrow +\infty} P(Q_{k,4} = 0) = 0$.

The complete proof is left as an open question for future works. ◀

Numerical results

Figure 21 shows the values of the matching rates of Φ_k^+ and Φ_k^- obtained by simulation, using the setting described in Section 2.5, as functions of the threshold $k \in \mathbb{N}$. The main take-away is that $\alpha(\Phi_k^+)$ and $\alpha(\Phi_k^-)$ seem to converge to $\frac{1}{2}$ and $-\frac{1}{2}$, respectively, thus supporting Conjecture 6.10. Figure 21b, which displays $\mu(\Phi_k^+)_{3,6} = \frac{1}{2} - \alpha(\Phi_k^+)$, suggests that the convergence is linear (with a logarithmic scale on the y-axis), meaning that the rate $\mu(\Phi_k^+)_{3,6}$ decreases like γ^k for some $0 < \gamma < 1$. Each curve is obtained by running a simulation consisting of 10^{10} steps (as specified in Section 2.5).



(a) Evolution of $\alpha(\Phi_k^+)$ and $\alpha(\Phi_k^-)$.

(b) Evolution of $\mu(\Phi_k^+)_{3,6} = \frac{1}{2} - \alpha(\Phi_k^+)$.

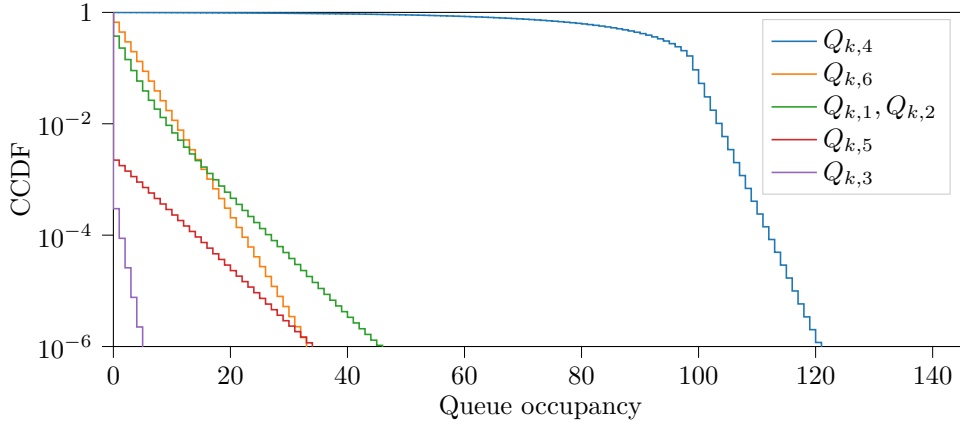
■ **Figure 21** Evolution of the matching rates under Φ_k^+ and Φ_k^- as functions of the threshold k , as defined in the intuition of proof of Conjecture 6.10. Figure 21a is in kernel coordinates, while Figure 21b zooms on $\mu(\Phi_k^+)_{3,6} = \frac{1}{2} - \alpha(\Phi_k^+)$.

Figure 22 shows the empirical distribution of the queue sizes under Φ_{100}^+ . One can observe in particular that:

- The distribution of $Q_{k,4}$ is centered around k . This illustrates the design of Φ_k^+ , which aims at “stabilizing” the number of class-4 items in the system around the value k . As a consequence, the probability that there is no class-4 item is low.
- Typically, the system does not contain items of classes 3 and 5, which is a consequence of the low probability of having no class-4 items. The presence of class-4 items precludes the presence of items of classes 3 and 5 because the policy is greedy.
- Because of the symmetric roles of classes 1 and 2, the distributions of $Q_{k,1}$ and $Q_{k,2}$ are identical, even though the presence of class-1 and class-2 items is mutually exclusive.

6.2 Arbitrary policies

We now focus on $\Pi_{\mathcal{P}} \subseteq \Pi_{\geq 0}$, the set of matching rate vectors that can be achieved by an arbitrary stable policy. A simple necessary and sufficient condition for a vertex of $\Pi_{\geq 0}$ to



■ **Figure 22** Empirical complementary cumulative distribution function (CCDF) of the queue sizes under policy Φ_k^+ for $k = 100$.

belong to $\Pi_{\mathcal{P}}$ was already derived in Proposition 5.7: if μ is a vertex of $\Pi_{\geq 0}$, then $\mu \in \Pi_{\mathcal{P}}$ if and only if the vertex μ is bijective (in the sense that its support graph is bijective). Proposition 6.11 uses this result to characterize which faces of the polytope $\Pi_{\geq 0}$ can be entirely covered by stable policies.

► **Proposition 6.11.** *For each face F of the polytope $\Pi_{\geq 0}$, we have $F \subseteq \Pi_{\mathcal{P}}$ if and only if all vertices of F are bijective. In particular, $\Pi_{\mathcal{P}} = \Pi_{\geq 0}$ if and only if the polytope $\Pi_{\geq 0}$ is simple and the matching problem (G, λ) is essential.*

Proof. The first statement is a consequence of the convexity result of Proposition 6.1 and the characterization of the vertices of the polytope $\Pi_{\geq 0}$ in Proposition 5.7. The second statement is a consequence of Proposition 5.13. ◀

Achieving injective-only vertices

Although Proposition 5.7 shows that injective-only vertices cannot be achieved by stable policies, we now introduce a sequence of stable policies that we conjecture yield a sequence of matching rate vectors that are arbitrarily close a vertex of the polytope $\Pi_{\geq 0}$, even if this vertex is injective-only. This conjecture is supported by numerical results shown later in this section.

► **Conjecture 6.12.** *Consider a vertex μ of $\Pi_{\geq 0}$ and let $\underline{G} = (V, \underline{E})$ denote its support graph. For each $k \in \mathbb{N}$, consider the following semi-filtering policy, denoted by $\Phi_k(\mu)$:*

- *If the size of the longest queue is less than k , apply the filtering match-the-longest policy adapted to \underline{G} with filter \underline{E} ;*
- *Otherwise, apply the greedy match-the-longest policy adapted to \underline{G} .*

$\Phi_k(\mu)$ is stable for each $k \in \mathbb{N}$ and $\lim_{k \rightarrow \infty} \mu(\Phi_k(\mu)) = \mu$.

Intuition of the proof. In essence, the approach is similar to the one that supports Conjecture 6.10: we take a possibly unstable policy that achieves the desired goal (a matching rate vector equal to μ), and we make it stable by reverting to a stable policy when the queue sizes become too large. If the threshold is large enough, most matches will be made under the unstable policy that achieves the desired goal.

We present here a sketch of proof for the case where \underline{G} is a bipartite graph, more precisely a tree (like the support graphs of the vertices in Example 3.13 when $\lambda_1 = \bar{\lambda}_2 = \bar{\lambda}_3 = \lambda_4$).

The tree case is the most difficult to handle. If \underline{G} is a unicyclic graph with an odd cycle, the problem is simplified by the stabilizability of (\underline{G}, λ) . If \underline{G} is made of multiple connected components, the proof needs to compare the model to a virtual model where each component acts independently and accounts for the interactions between the components. For example, when one the queue is longer than k , all components may interact when decisions follow the match-the-longest policy on the graph G .

The stability of the model $(G, \lambda, \Phi_k(\mu))$ for each $k \in \mathbb{N}$ can be proved using the Lyapunov-Foster theorem [8, Theorem 1.1 in Chapter 5], using the fact that, outside a finite set of states (those where all queues are shorter than k), the policy $\Phi_k(\mu)$ behaves like the stable match-the-longest policy on the graph G . Consider a random vector $Q = (Q_1, Q_2, \dots, Q_n)$ distributed like the vector of queue sizes in the matching model $(G, \lambda, \Phi_k(\mu))$ in stationary regime, and let $Q_+ = \sum_{i \in V_+} Q_i$ and $Q_- = \sum_{i \in V_-} Q_i$ denote the total queue sizes in the parts V_+ and V_- of the (bipartite) graph \underline{G} . Also let $L_+ = \max(Q_i, i \in V_+)$, $L_- = \max(Q_i, i \in V_-)$, and $L = L_+$ if $L_+ \geq L_-$ and $L = -L_-$ if $L_+ < L_-$, so that $|L|$ is the size of the longest queue and the sign of L indicates the part (V_+ or V_-) to which this queue belongs. The random variable L will play a crucial role in the proof. Lastly, we let $p_\ell = \mathbb{P}(L = \ell)$, $\ell \in \{\dots, -2, -1, 0, 1, 2, \dots\}$, denote its distribution. The dependency of these random variables on k (via the policy $\Phi_k(\mu)$) is left implicate to simplify notation.

To show that $\lim_{k \rightarrow \infty} \mu(\Phi_k(\mu)) = \mu$, we need to upperbound the probability that a match is performed along an edge that does not belong to the support of the vertex μ . To do this, we upperbound the probability of applying the greedy match-the-longest policy on the whole graph G , i.e., the probability $\sum_{|\ell| \geq k} p_\ell$ that the size of the longest queue is k or more.

Intuitively, we expect that the stability of the match-the-longest policy implies a uniform drift of the size of the longest queue towards the origin whenever this size is larger than k , provided that k is large enough. Formally, we conjecture that there exist $K \in \mathbb{N}$ and $0 < \rho < 1$ such that, for each $k \geq K$, we have under Φ_k that $p_{\ell+1} \leq \rho p_\ell$ for each $\ell \geq k$ and $p_{\ell-1} \leq \rho p_\ell$ for each $\ell \leq -k$. In particular, $\sum_{\ell \geq k} p_\ell \leq \frac{1}{1-\rho} p_k$ and $\sum_{\ell \leq -k} p_\ell \leq \frac{1}{1-\rho} p_{-k}$. This implies that, if p_k and p_{-k} go to 0 as k goes to infinity, the probability that the filter on \underline{E} is disabled goes to 0 as well.

To control p_k and p_{-k} , we now need to consider the cases where the size of the longest queue is at most k . We argue that, as long as $|L| \leq k$, the dynamics of L is mainly controlled by an unbiased random walk between the two parts of the bipartite graph. More precisely, as long as $|L| \leq k$, the difference $Q_+ - Q_-$ behaves like an unbiased random walk because any arrival of a class in V_+ increases the difference by one, any arrival of a class in V_- decreases the difference by one, and the arrival rates of V_+ and V_- are equal (see Equation (24)). We conjecture that, when $|Q_+ - Q_-|$ is large enough (while still satisfying $|L| \leq k$), a distinctive structure emerges due to the filtered match-the-longest policy: the queues in one part of the graph have approximately equal sizes, while the queues in the other part are empty. In particular, we conjecture that the typical states of the system (when $|L| \leq k$) verify (in some sense) $L \approx \frac{Q_+ - Q_-}{|V_+|}$ if $Q_+ > Q_-$ and $L \approx \frac{Q_+ - Q_-}{|V_-|}$ if $Q_+ < Q_-$. In a true unbiased random walk over a bounded range of integers, all states have equal probabilities. By comparing the evolutions of L and $Q_+ - Q_-$ and taking into account the border cases ($|L| = k$), we conjecture that there exist $K' \in \mathbb{N}$ and $c > 0$ such that, if $k \geq K'$, then $p_\ell \geq c \max(p_k, p_{-k})$ for all $\ell \in \{-k, \dots, -1, 0, 1, \dots, k\}$. In particular, both p_k and p_{-k} should be upperbounded by $(c(2k+1))^{-1}$ when k is large enough.

By combining the two cases $|L| \geq k$ and $|L| \leq k$, we get $\sum_{|\ell| \geq k} p_\ell = O(1/k)$. In other words, the probability that the size of the longest queue is greater than k tends to 0 as k goes to infinity. As matchings outside \underline{E} occur only when the size of the longest queue is greater

than k , we conclude that the matching rate of an edge outside \underline{E} goes to 0, which by continuity of the conservation principle and injectivity of \underline{G} implies that $\lim_{k \rightarrow \infty} \mu(\Phi_k(\mu)) = \mu$. ◀

One drawback of the semi-filtering policies introduced in Conjecture 6.12 is that, as $k \rightarrow +\infty$ and we approach an injective-only vertex, the average size of the longest queue tends to grow. Therefore, there is a trade-off between approaching an injective-only vertex and minimizing the waiting time. This issue is similar to the threshold-based greedy policies introduced in Conjecture 6.10, where the vertices of $\Pi_{\geq 0}$ could be reached asymptotically by increasing the average size of a queue.

► **Remark 6.13.** In [29], a family of policies is introduced to optimize a reward function on edges. The goal is similar to ours, with two main differences: first, the aim is to optimize a reward function on matching rates without providing a description of the limit rate vector, while we show that the limit rate vector is a vertex of $\Pi_{\geq 0}$; second, each policy of the family makes decisions based on a utility function that is a convex combination of the edge-dependent reward and the queue size, and matching decisions are based on virtual queues that can become negative. The policies are indexed by the coefficient β of the convex combination, which determines the trade-off between maximizing the reward and minimizing the queue size, and which is the counterpart of the threshold k . Yet, the authors also notice in their simulations that the queue sizes grow as $1/\beta$, which is the counterpart of our k .

If Conjecture 6.12 holds, which we strongly believe and will support with numerical results, we obtain a simple characterization of the matching rates that can be achieved by a stable policy.

► **Corollary 6.14.** *Any positive solution of the conservation equation (CE-2) can be obtained by a stable policy. In other words,*

$$\Pi_{>0} \subseteq \Pi_{\mathcal{P}} \subseteq \Pi_{\geq 0}.$$

Proof. We establish that $\Pi_{\mathcal{P}} \subseteq \Pi_{\geq 0}$ by recalling that any stable policy yields a matching rate vector that satisfies (CE) and has non-negative coordinates.

We now prove that $\Pi_{>0} \subseteq \Pi_{\mathcal{P}}$. The closure of $\Pi_{\mathcal{P}}$ is convex, as $\Pi_{\mathcal{P}}$ itself is convex. According to Conjecture 6.12, the closure of $\Pi_{\mathcal{P}}$ contains all the vertices of $\Pi_{\geq 0}$, thereby encompassing $\Pi_{\geq 0}$ as well by convexity. Consequently, still by virtue of convexity, $\Pi_{\mathcal{P}}$ contains the interior⁹ of its closure, which includes the interior of $\Pi_{\geq 0}$, denoted as $\Pi_{>0}$. ◀

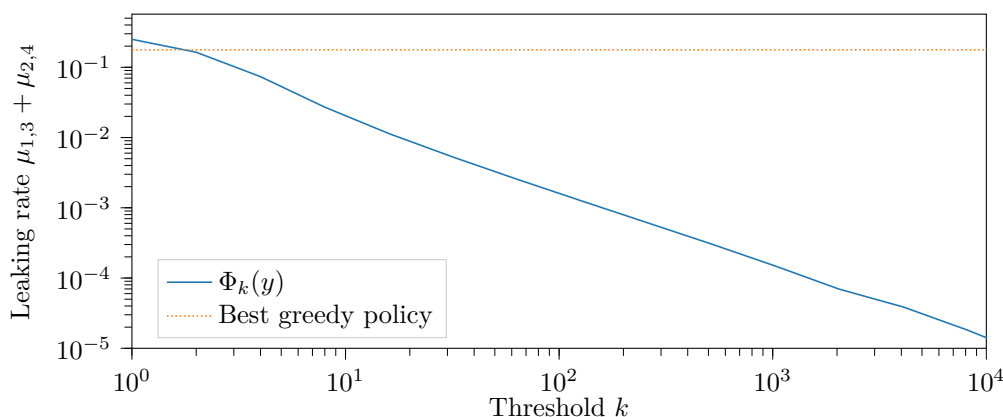
Numerical results

To provide a numerical support for Conjecture 6.12, we evaluated $\Phi_k(\mu)$ for some injective-only vertices. We first considered the the diamond problem (30) with $\lambda_1 = \bar{\lambda}_2 = \bar{\lambda}_3 = \lambda_4 = \frac{1}{4}$ and $\beta = \frac{1}{8}$. The two vertices of $\Pi_{\geq 0}$ have injective-only support graphs that are identical up to permutation of nodes 2 and 3. We choose to focus on the vertex $\mu = (\frac{1}{8}, 0, \frac{1}{8}, \frac{1}{8})$: its support consists of the line $\{1, 2\}, \{2, 3\}, \{3, 4\}$ (a simple example of a tree), so that the edges $\{1, 3\}$ and $\{2, 4\}$ should be avoided. Each curve is obtained by running a simulation consisting of 10^{10} steps (as specified in Section 2.5).

Figure 23 shows the convergence of $\mu(\Phi_k(\mu))$ to μ by measuring as a function of k the *leaking rate* $\mu_{1,3}(\Phi_k(\mu)) + \mu_{2,4}(\Phi_k(\mu))$, i.e., the total matching rates observed along the

⁹ Here the notion of *interior* is relative to the canonic topology of the d -dimensional affine space of the solutions of the conservation equation.

forbidden edges $\{1, 3\}$ and $\{2, 4\}$. The convergence seems to be logarithmic, i.e., like $1/k$. This is in line with the intuition of the proof of Conjecture 6.12, which states that the probability to enable forbidden edges should behave like $O(1/k)$. The leaking rate achievable by Φ_+ , the greedy edge-priority policy defined in Proposition 6.8, is displayed for comparison purposes. Φ_+ minimizes the leaking rate among all greedy policies. As expected, $\Phi_k(\mu)$ yields a lower leaking rate than Φ_+ for large values of k . However, it can be observed that $\Phi_0(\mu)$ exhibits a higher leaking rate. This is actually expected: when $k = 0$, filtering is completely disabled, and thus $\Phi_0(\mu)$ represents the *match-the-longest* (greedy) policy, which does not attempt to reduce matchings along the forbidden edges.



■ **Figure 23** Total matching rate along the forbidden edges under policy $\Phi_k(\mu)$ in the diamond problem of (30), with $\lambda_1 = \bar{\lambda}_2 = \bar{\lambda}_3 = \lambda_4 = \frac{1}{4}$ and $\beta = \frac{1}{8}$. The vertex μ has support $\{1, 2\}, \{2, 3\}, \{3, 4\}$, so that the edges to avoid are $\{1, 3\}$ and $\{2, 4\}$. The performance of the greedy edge-priority policy Φ_+ is displayed for comparison.

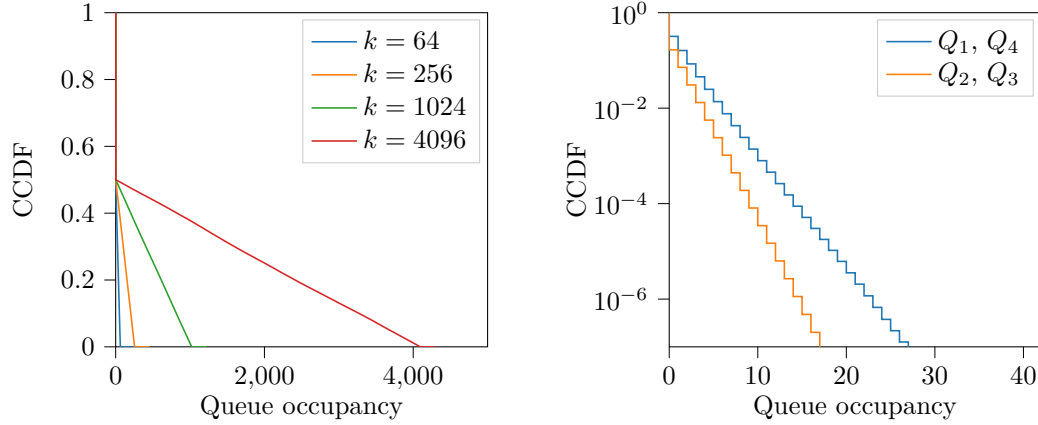
As mentioned previously, approaching an injective-only vertex should cause the queues to become arbitrarily large. To support this assertion, we studied by simulation the empirical distributions of the queue sizes. Although we do not believe that the sizes of all four classes of the model $(D, \lambda, \Phi_k(\mu))$ follow the same distribution, the difference was not visible in our observations because it is only of second order compared to the overall increase of queue size as k increases. Figure 24a presents, for various values of k , the unique observed distribution. For a given class i , the following comments can be made:

- The queue is empty half of time. Referring to the intuition of the proof of Conjecture 6.12, this is related to our observation that, as long as the maximum queue size is below k , there is always one part of the bipartite subgraph whose queues are empty, and the difference $Q_+ - Q_-$ in total queue size between the two parts behaves like an unbiased random walk. This is due to the fact that the total arrival rates within each part are equal to each other. See Example 3.12 and Remark 5.8 for more insights.
- Otherwise, the size of queue i appears to follow a uniform distribution on $\{1, 2, \dots, k\}$.
- The probability that the size of queue i exceeds k seems to be negligible (as observed in Figure 23, we conjecture a decrease in $O(1/k)$).

These observations align with the intuition of proof that supports Conjecture 6.12. In particular, most of the time, the queue sizes are short (i.e., below k) and we apply a filtering policy, so that the evolution of the queue sizes resembles an unbiased random walk truncated by k . The difficulty in distinguishing between the distributions of different queue sizes can be attributed to the fact that all distributions are primarily influenced by this unbiased random

walk. In comparison, Figure 24b presents the empirical CCDF of the queue sizes under Φ_+ , yielding the following insights:

- The queue sizes are smaller under Φ_+ compared to $\Phi_k(\mu)$.
- Due to symmetry in the matching model, Q_1 and Q_4 follow the same distribution, as do Q_2 and Q_3 . These two distributions are clearly distinguished by the simulations.



(a) Empirical queue size distribution under $\Phi_k(\mu)$. (b) Empirical queue size distributions under Φ_+ .

■ **Figure 24** Queue distributions of $\Phi_k(\mu)$ and Φ_+ for the diamond problem (30).

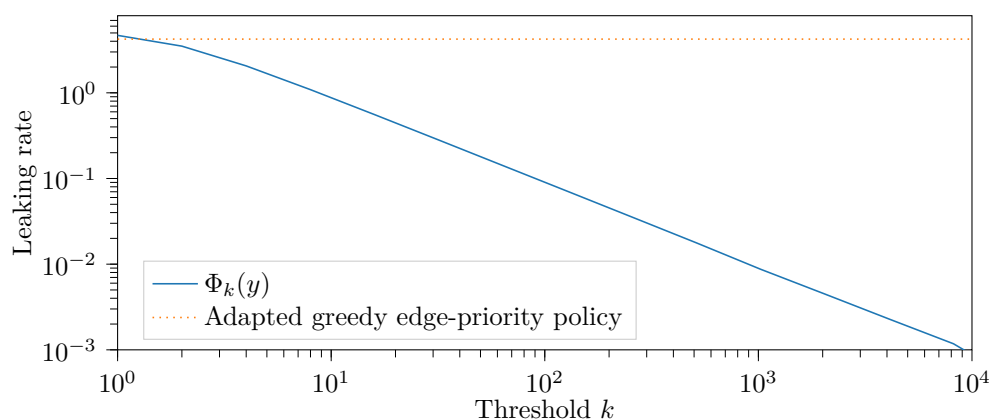
We made the same experiments on the codomino matching problem of Example 5.11. We choose the injective-only vertex μ depicted in Figure 15f: its support consists of three non-adjacent edges $\{1, 6\}, \{2, 3\}, \{4, 5\}$, that is, three connected components that are all bipartite. Edges $\{1, 2\}, \{2, 6\}, \{3, 4\}, \{3, 5\}$, and $\{5, 6\}$ should be avoided.

Figure 25 measures the *leaking rate* of $\Phi_k(\mu)$ along the five forbidden edges as k increases. As before, the convergence seems to be logarithmic, i.e., of order $1/k$. We also designed for comparison a greedy policy similar to Φ_+ but adapted to the codomino. This policy, which we also call Φ_+ by convenience, is the edge-priority policy adapted to the codomino graph whereby: edges $\{1, 6\}, \{2, 3\}, \{4, 5\}$ have the highest priority; edges $\{2, 6\}$ and $\{3, 5\}$ have the lowest priority. We leave it to the reader to verify that this policy is properly defined in the sense that all possible matching decisions are unambiguous. As for the diamond problem, we observe that $\Phi_k(\mu)$ yields a lower leaking rate than Φ_+ except for $k = 0$, where it behaves like the *match-the-longest* policy and does not take the vertex μ into account at all.

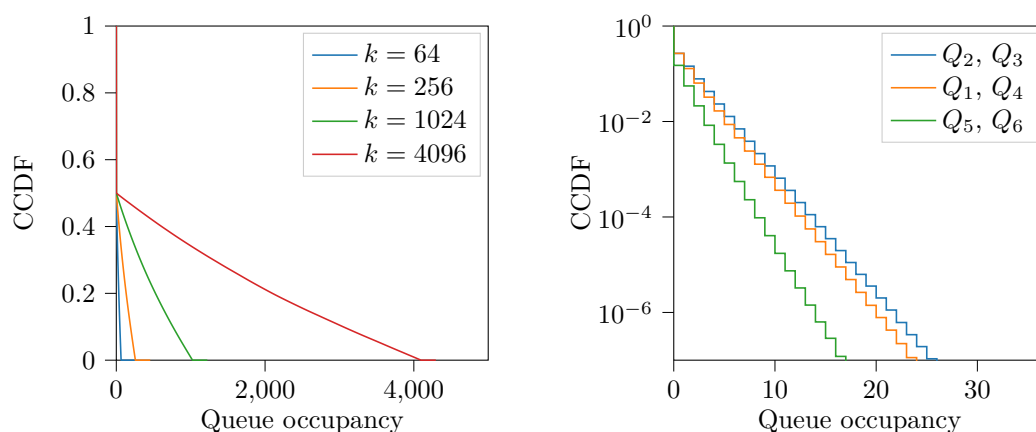
Figure 26 presents the empirical distribution of the queue sizes. Like for the diamond problem, the simulations do not show visible differences between the distributions of different classes. Figure 26a shows the (unique) empirical distribution for various values of k . We observe again that the queues are empty half of time, and that otherwise they follow a (relatively) uniform distribution on $\{0, 1, 2, \dots, k\}$. In comparison, the queues under policy Φ_+ , displayed in Figure 26b, are more compact. That Q_1, Q_2 , and Q_5 have the same distribution as Q_4, Q_3 , and Q_6 , respectively, is a consequence of the model symmetry.

References

- 1 Ivo Adan, Ana Bušić, Jean Mairesse, and Gideon Weiss. Reversibility and Further Properties of FCFS Infinite Bipartite Matching. *Mathematics of Operations Research*, 43(2):598–621, December 2017. Publisher: INFORMS. doi:10.1287/moor.2017.0874.



■ **Figure 25** Total matching rate along the forbidden edges under policy $\Phi_k(\mu)$ for the codomino problem defined from Example 5.11. μ is the vertex depicted in Figure 15f. The performance of Φ_+ , an edge-priority policy that prioritizes the support of μ , is displayed for comparison.



(a) Empirical queue size distribution under $\Phi_k(\mu)$.

(b) Empirical queue size distributions under Φ_+ .

■ **Figure 26** Queue distributions of $\Phi_k(\mu)$ and Φ_+ for the codomino problem described in Example 5.11.

- 2 Ivo Adan and Gideon Weiss. Exact FCFS matching rates for two infinite multitype sequences. *Operations Research*, 60(2):475–489, 04 2012. Publisher: INFORMS. doi:10.1287/opre.1110.1027.
- 3 Angelos Aveklouris, Levi DeValve, and Amy R. Ward. Matching Impatient and Heterogeneous Demand and Supply. *arXiv:2102.02710 [cs, eess, math]*, December 2021. arXiv: 2102.02710. URL: <http://arxiv.org/abs/2102.02710>.
- 4 Egon Balas and William R. Pulleyblank. The perfectly matchable subgraph polytope of a bipartite graph. *Networks*, 13(4):495–516, 1983. doi:10.1002/net.3230130405.
- 5 Jocelyn Begeot, Irène Marcovici, and Pascal Moyal. Stability regions of systems with compatibilities and ubiquitous measures on graphs. *Queueing Systems*, 103(3):275–312, 2023. doi:10.1007/s11134-023-09872-0.
- 6 Jocelyn Begeot, Irène Marcovici, Pascal Moyal, and Youssef Rahme. A general stochastic matching model on multigraphs. *arXiv:2011.05169 [math]*, November 2020. arXiv: 2011.05169. URL: <http://arxiv.org/abs/2011.05169>.

- 7 Péter Biró, Joris van de Klundert, David Manlove, William Petteřsson, Tommy Andersson, Lisa Burnapp, Pavel Chromy, Pablo Delgado, Piotr Dworczak, Bernadette Haase, Aline Hemke, Rachel Johnson, Xenia Klimentova, Dirk Kuypers, Alessandro Nanni Costa, Bart Smeulders, Frits Spieksma, María O. Valentín, and Ana Viana. Modelling and optimisation in european kidney exchange programmes. *European Journal of Operational Research*, 291(2):447–456, 2021. doi:10.1016/j.ejor.2019.09.006.
- 8 Pierre Bremaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Texts in Applied Mathematics. Springer-Verlag, New York, 1999. URL: <https://www.springer.com/gp/book/9780387985091>, doi:10.1007/978-1-4757-3124-8.
- 9 Ana Bušić, Varun Gupta, and Jean Mairesse. Stability of the Bipartite Matching Model. *Advances in Applied Probability*, 45(2):351–378, June 2013. Publisher: Cambridge University Press. doi:10.1239/aap/1370870122.
- 10 Ana Bušić and Sean Meyn. Approximate optimality with bounded regret in dynamic matching models. *arXiv:1411.1044 [cs, math]*, June 2016. arXiv: 1411.1044. URL: <http://arxiv.org/abs/1411.1044>.
- 11 Arnaud Cadas, Ana Bušić, and Josu Doncel. Optimal control of dynamic bipartite matching models. In *Proceedings of the 12th EAI International Conference on Performance Evaluation Methodologies and Tools, VALUETOOLS 2019, Universitat de les Illes Balears, Palma de Mallorca, Spain, March 12-15, 2019*, pages 39–46. ACM, 2019. doi:10.1145/3306309.3306317.
- 12 Arnaud Cadas, Josu Doncel, Jean-Michel Fourneau, and Ana Bušić. Flexibility can hurt dynamic matching system performance. *arXiv:2009.10009 [cs, math]*, September 2020. arXiv: 2009.10009. URL: <http://arxiv.org/abs/2009.10009>.
- 13 René Caldentey, Edward H. Kaplan, and Gideon Weiss. FCFS infinite bipartite matching of servers and customers. *Advances in Applied Probability*, 41(3):695–730, 09 2009. Publisher: Cambridge University Press. doi:10.1239/aap/1253281061.
- 14 Céline Comte and Jaap Storm. Global and local stability in parallel service systems with redundancy. Unpublished.
- 15 Céline Comte. Stochastic non-bipartite matching models and order-independent loss queues. *Stochastic Models*, 38(1):1–36, January 2022. doi:10.1080/15326349.2021.1962352.
- 16 Céline Comte and Jan-Pieter Dorsman. Performance Evaluation of Stochastic Bipartite Matching Models. In *Performance Engineering and Stochastic Modeling*, Lecture Notes in Computer Science, pages 425–440. Springer International Publishing, 2021. doi:10.1007/978-3-030-91825-5_26.
- 17 Dragoš Cvetkovic, Peter Rowlinson, and Slobodan Simic. *Spectral Generalizations of Line Graphs: On Graphs with Least Eigenvalue -2*. London Mathematical Society Lecture Note Series. Cambridge University Press, 2004. doi:10.1017/CB09780511751752.
- 18 J. G. Dai. On positive harris recurrence of multiclass queueing networks: A unified approach via fluid limit models. *The Annals of Applied Probability*, 5(1):49–77, 1995. Publisher: Institute of Mathematical Statistics. doi:10.1214/aoap/1177004828.
- 19 Reinhard Diestel. *Graph Theory*. Springer Publishing Company, Incorporated, 5th edition, August 2017.
- 20 Michael Doob. An interrelation between line graphs, eigenvalues, and matroids. *Journal of Combinatorial Theory, Series B*, 15(1):40–50, 08 1973. doi:10.1016/0095-8956(73)90030-0.
- 21 Itai Gurvich and Amy Ward. On the Dynamic Control of Matching Queues. *Stochastic Systems*, 4(2):479–523, October 2014. Publisher: INFORMS. doi:10.1287/13-SSY097.
- 22 Matthieu Jonckheere, Pascal Moyal, Claudia Ramírez, and Nahuel Soprano-Loto. Generalized max-weight policies in stochastic matching. *arXiv:2011.04535 [math]*, November 2020. arXiv: 2011.04535. URL: <http://arxiv.org/abs/2011.04535>.
- 23 N. Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing, STOC '84*, pages 302–311, New York, NY, USA, December 1984. Association for Computing Machinery. doi:10.1145/800057.808695.

- 24 F. P. Kelly. *Reversibility and Stochastic Networks*. Cambridge University Press, 2011.
- 25 László Lovász and M. D. Plummer. *Matching Theory*. American Mathematical Soc., 2009.
- 26 Jean Mairesse and Pascal Moyal. Stability of the stochastic matching model. *Journal of Applied Probability*, 53(4):1064–1077, 12 2016. Publisher: Cambridge University Press. doi:10.1017/jpr.2016.65.
- 27 Fabien Mathieu. Stochastic Matching, 2022. URL: https://balouf.github.io/stochastic_matching/index.html.
- 28 Pascal Moyal, Ana Bušić, and Jean Mairesse. A product form for the general stochastic matching model. *Journal of Applied Probability*, 58(2):449–468, 06 2021. doi:10.1017/jpr.2020.100.
- 29 Mohammadreza Nazari and Alexander L. Stolyar. Reward maximization in general dynamic matching systems. *Queueing Systems*, 91(1):143–170, February 2019. doi:10.1007/s11134-018-9593-y.
- 30 J. R. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 1997. doi:10.1017/CB09780511810633.
- 31 Youssef Rahme and Pascal Moyal. A stochastic matching model on hypergraphs. *Advances in Applied Probability*, 53(4):951–980, December 2021. doi:10.1017/apr.2021.8.
- 32 Thirupathaiiah Vasantam and Don Towsley. A throughput optimal scheduling policy for a quantum switch. In *Quantum Computing, Communication, and Simulation II*, volume 12015, pages 14–23. SPIE, 2022. doi:10.1117/12.2616950.
- 33 Günter M. Ziegler. *Lectures on polytopes*. Springer-Verlag, New York, 1995. URL: http://www.worldcat.org/search?qt=worldcat_org_all&q=9780387943657.
- 34 Martin Zubeldia, Prakirt R. Jhunjhunwala, and Siva Theja Maguluri. Matching queues with abandonments in quantum switches: Stability and throughput analysis, 2022. URL: <http://arxiv.org/abs/2209.12324>, arXiv:2209.12324[cs], doi:10.48550/arXiv.2209.12324.

A Minimal stability region for greedy matching policies

The following result gives a sufficient stability condition for greedy matching policies. The proof relies on a linear Lyapunov function. This result can be seen as the counterpart of [9, Proposition 5.1] for non-bipartite matching models.

► **Proposition A.1.** *Consider a matching problem (G, λ) with a connected graph G . If*

$$\sum_{i \in V(\mathcal{I})} \lambda_i > \frac{1}{2} \sum_{i \in V} \lambda_i, \quad \mathcal{I} \in \mathbb{I}, \quad (36)$$

then the matching model (G, λ, Φ) is stable for every greedy matching policy Φ .

Proof. Consider a matching problem (G, λ) that satisfies (36) and a greedy matching policy Φ adapted to the graph G . Since the Markov chain $(S_t, t \in \mathbb{N})$ associated with the matching model (G, λ, Φ) depends on the vector λ only up to a positive multiplicative constant, we can assume without loss of generality that $\sum_{i \in V} \lambda_i = 1$. Let \mathcal{S} denote the state space of this Markov chain and $|\cdot|$ the corresponding queue-size function. We consider the Lyapunov function $F : \mathcal{S} \rightarrow \mathbb{R}$ defined by $F(s) = \sum_{i \in V} |s|_i$ (that is, $F(s)$ is the number of unmatched items in state s) for each $s \in \mathcal{S}$. For each $t \in \mathbb{N}$ and $s \in \mathcal{S}$, we have

$$\mathbb{E}(F(S_{t+1}) \mid S_t = s) - F(s) = \sum_{i \in V \setminus V(\mathcal{I})} \lambda_i - \sum_{i \in V(\mathcal{I})} \lambda_i = - \left(\sum_{i \in V(\mathcal{I})} \lambda_i - \sum_{i \in V \setminus V(\mathcal{I})} \lambda_i \right),$$

where $\mathcal{I} = \{i \in V : |s|_i \geq 1\}$ is the set of classes of unmatched items in state s . Importantly, if $F(s) > 0$ (that is, $s \neq \emptyset$), then \mathcal{I} is an independent set of the compatibility graph G because it is non-empty and the policy Φ is greedy. It follows that, for each $s \in \mathcal{S} \setminus \{\emptyset\}$,

$$\mathbb{E}(F(S_{t+1}) \mid S_t = s) - F(s) \leq -\varepsilon, \quad \text{with } \varepsilon = \min_{\mathcal{I} \in \mathbb{I}} \left(\sum_{i \in V(\mathcal{I})} \lambda_i - \sum_{i \in V \setminus V(\mathcal{I})} \lambda_i \right).$$

Equation (36) implies that $\varepsilon > 0$. Using the Lyapunov-Foster theorem [8, Theorem 1.1 in Chapter 5], we conclude that the matching model (G, λ, Φ) is stable. \blacktriangleleft

As one would expect, any matching problem (G, λ) that satisfies (36) is stabilizable in the sense of Definition 2.4. Indeed, (36) implies Proposition 3.7(ii) because $\mathcal{I} \subseteq V \setminus V(\mathcal{I})$ for each $\mathcal{I} \in \mathbb{I}$. Corollary A.2 below shows that, conversely, whether a stabilizable matching problem satisfies (36) depends on the structure of the graph G : conditions (i) and (ii) exhibit compatibility graphs G such that (36) is satisfied whenever the matching problem (G, λ) is stabilizable, while conditions (iii) and (iv) exhibit stabilizable compatibility graphs G for which (36) is never satisfied.

► Corollary A.2. *Consider a matching problem (G, λ) .*

Under the following two conditions, the stabilizability of the matching model (G, λ) implies that (36) is satisfied, and therefore that the matching model (G, λ, Φ) is stable under every greedy policy Φ adapted to G :

- (i) *G is a complete graph with $n \geq 3$ nodes.*
- (ii) *G is the diamond graph of Example 3.13.*

Under the following conditions, (36) is never satisfied:

- (iii) *The graph G has diameter greater than or equal to 3.*
- (iv) *The graph G contains a leaf (that is, a node with degree 1).*

Proof. We first need to prove that, under either condition (i) or condition (ii), the matching model (G, λ) is stabilizable if and only if (36) is satisfied. We proceed by verifying that, under either of these two conditions, Proposition 3.7(ii) and (36) are equivalent:

- (i) First assume that condition (i) is satisfied. The independent sets of a complete graph K_n are the singletons. Using this observation, we can verify that Proposition 3.7(ii) and (36) are both equivalent to $\lambda_i < \frac{1}{2} \sum_{i' \in V} \lambda_{i'}$ for each $i \in V$.
- (ii) Now assume that condition (ii) is satisfied, that is, G is the diamond graph. The conclusion follows by recalling that Proposition 3.7(ii) simplifies to (13), and then by observing that (13) and (36) are equivalent.

To prove that (36) cannot be satisfied under either condition (iii) or (iv), we proceed by contradiction:

- (iii) First assume that condition (iii) is satisfied, and let i and j denote two nodes that are at distance 3 or more. In particular, the sets V_i and V_j are disjoint. If (36) is satisfied, then applying this equation to both $\{i\}$ and $\{j\}$ and summing the inequalities yields $\sum_{i' \in V_i \cup V_j} \lambda_{i'} > \sum_{i' \in V} \lambda_{i'}$, which is a contradiction since $V_i \cup V_j \subseteq V$. Hence, (36) cannot be satisfied by both $\{i\}$ and $\{j\}$.
- (iv) Now assume that condition (iv) is satisfied. Let i denote a leaf node of G and j the (only) neighbor of i . Then again, applying (36) to both $\{i\}$ and $\{j\}$ and summing the inequalities yields $\sum_{i' \in V_j \cup \{j\}} \lambda_{i'} > \sum_{i' \in V} \lambda_{i'}$, which is again a contradiction. \blacktriangleleft

B Proof of Proposition 5.3

The proof is based mainly on the notion of *cycle space* in a graph. We briefly summarize the concepts that are useful to understand the proof (see [19, Section 1.9] for details).

A *spanning subgraph* of a graph $G = (V, E)$ is a subgraph $G' = (V, E')$ with $E' \subseteq E$. Importantly, G and G' have the same set of nodes. A subgraph is *Eulerian* if every vertex has an even degree (possibly zero). In particular, if E' is a set of edges that form a cycle in G , then the graph (V, E') is Eulerian. The *cycle space* of G is the vector space made of all Eulerian spanning subgraphs of G , using the symmetric difference of the edge sets for addition and the two-element field for scalar multiplication. Equivalently, the cycle space can be described as a vector space of the finite field $\mathbb{Z}/2\mathbb{Z}$: each vector $g = (g_1, g_2, \dots, g_m)$ in this vector space satisfies $\sum_{k \in E_i} g_k = 0$ (modulo 2) for each $i \in V$, and the addition and multiplication are the usual operations in $\mathbb{Z}/2\mathbb{Z}$. For example, if G is the codomino graph of Figure 13, the spanning subgraphs G_1 , G_2 , and G_3 of G with edge sets $E_1 = \{\{1, 2\}, \{1, 6\}, \{2, 3\}, \{3, 4\}, \{4, 5\}, \{5, 6\}\}$, $E_2 = \{\{2, 3\}, \{2, 6\}, \{3, 5\}, \{5, 6\}\}$, and $E_3 = \{\{1, 2\}, \{1, 6\}, \{2, 6\}\}$, respectively, belong to the cycle space of G . The edge set of the addition $G_1 + G_2$ is the set $\{\{1, 2\}, \{1, 6\}, \{2, 6\}, \{3, 4\}, \{3, 5\}, \{4, 5\}\}$ of edges that are either in E_1 or in E_2 , but not in both. Similarly, the edge set of $G_1 + G_3$ is $\{\{2, 3\}, \{2, 6\}, \{3, 4\}, \{4, 5\}, \{5, 6\}\}$, and the edge set of $G_2 + G_3$ is $\{\{1, 2\}, \{1, 6\}, \{2, 3\}, \{3, 5\}, \{5, 6\}\}$. One can verify that $\{G_1, G_2, G_3\}$ forms a basis of the cycle space of G . Importantly, in general, the dimension of the cycle space is $m - n + 1$.

Proof that Algorithm 1 terminates. We first prove the existence of edge a defined on line 3 of Algorithm 1. By definition of a spanning tree, \mathcal{T} contains $n - 1$ edges and, for each edge $a \in E \setminus \mathcal{T}$, $\mathcal{T} \cup \{a\}$ contains a unique cycle. The $m - n + 1 = |E \setminus \mathcal{T}|$ cycles thus obtained are independent in the sense that each cycle contains at least one edge (a) that is not contained in the other cycles. Therefore, these $m - n + 1$ cycles form a basis of the cycle space of G . Since a linear combination of even cycles cannot produce a subgraph consisting of a single odd cycle¹⁰, and since G contains an odd cycle (as it is non-bipartite), then at least one of the $m - n + 1$ basis cycles is odd.

We now verify that, for each $s \in E \setminus (\mathcal{T} \cup \{a\})$, $\mathcal{T} \cup \{a, s\}$ contains either (i) an even cycle C_ℓ or (ii) a kayak paddle $KP_{\ell, r, p}$ with two odd cycles. By construction, $\mathcal{T} \cup \{a\}$ contains a unique cycle C_r , which is odd, and $\mathcal{T} \cup \{s\}$ contains a unique cycle C_ℓ . $\mathcal{T} \cup \{a, s\}$ contains both C_r and C_ℓ . We now proceed by elimination:

- If C_ℓ is even, then C_ℓ is an even cycle included into $\mathcal{T} \cup \{a, s\}$, and we are in case (i).
- If C_ℓ is odd and shares at least one edge with C_r , then the symmetric difference of C_r and C_ℓ is an even cycle, and it is again included into $\mathcal{T} \cup \{a, s\}$, so we are again in case (i).
- If C_ℓ is odd and C_r and C_ℓ have no edge in common, then we are in case (ii). ◀

Proof that Algorithm 1 returns the correct result. We finally prove that the family \mathcal{B} returned by Algorithm 1: (i) has cardinality $m - n$, (ii) is linearly independent, and (iii) is included into the kernel of A . We prove each item one after another:

- (i) The family \mathcal{B} has cardinality $m - n$: It suffices to observe that \mathcal{B} has same cardinality as $E \setminus (\mathcal{T} \cup \{a\})$, which we already mentioned has cardinality $m - n$.

¹⁰ A linear combination of even cycles may produce a subgraph consisting of an even number of disjoint odd cycles, as illustrated by $G_1 + G_2$ in the example above, but not a subgraph consisting of one odd cycle. Indeed, the symmetric difference of two edge sets contains an even number of edges if both edge sets contain an even number of edges.

- (ii) The family \mathcal{B} is linearly independent: For each $s \in E \setminus (\mathcal{T} \cup \{a\})$, the basis vector constructed from edge s is the only vector in \mathcal{B} whose support contains this edge.
- (iii) The family \mathcal{B} is included into the kernel of A : Let $b \in \mathcal{B}$. Our goal is to prove that b belongs to the kernel of A , i.e., that $\sum_{k \in E_i} b_k = 0$ for each $i \in V$. First observe that, for each $i \in V$, we have $\sum_{k \in E_i} b_k = \sum_{k \in E_i \cap S} b_k$, where S is the support of the vector b . In particular, we have immediately $\sum_{k \in E_i} b_k = 0$ for each $i \in V$ such that $E_i \cap S = \emptyset$. Now consider a node $i \in V$ such that $E_i \cap S \neq \emptyset$. We make a case disjunction depending on the support S of b :

- If S is an even cycle C_ℓ , then $E_i \cap S = \{k_1, k_2\}$, where k_1 and k_2 are two consecutive edges of the cycle C_ℓ . Line 10 in the algorithm implies that $b_{k_1} = -b_{k_2} \in \{1, -1\}$. It follows that $\sum_{k \in E_i} b_k = b_{k_1} + b_{k_2} = 0$.
- If S is a kayak paddle $KP_{\ell,r,p}$ with odd cycles C_ℓ and C_r and central path P_p , $p \in \mathbb{N}$, we again distinguish several cases:

Node i does not belong to the central path: If $E_i \cap S \subseteq C_\ell$ or $E_i \cap S \subseteq C_r$, we conclude as before.

Node i does not belong to a cycle: If $E_i \cap S \subseteq P_p$, then $E_i \cap S = \{k_1, k_2\}$ where k_1 and k_2 are two consecutive edges of the path P_p . Line 22 of the algorithm implies that $b_{k_1} = -b_{k_2} \in \{2, -2\}$. It follows that $\sum_{k \in E_i} b_k = b_{k_1} + b_{k_2} = 0$.

Node i belongs to a cycle and the central path: The only remaining case is when $E_i \cap S$ intersects several sets among C_ℓ , C_r , and P_p . If $p = 0$, that is, if the central path is the node i , then $E_i \cap S = \{k_1, k_2, k_3, k_4\}$, where k_1 and k_2 (resp. k_3 and k_4) are two consecutive edges in C_ℓ (resp. C_r). Lines 18 and 26 yield $b_{k_1} = b_{k_2} = -1$ and $b_{k_3} = b_{k_4} = 1$, which implies that $\sum_{k \in E_i} b_k = b_{k_1} + b_{k_2} + b_{k_3} + b_{k_4} = 0$. If $p \geq 1$, then $E_i \cap S = \{k_1, k_2, k_3\}$, where k_1 and k_2 are two consecutive edges of either C_ℓ or C_r , and k_3 is an edge in P_p . Lines 18, 22, and 26 of the algorithm imply that $b_{k_1} = b_{k_2} \in \{1, -1\}$ and that $b_{k_3} = -2b_{k_1}$, so that we conclude again that the desired sum is zero. \blacktriangleleft

C Proofs of the results of Section 6.1

C.1 Proof of Proposition 6.1

Consider a stabilizable matching problem (G, λ) and let $\Pi_{\mathcal{P}}$ (resp. Π_G) denote the set of vectors of matching rates achievable by stable policies (resp. by stable greedy policies) adapted to the compatibility graph G .

Convexity of $\Pi_{\mathcal{P}}$

Consider two (extended) policies Φ_1 and Φ_2 that stabilize the matching problem (G, λ) . The state-space, queue-size function, and empty state of Φ_1 (resp. Φ_2) are denoted by $\mathcal{S}_1, |\cdot|_1$, and \emptyset_1 (resp. $\mathcal{S}_2, |\cdot|_2$, and \emptyset_2). Given $0 < \beta < 1$, our goal is to build a matching policy Φ_β that also stabilizes the matching problem (G, λ) and satisfies $\mu(\Phi_\beta) = \beta\mu(\Phi_1) + (1 - \beta)\mu(\Phi_2)$.

Intuitively, the policy Φ_β will consistently follow the decisions of either Φ_1 or Φ_2 as long as the system is non-empty, and switch between these two policies each time the system becomes empty (therefore, at renewal times). The probability of choosing either policy after visiting the empty state will be chosen to achieve the desired matching rate vector on the long run. More formally, the policy Φ_β will have the following characteristics:

- State space $\mathcal{S}_\beta = (\mathcal{S}_1 \times \{\emptyset_2\}) \cup (\{\emptyset_1\} \times \mathcal{S}_2)$.
- Queue-size function $|\cdot|_\beta$ defined by $|(s_1, s_2)|_\beta = |s_1|_1 + |s_2|_2$ for each $(s_1, s_2) \in \mathcal{S}_\beta$.

■ Empty state: $\emptyset_\beta = (\emptyset_1, \emptyset_2)$.

For some $0 < \gamma_\beta < 1$ that will be specified later, the matching policy Φ_β is defined as follows:

$$\begin{aligned}\Phi_\beta((\emptyset_1, \emptyset_2), i, j, (s_1, \emptyset_2)) &= \gamma_\beta \Phi_1(\emptyset_1, i, j, s_1), & s_1 \in \mathcal{S}_1, \quad i \in V, \quad j \in V \cup \{\perp\}, \\ \Phi_\beta((\emptyset_1, \emptyset_2), i, j, (\emptyset_1, s_2)) &= (1 - \gamma_\beta) \Phi_2(\emptyset_2, i, j, s_2), & s_2 \in \mathcal{S}_2, \quad i \in V, \quad j \in V \cup \{\perp\}, \\ \Phi_\beta((s_1, \emptyset_2), i, j, (s'_1, \emptyset_2)) &= \Phi_1(s_1, i, j, s'_1), & s_1 \in \mathcal{S}_1 \setminus \{\emptyset_1\}, \quad s'_1 \in \mathcal{S}_1, \\ \Phi_\beta((\emptyset_1, s_2), i, j, (\emptyset_1, s'_2)) &= \Phi_2(s_2, i, j, s'_2), & s_2 \in \mathcal{S}_2 \setminus \{\emptyset_2\}, \quad s'_2 \in \mathcal{S}_2.\end{aligned}$$

The first two equations say that, when the system is empty, we next apply Φ_1 with probability γ_β and Φ_2 with probability $1 - \gamma_\beta$. The third (resp. fourth) equation says that, once we start applying policy Φ_1 (resp. Φ_2), we keep applying this policy until we re-visit the empty state.

The stability of the policies Φ_1 and Φ_2 implies that of the policy Φ_β . According to the elementary renewal theorem for renewal reward processes, the vector giving the long-run average matching rates under the policy Φ_β is given by

$$\begin{aligned}\mu(\Phi_\beta) &= \frac{\gamma_\beta T_1 \mu(\Phi_1) + (1 - \gamma_\beta) T_2 \mu(\Phi_2)}{\gamma_\beta T_1 + (1 - \gamma_\beta) T_2}, \\ &= \frac{\gamma_\beta T_1}{\gamma_\beta T_1 + (1 - \gamma_\beta) T_2} \mu(\Phi_1) + \frac{(1 - \gamma_\beta) T_2}{\gamma_\beta T_1 + (1 - \gamma_\beta) T_2} \mu(\Phi_2).\end{aligned}$$

where T_1 (resp. T_2) denotes the mean number of matches between two successive visits to the empty state \emptyset_1 (resp. \emptyset_2) in the matching model (G, λ, Φ_1) (resp. (G, λ, Φ_2)). We let the reader verify that the following value of γ_β yields $\mu(\Phi_\beta) = \beta \mu(\Phi_1) + (1 - \beta) \mu(\Phi_2)$:

$$\gamma_\beta = \frac{\beta T_2}{(1 - \beta) T_1 + \beta T_2}.$$

Convexity of Π_G

It suffices to observe that, if the policies Φ_1 and Φ_2 are greedy, so is the policy Φ_β .

C.2 Proof of the results of Section 6.1.1 for the complete graph

Proposition 6.4(i)

Consider a (possibly random) size-based greedy policy $\Phi : \mathcal{Q}_G(K_n) \times V \times (V \cup \{\perp\}) \rightarrow [0, 1]$, as defined in Remark 2.1, where $\mathcal{Q}_G(K_n)$ is given in (25). Given $(q, i) \in \mathcal{Q}_G(K_n) \times V$, the definition of $\mathcal{Q}_G(K_n)$ implies that $\{j \in V_i : q_j \geq 1\}$ is either a singleton or the emptyset. In the former case, letting j denote the unique element of the singleton, we have $\Phi(q, i, \perp) = 1$ if $i = j$, while the greediness of Φ implies that $\Phi(q, i, j) = 1$ if $i \neq j$. In the latter case, we have directly $\Phi(q, i, \perp) = 1$. In all cases, the matching decision is deterministic and makes the same decisions as in (26).

Proposition 6.4(ii)

The same argument can be repeated for an arbitrary greedy policy Φ with state space $(\mathcal{S}, |\cdot|)$. Given $(s, i) \in \mathcal{S} \times V$, we know that $|s| \in \mathcal{Q}_G(K_n)$, so that $\{j \in V_i : |s|_j \geq 1\}$ is either a singleton or the emptyset. In the former case, letting j denote the unique element of the singleton, we have $\sum_{s' \in \mathcal{S}} \Phi(s, i, \perp, s') = 1$ if $i = j$, while the greediness of Φ implies that $\sum_{s' \in \mathcal{S}} \Phi(s, i, j, s') = 1$ if $i \neq j$. In the latter case, we have directly $\sum_{s' \in \mathcal{S}} \Phi(s, i, \perp, s') = 1$. In all cases, we have $\sum_{s' \in \mathcal{S}} \Phi(s, i, \Phi_G(|s|, i), s') = 1$.

Proposition 6.4(iii)

This statement follows directly from Proposition 6.4(ii) and from the fact that, under any matching policy, the dynamics of the queue-size process is fully specified by the sequence of incoming item classes and the matching decisions via (2).

Proposition 6.5(i)

Proposition 6.5(i) is a consequence of Corollary A.2(i).

Proposition 6.5(ii)

We can focus without loss of generality on the first-come-first-matched policy, as Proposition 6.4(iii) implies that all greedy matching policies yield the same vector of matching rates. Equation (27) follows by observing that, for each edge $e_k = \{i, j\} \in E$, a match between classes i and j happens in one of the following two cases:

- a class- i item arrives while queue j is non-empty, which happens at rate $\lambda_i p_j$;
- a class- j item arrives while queue i is non-empty, which happens at rate $\lambda_j p_i$.

Equations (28) and (29) follow from [15, Proposition 5]. Indeed, for each $i \in V$, applying [15, Equation (10)] to the independent set $\{i\}$ yields (28), and the value of p_\emptyset given in (29) follows from the normalizing equation. This result may also be obtained more directly by observing that, for each $i \in V$, the restriction of the transition diagram of the Markov chain (K_n, λ, Φ_G) to the states where all queues but queue i are empty is similar to a (discrete-time) birth-and-death process with birth probability λ_i and death probability $\sum_{j \neq i} \lambda_j$.

C.3 Proof of the results of Section 6.1.2 for the diamond graph**Proposition 6.6(i)**

We leave it to the reader to verify that $\mathcal{Q}_G(K_3)$ is the image of $\mathcal{Q}_G(D)$ by the application $q = (q_1, q_2, q_3, q_4) \mapsto \langle q \rangle = (q_1 + q_4, q_2, q_3)$. This means in particular that $q_1 + q_4$, q_2 , and q_3 cannot be positive simultaneously if $q \in \mathcal{Q}_G(D)$. Since the diamond graph D has only four nodes, we can then conclude by enumerating all relevant cases, depending on the support of $\langle q \rangle = (q_1 + q_4, q_2, q_3)$. For example, if $q_1 + q_4 \geq 1$ and $q_2 = q_3 = 0$, then we have immediately $\Phi(q, i) = \perp$ if $i \in \{1, 4\}$, and the greediness of the policy Φ implies that $\Phi(q, i) \in \{1, 4\}$ if $i \in \{2, 3\}$; in other words, we have $\langle \Phi(q, i) \rangle = \perp$ if $\langle i \rangle = 1$ and $\langle \Phi(q, i) \rangle = 1$ if $\langle i \rangle \in \{2, 3\}$. Similarly, if $q_1 + q_4 = q_3 = 0$ and $q_2 \geq 1$, then $\Phi(q, i) = \perp$ if $i = 2$ and $\Phi(q, i) = 2$ if $i \in \{1, 3, 4\}$, that is, $\langle \Phi(q, i) \rangle = \perp$ if $\langle i \rangle = 2$ and $\langle \Phi(q, i) \rangle = 2$ if $\langle i \rangle \in \{1, 3\}$. In all cases, we can verify that $\langle \Phi(q, i) \rangle$ is equal to $\Phi_G(K_3)(\langle q \rangle, \langle i \rangle)$.

Alternatively, by taking a step back, we can prove Proposition 6.6(i) more directly by observing that, for each $(q, i) \in \mathcal{Q}_G(D) \times V$, (i) the support of $\langle q \rangle$ is a singleton $\{j\}$ whenever $q \neq 0$, and (ii) whether or not $j \in V_i$ depends on i only via $\langle i \rangle$.

Proposition 6.6(ii)

The same argument can be repeated for an arbitrary greedy policy Φ with state space $(\mathcal{S}, |\cdot|)$. As in the proof of Proposition 6.4(ii), the key argument consists of observing that we still have $q = |s| \in \mathcal{Q}_G(D)$ for each $(s, i) \in \mathcal{S} \times V$, so that $q_1 + q_4$, q_2 , and q_3 cannot be positive simultaneously.

Proposition 6.6(iii)

The conclusion follows in much the same way as in the proof of Proposition 6.4(iii), by injecting statements (i) and (ii) from Proposition 6.6 into (2).

Proposition 6.7(i)

This is a consequence of Corollary A.2(ii).

Proposition 6.7(ii)

The equation $\mu_{2,3} = \beta = \frac{1}{2}(\lambda_2 + \lambda_3 - \lambda_1 - \lambda_4)$ is a direct consequence of (CE). The inequalities (32) for $\mu_{1,2}$, $\mu_{1,3}$, $\mu_{2,4}$, and $\mu_{3,4}$ follow by observing that, for each edge $\{i, j\} \in \{\{1, 2\}, \{1, 4\}, \{2, 3\}, \{3, 4\}\}$, a match between classes i and j happens at least in one of the following cases:

- a class- i item arrives while the system contains unmatched items that all belong to class j ,
 - a class- j item arrives while the system contains unmatched items that all belong to class i .
- These events occur at rates $\lambda_i p_j$ and $\lambda_j p_i$, respectively. Equations (32) are not equalities in general because the above list is not exhaustive. For example, depending on the greedy policy, a match between classes 1 and 2 may happen if a class-2 item arrives while the system contains unmatched items of class 1 and unmatched items of class 4.

Proposition 6.7(iii)

The expressions for p_\emptyset , p_2 , p_3 , and $p_{1,4}$ follow directly by combining Proposition 6.6(iii) with Equations (28) and (29) in Proposition 6.5. We now derive the lower bound (34) for p_1 . The one for p_4 follows by symmetry.

First assume that the greedy policy Φ is deterministic and size-based, so that it satisfies Proposition 6.6(i). For each $q \in \mathcal{Q}_G(D)$, let π_q denote the probability that the Markov chain (D, λ, Φ) is in state q in stationary regime. The probability that we want to lower-bound is

$$p_1 = \sum_{\ell=1}^{+\infty} \pi_{\ell \mathbb{1}_1}. \quad (37)$$

Now let $\ell \in \mathbb{N}_{>0}$ and consider the balance equation for state $\ell \mathbb{1}_1$, given by

$$(\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4)\pi_{\ell \mathbb{1}_1} = \lambda_1 \pi_{(\ell-1)\mathbb{1}_1} + (\lambda_2 + \lambda_3)\pi_{(\ell+1)\mathbb{1}_1} + C, \quad (38)$$

where C is a non-negative real that depends on the model parameters, the integer ℓ , and the policy Φ , and that accounts for the flow to state $\ell \mathbb{1}_1$ from state $\ell \mathbb{1}_1 + \mathbb{1}_4$, if any. It follows that $(\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4)\pi_{\ell \mathbb{1}_1} > \lambda_1 \pi_{(\ell-1)\mathbb{1}_1}$. An inductive argument allows us to conclude that

$$\pi_{\ell \mathbb{1}_1} > \left(\frac{\lambda_1}{\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4} \right)^\ell p_\emptyset, \quad \ell \in \mathbb{N}_{>0}. \quad (39)$$

Injecting this inequality into (37) allows us to conclude:

$$p_1 > \sum_{\ell=1}^{+\infty} \left(\frac{\lambda_1}{\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4} \right)^\ell p_\emptyset = \frac{\lambda_1}{\lambda_2 + \lambda_3 + \lambda_4} p_\emptyset. \quad (40)$$

If Φ is an extended policy with state space $(\mathcal{S}, |\cdot|)$, we can still write (37)–(40) and reach the same conclusion. The only difference is that π can no longer be defined as the stationary distribution of a Markov chain: instead, we define $\pi_q = \sum_{s \in \mathcal{S}: |s|=q} \varpi_s$ for each $q \in \mathcal{Q}_G(D)$, where ϖ_s is the probability that the Markov chain (D, λ, Φ) is in state s in stationary regime, for each $s \in \mathcal{S}$.

Proposition 6.8(i)–(iia)–(iib)

The diamond graph D has $n = 4$ nodes and $m = 5$ edges. Therefore, according to Propositions 5.1 and 6.1, the sets $\Pi_{\geq 0}$, $\Pi_{> 0}$, and $\Pi_{\mathcal{G}}$ have dimension $d = m - n = 1$, meaning that they are intervals in \mathbb{R} . The equations for the intervals $\Pi_{\geq 0}$ and $\Pi_{> 0}$ follow directly from the change-of-basis equation $\mu = (2\lambda_1\bar{\lambda}_2 + \alpha, 2\lambda_1\bar{\lambda}_3 - \alpha, \beta, 2\bar{\lambda}_2\lambda_4 - \alpha, 2\bar{\lambda}_3\lambda_4 + \alpha)$ from (30). That $\Pi_{\mathcal{G}}$ is also an interval is a consequence of Proposition 6.1. The (non-strict) inequality $-2\min(\lambda_1\bar{\lambda}_2, \bar{\lambda}_3\lambda_4) \leq \alpha_- \leq \alpha_+ \leq 2\min(\lambda_1\bar{\lambda}_3, \bar{\lambda}_2\lambda_4)$ is a consequence of Proposition 6.2, which states that $\Pi_{\mathcal{G}} \subseteq \Pi_{> 0}$. The first and third inequalities are also strict because α_+ and α_- belong to $\Pi_{\mathcal{G}}$ (see below), while $2\min(\lambda_1\bar{\lambda}_3, \bar{\lambda}_2\lambda_4)$ and $-2\min(\lambda_1\bar{\lambda}_2, \bar{\lambda}_3\lambda_4)$ do not belong to $\Pi_{> 0}$. That $\Pi_{\mathcal{G}}$ is a closed interval of the form $\Pi_{\mathcal{G}} = [\alpha_-, \alpha_+]$ and that α_+ and α_- are as given in Proposition 6.8(iia)–(iib) are consequences of Lemma C.1 below, which will be proved by a coupling argument later in this appendix.

► **Lemma C.1.** *Consider the edge-priority policy Φ_+ adapted to the diamond graph D whereby edges $\{1, 2\}$ and $\{3, 4\}$ have the highest priority. For each greedy policy Φ adapted to the compatibility graph D , we have*

$$\begin{aligned} \mu_{1,2}(\Phi) &\leq \mu_{1,2}(\Phi_+), & \mu_{3,4}(\Phi) &\leq \mu_{3,4}(\Phi_+), \\ \mu_{1,3}(\Phi) &\geq \mu_{1,3}(\Phi_+), & \mu_{2,4}(\Phi) &\geq \mu_{2,4}(\Phi_+), \end{aligned} \quad (41)$$

Equivalently, in kernel coordinates, we have $\alpha(\Phi) \leq \alpha(\Phi_+)$.

Proof. See later in this appendix. ◀

Proposition 6.8(iic)

Most of the quantities we consider in this proof are functions of β , but this dependency is left implicit to simplify notation. In particular, we let μ denote (the edge coordinates of) the matching rate vector in the matching model (D, λ, Φ_+) , where $\lambda = (\lambda_1, \bar{\lambda}_2 + \beta, \bar{\lambda}_3 + \beta, \lambda_4)$, with $\lambda_1 + \lambda_4 = \bar{\lambda}_2 + \bar{\lambda}_3 = \frac{1}{2}$, and Φ_+ is the greedy edge-priority policy defined in Lemma C.1. By combining (32) with the conservation equation $\mu_{1,2} + \mu_{1,3} = \lambda_1$, we obtain the following lower and upper bounds for $\mu_{1,2}$:

$$\lambda_1 p_2 + \lambda_2 p_1 \leq \mu_{1,2} \leq \lambda_1 - (\lambda_1 p_3 + \lambda_3 p_1). \quad (42)$$

In addition, injecting the definition (30) of λ into (33) and (34) shows that, in the model (D, λ, Φ_+) , we have

$$\begin{aligned} p_2 &= p_{\emptyset} \frac{\bar{\lambda}_2 + \beta}{2\bar{\lambda}_3}, & p_3 &= p_{\emptyset} \frac{\bar{\lambda}_3 + \beta}{2\bar{\lambda}_2}, & p_{1,4} &= p_{\emptyset} \frac{1}{3\beta}, \\ p_1 &> p_{\emptyset} \frac{\lambda_1}{\frac{1}{2} + \lambda_4 + 2\beta}, & p_4 &> p_{\emptyset} \frac{\lambda_4}{\frac{1}{2} + \lambda_1 + 2\beta}, \end{aligned} \quad (43)$$

with, by the normalization equation $p_{\emptyset} + p_2 + p_3 + p_{1,4} = 1$,

$$p_{\emptyset} = \left(1 + \frac{1}{3\beta} + \frac{\bar{\lambda}_2 + \beta}{2\bar{\lambda}_3} + \frac{\bar{\lambda}_3 + \beta}{2\bar{\lambda}_2} \right)^{-1}. \quad (44)$$

Taking the limit of (43) and (44) as $\beta \rightarrow +\infty$, we conclude that both the lower-bound and the upper-bound in (42) tend to $2\lambda_1\bar{\lambda}_2$ as $\beta \rightarrow +\infty$. Then combining (42) with the squeeze theorem allows us to conclude that $\mu_{1,2}$ also tends to $2\lambda_1\bar{\lambda}_2$ as $\beta \rightarrow +\infty$. By symmetry, we obtain directly $\mu_{1,3} \rightarrow 2\lambda_1\bar{\lambda}_3$, $\mu_{2,4} \rightarrow 2\bar{\lambda}_2\lambda_4$, and $\mu_{3,4} \rightarrow 2\bar{\lambda}_3\lambda_4$ as $\beta \rightarrow +\infty$. According to the change-of-basis equation $\mu = (2\lambda_1\bar{\lambda}_2 + \alpha, 2\lambda_1\bar{\lambda}_3 - \alpha, \beta, 2\bar{\lambda}_2\lambda_4 - \alpha, 2\bar{\lambda}_3\lambda_4 + \alpha)$ from (30), this means that $\alpha_+ \rightarrow 0$ as $\beta \rightarrow +\infty$.

Lemma C.1

Consider a greedy policy Φ adapted to the graph D . We will prove the inequality relations in (41) using a coupling argument. More specifically, we will compare the matching models (D, I, Φ) and (D, I, Φ_+) , where $I = (I_t, t \in \mathbb{N})$ is a sequence of i.i.d. classes, such that $I_t = i$ with probability $\lambda_i/(\lambda_1 + \lambda_2 + \lambda_4 + \lambda_3)$ for each $t \in \mathbb{N}$ and $i \in V$.

Considering the model (D, I, Φ) , we let Q_t denote the vector of queue sizes at time t , $L_{t,i}$ the number of class- i items among the first t arrivals, and $M_{t,\{i,j\}}$ (or $M_{t,i,j}$ for short) the number of times that classes i and j are matched over the first t arrivals, for each $t \in \mathbb{N}$ and $i, j \in V$, as defined in (2)–(5). We introduce similar notation for the model (D, I, Φ^+) , the only difference being that all quantities have superscript $+$. Since both models have the same sequence of incoming items, we have $L_{t,i} = L_{t,i}^+$ for each $t \in \mathbb{N}$ and $i \in V$. As usual, we also assume that $Q_0 = Q_0^+ = 0$. Neither $(Q_t, t \in \mathbb{N})$ nor $(Q_t^+, t \in \mathbb{N})$ need be Markov chains for our argument to hold.

Our end goal is to prove that the following inequalities are satisfied at each time $t \in \mathbb{N}$:

$$M_{t,1,2} \leq M_{t,1,2}^+, \quad (45-1,2)$$

$$M_{t,3,4} \leq M_{t,3,4}^+, \quad (45-3,4)$$

$$M_{t,1,3} \geq M_{t,1,3}^+, \quad (45-1,3)$$

$$M_{t,2,4} \geq M_{t,2,4}^+. \quad (45-2,4)$$

Injecting these inequalities into the definition (9) of the matching rates yields the inequalities (41). We will prove (45) by induction over time $t \in \mathbb{N}$. The following equations will be instrumental to prove the induction step:

$$Q_{t,1} + Q_{t,4} = Q_{t,1}^+ + Q_{t,4}^+, \quad Q_{t,2} = Q_{t,2}^+, \quad Q_{t,3} = Q_{t,3}^+, \quad (46)$$

$$L_{t,1} = Q_{t,1} + M_{t,1,2} + M_{t,1,3} = Q_{t,1}^+ + M_{t,1,2}^+ + M_{t,1,3}^+, \quad (47)$$

$$L_{4,t} = Q_{t,4} + M_{t,2,4} + M_{t,3,4} = Q_{t,4}^+ + M_{t,2,4}^+ + M_{t,3,4}^+. \quad (48)$$

Equation (46) follows from Proposition 6.6(iii), while (47) and (48) follow from our assumption that the arrivals in both models are coupled. Furthermore, given the definition of $Q_{\mathcal{G}}(D)$ in (31), we know that only one integer among $Q_{t,1} + Q_{t,4}$, $Q_{t,2}$, and $Q_{t,3}$ can be positive, for each $t \in \mathbb{N}$.

We now proceed to the induction step. Let $t \in \mathbb{N}$. We now prove that, assuming that the inequalities (45) are satisfied at time t , these inequalities are also satisfied at time $t + 1$. We distinguish several cases depending on the value of I_t :

Case $I_t = 1$: We have directly $M_{t+1,i,j} = M_{t,i,j}$ and $M_{t+1,i,j}^+ = M_{t,i,j}^+$ for $(i, j) \in \{(2, 4), (3, 4)\}$,

hence the induction assumption implies that (45-2,4) and (45-3,4) hold at time $t + 1$.

Since the policy Φ is greedy, we only have three mutually-exclusive cases:

Case $Q_{t,2} \geq 1$: The class-1 item is matched with a class-2 item already present, and we obtain $M_{t+1,1,2} = M_{t,1,2} + 1$ and $M_{t+1,1,3} = M_{t,1,3}$.

Case $Q_{t,3} \geq 1$: The class-1 item is matched with a class-3 item already present, and we obtain $M_{t+1,1,2} = M_{t,1,2}$ and $M_{t+1,1,3} = M_{t,1,3} + 1$.

Case $Q_{t,2} = Q_{t,3} = 0$: The class-1 item is left unmatched, and we obtain $M_{t+1,1,2} = M_{t,1,2}$ and $M_{t+1,1,3} = M_{t,1,3}$.

Since the policy Φ^+ is also greedy, we can repeat the same argument for the quantities associated with Φ^+ . Combining this observation with (46) yields $M_{t+1,1,2} - M_{t+1,1,2}^+ = M_{t,1,2} - M_{t,1,2}^+$ and $M_{t+1,1,3} - M_{t+1,1,3}^+ = M_{t,1,3} - M_{t,1,3}^+$. Hence, the induction assumption implies directly that (45-1,2) and (45-1,3) are satisfied at time $t + 1$.

Case $I_t = 2$: We have directly $M_{t+1,i,j} = M_{t,i,j}$ and $M_{t+1,i,j}^+ = M_{t,i,j}^+$ for $(i,j) \in \{(1,3), (3,4)\}$, hence the induction assumption implies that (45-1,3) and (45-3,4) hold at time $t+1$. Proving that (45-1,2) and (45-2,4) also hold at time $t+1$ is more intricate, and we will distinguish three mutually-exclusive cases depending on the values of $Q_{t,1}$, $Q_{t,4}$, $Q_{t,1}^+$, and $Q_{t,4}^+$:

Case $Q_{t,1}^+ + Q_{t,4}^+ = 0$: Under both policies, the class-2 item is either matched with a class-3 item or added to the queue. In particular, we obtain $M_{t+1,i,j} = M_{t,i,j}$ and $M_{t+1,i,j}^+ = M_{t,i,j}^+$ for $(i,j) \in \{(1,2), (2,4)\}$, so that (45-1,2) and (45-2,4) are again satisfied at time $t+1$ thanks to the induction assumption.

Case $Q_{t,1}^+ + Q_{t,4}^+ \geq 1$: We again subdivide this case into three mutually-exclusive cases:

Case $Q_{t,1}^+ \geq 1$ and $Q_{t,4}^+ \geq 1$: We have $M_{t+1,1,2}^+ = M_{t,1,2}^+ + 1$ and $M_{t+1,2,4}^+ = M_{t,2,4}^+$ by definition of the policy Φ^+ , while for the policy Φ we only know that $M_{t+1,1,2} \in \{M_{t,1,2}, M_{t,1,2} + 1\}$ and $M_{t+1,2,4} \in \{M_{t,2,4}, M_{t,2,4} + 1\}$. We can verify that (45-1,2) and (45-2,4) hold at time $t+1$ thanks to the induction assumption.

Case $Q_{t,1}^+ = 0$ and $Q_{t,4}^+ \geq 1$: By greediness, the policy Φ^+ matches the incoming class-2 item with a class-4 item, and we obtain $M_{t+1,1,2}^+ = M_{t,1,2}^+$ and $M_{t+1,2,4}^+ = M_{t,2,4}^+ + 1$. If the policy Φ makes the same decision, then we also have $M_{t+1,1,2} = M_{t,1,2}$ and $M_{t+1,2,4} = M_{t,2,4} + 1$, hence (45-1,2) and (45-2,4) hold at time $t+1$ thanks to the induction assumption. Otherwise, the policy Φ matches the class-2 item with a class-1 item, meaning that $M_{t+1,1,2} = M_{t,1,2} + 1$ and $M_{t+1,2,4} = M_{t,2,4}$. Importantly, this is only possible if $Q_{t,1} \geq 1$. We now prove (45-1,2) and (45-2,4) as follows:

- Proving (45-1,2) boils down to proving $M_{t,1,2}^+ \geq M_{t,1,2} + 1$. We have successively:

$$M_{t,1,2}^+ - M_{t,1,2} = (Q_{t,1} - Q_{t,1}^+) + (M_{t,1,3} - M_{t,1,3}^+) \geq 1 + 0 = 1,$$

where the equality follows from (47) and the inequality follows from the induction assumption and the fact that $Q_{t,1} \geq 1$ and $Q_{t,1}^+ = 0$.

- Proving (45-2,4) boils down to proving $M_{t,2,4}^+ + 1 \leq M_{t,2,4}$. We have successively

$$M_{t,2,4} - M_{t,2,4}^+ = (Q_{t,4}^+ - Q_{t,4}) + (M_{t,3,4}^+ - M_{t,3,4}) \geq 1 + 0 = 1,$$

where the equality follows from (48) and the inequality follows from the induction assumption and the observation that $Q_{t,4}^+ - Q_{t,4} = Q_{t,1} - Q_{t,1}^+ \geq 1$.

Intuitively, the only way that an incoming class-2 item is matched at time t with a class-1 item under the policy Φ and with a class-4 item under the policy Φ^+ is if, in the past, the policy Φ^+ had made one more match along edge $\{1,2\}$ and one less match along edge $\{2,4\}$ compared to the policy Φ .

Case $Q_{t,1}^+ \geq 1$ and $Q_{t,4}^+ = 0$: This case is symmetrical to the previous case.

Case $I_t = 3$: This case is symmetrical to the case $I_t = 2$.

Case $I_t = 4$: This case is symmetrical to the case $I_t = 1$.