



HAL
open science

Generalized max-weight policies in stochastic matching

Matthieu Jonckheere, Pascal Moyal, Claudia Ramírez, Nahuel Soprano-Loto

► **To cite this version:**

Matthieu Jonckheere, Pascal Moyal, Claudia Ramírez, Nahuel Soprano-Loto. Generalized max-weight policies in stochastic matching. *Stochastic Systems*, 2023, 13 (1), pp.40-58. 10.1287/stsy.2022.0098 . hal-03354805

HAL Id: hal-03354805

<https://hal.science/hal-03354805>


Submitted on 5 Dec 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.



Distributed under a Creative Commons Attribution 4.0 International License



Stochastic Systems

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Generalized Max-Weight Policies in Stochastic Matching

Matthieu Jonckheere, Pascal Moyal, Claudia Ramírez, Nahuel Soprano-Loto

To cite this article:

Matthieu Jonckheere, Pascal Moyal, Claudia Ramírez, Nahuel Soprano-Loto (2023) Generalized Max-Weight Policies in Stochastic Matching. *Stochastic Systems* 13(1):40-58. <https://doi.org/10.1287/stsy.2022.0098>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Generalized Max-Weight Policies in Stochastic Matching

 Matthieu Jonckheere,^a Pascal Moyal,^b Claudia Ramírez,^c Nahuel Soprano-Loto^{c,d,*}
^aLAAS-CNRS (Laboratoire d'analyse et d'architecture des systèmes)-(Centre national de la recherche scientifique), 31400 Toulouse, France;

^bFaculté des Sciences et Technologies, IECL/INRIA PASTA (Institut Élie Cartan de Lorraine)/(Institut national de recherche en sciences et technologies du numérique, Processus Aléatoires Spatio-Temporels et Applications), Site de Nancy, Faculté des Sciences et Technologies Campus, 54506 Vandoeuvre-lès-Nancy, France; ^cAristas S.R.L., CABA (Ciudad Autónoma de Buenos Aires) C1414CLO, Argentina;

^dDepartamento de Matemática, Facultad de Ingeniería, Universidad de Buenos Aires, CABA C1063ACV, Argentina

*Corresponding author

Contact: matthieu.jonckheere@laas.fr,  <https://orcid.org/0000-0003-3576-5866> (MJ); pascal.moyal@univ-lorraine.fr,

 <https://orcid.org/0000-0002-6638-5551> (PM); c.ramirez@aristas.com.ar (CR); nsoprano@fi.uba.ar,

 <https://orcid.org/0000-0002-5120-3903> (NS-L)

Received: July 14, 2021

Revised: January 29, 2022; May 27, 2022

Accepted: June 6, 2022

Published Online in Articles in Advance:
August 1, 2022

<https://doi.org/10.1287/stsy.2022.0098>
Copyright: © 2022 The Author(s)

Abstract. We consider a matching system where items arrive one by one at each node of a compatibility network according to Poisson processes and depart from it as soon as they are matched to a compatible item. The matching policy considered is a generalized max-weight policy where decisions can be noisy. Additionally, some of the nodes may have impatience, that is, leave the system before being matched. Using specific properties of the max-weight policy, we construct a simple quadratic Lyapunov function. This allows us to establish stability results, to prove exponential convergence speed toward the stationary measure, and to give explicit bounds for the stationary mean of the largest queue size in the system. We finally illustrate some of these results using simulations on toy examples.



Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as "Stochastic Systems. Copyright © 2022 The Author(s). <https://doi.org/10.1287/stsy.2022.0098>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>."

Funding: N. Soprano-Loto is partially supported by Aristas, MathAmSud [Grant RSPSM (Random Structures and Processes in Statistical Mechanics)], and ANR (Agence Nationale de la Recherche) [Grant PRC (Projets de Recherche Collaborative) MATCHES (CE40)].

Keywords: stochastic matching • max-weight • stability • Lyapunov function

1. Introduction

The so-called *bipartite stochastic matching model* was introduced in Caldentey et al. (2009) (see also Adan and Weiss (2012) and Adan et al. (2017)) as a variant of skill-based service systems. In this model, two infinite ordered sequences (a sequence of servers and a sequence of customers) are considered, and customer/server are matched in first come, first served order, that is, from left to right, according to compatibilities that are prescribed by a given bipartite graph. In the so-called *extended bipartite matching model* (Bušić et al. 2013), customer/server couples enter the system at each time point. A matching occurs when an arriving customer (respectively, server) finds a compatible server (respectively, customer) present in the system, in which case both leave the system instantaneously. Otherwise, items remain in the system waiting for compatible arrivals (in particular, there are no service times). The field of applications of such models is very wide, from housing to job applications, from transplant protocols to blood banks, peer-to-peer systems to dating websites, ride sharing, and so on. The mathematical settings of such models are the following: there is a set C of customer classes and a set S of server classes, and a bipartite graph on the bipartition $C \cup S$ indicates the compatibilities: a customer class $c \in C$ and a server class $s \in S$ are compatible if and only if there is an edge between nodes c and s in the compatibility graph. Customer/server couples enter the system in discrete time (or are read one by one from left to right, in the prior models), and the classes of the couples in the order of arrivals (respectively, read from left to right) are independent and identically distributed (i.i.d.) from a measure μ on $C \times S$. In Caldentey et al. (2009), Adan and Weiss (2012), and Adan et al. (2017), it is assumed that $\mu := \mu_C \otimes \mu_S$ on $C \times S$; that is, the class of the incoming customer and server are independent of one another, whereas this assumption is relaxed in Bušić et al. (2013). In case of multiple compatibilities, it is the role of the *matching policy* to determine what match is performed. In Adan et al. (2017), it is shown that, under a natural stability condition, the stationary probability of the system under first

come, first served has a remarkable product form. Interestingly, this result can then be adapted to various skill-based queueing models as well and in particular those applying (various declinations of) the so-called first come first matched (FCFM)-assign the longest idling server (ALIS) service discipline (Adan and Weiss 2014). For various extensions of bipartite matching models, see also Adan et al. (2018). In Bušić et al. (2013), the stability condition in Adan et al. (2017) is shown to be also sufficient for the stability of extended bipartite models ruled by the “match the longest” policy, consisting of always matching incoming items to the compatible class having the longest queue. In Moyal et al. (2018), a coupling from the past result is obtained, showing the existence of a unique bi-infinite matching in various cases of extended bipartite matching models and for a broader class of matching policies than FCFM, thereby generalizing various results of Adan et al. (2017).

In many practical systems, it is often more realistic to assume that arrivals are simple instead of pairwise. Moreover, the bipartition of the compatibility graph into classes of *servers* and *customers* only suits applications in which this bipartition is natural (donor/receiver, house/applicant, job/applicant, etc.). However, in many cases, the context requires that the compatibility graph take a general (i.e., not necessarily bipartite) form. For instance, in dating websites, it is a priori not possible to split items into two sets of classes with no possible matches within those sets. Similarly, in kidney exchange programs, intraincompatible donor/receiver couples enter the system looking for a compatible couple to perform a crossed transplant. Then, it is convenient to represent donor/receiver couples as *single* items, and compatibility between couples means that a kidney exchange can be performed between the two couples (the donor of the first couple can give to the receiver of the second, and the donor of the second can give to the receiver of the first). In particular, if one considers blood types as a primary compatibility criterion, the compatibility graph between couples is naturally nonbipartite. Motivated by such applications, among others, Mairesse and Moyal (2016) proposed a generalization of the previous matching models, termed *general stochastic matching model*, in which items enter one by one, and the compatibility graph G is general. For a given compatibility graph G and a given matching policy Φ , the *stability region* $\text{STAB}(G, \Phi)$ of this system is defined as the set of those measures μ deciding the class of the incoming items, such that the Markov chain describing the evolution of the system is positive recurrent. Mairesse and Moyal (2016) identified a universal necessary condition for stability, that is, a set $\text{NCOND}(G)$ that includes $\text{STAB}(G, \Phi)$ for any matching policy Φ and any graph G . Following Bušić et al. (2013), Mairesse and Moyal (2016) also showed that match the longest is *maximal stable*; that is, it has a maximal stability region, as the two sets $\text{NCOND}(G)$ and $\text{STAB}(G, \Phi)$ coincide. Moyal and Perry (2017) then showed that, aside for a particular class of graphs, the random class-uniform policy (choosing the match in a compatible class of the incoming item, uniformly at random among those having a nonempty queue) is never maximal stable and that priority policies are in general *not* maximal stable. Then, Moyal et al. (2020) showed that the policy first come, first matched is maximal stable and that, similarly to Adan et al. (2017), the stationary probability of the system can be written in a product form. Various extensions of the general stochastic matching model can also be found in Büke and Chen (2015, 2017) regarding systems with bipartite complete compatibility graphs but in which each match can be enacted with a nominal probability, and in Rahmé and Moyal (2019) to stochastic matching systems on hypergraphs, thereby matching items by groups of two or more. In a recent line of study, stochastic matching models have also been studied from the point of view of stochastic optimization (Gurvich and Ward 2014, Bušić and Meyn 2016, Nazari and Stolyar 2019). In particular, the max-weight policy introduced in Tassioulas and Ephremides (1992) is a natural policy allowing a natural tradeoff between efficiency and stability.

Here we consider a matching system where items arrive one by one at each node of a compatibility network according to Poisson processes and depart from it as soon as they are matched to a compatible item. The matching policy considered is a generalized max-weight policy where decisions can be subject to noise. Additionally, some of the nodes may have impatience, that is, leave the system before being matched.

1.1. Motivation

In many applications, renegeing of items needs to be taken into account: for instance, applicants may renege from the systems, whereas jobs/houses/university spots may no longer be available. Similarly, in organ transplants, patients may renege from the system because of death or change in their diagnosis, whereas organs have a short life duration outside a body and need to be transplanted within a very short period of time—see a first approach of such systems in Boxma et al. (2011).

The present work is dedicated to an extension of general stochastic matching models to the case where (i) classes of items can be impatient, in that they renege from the system if they do not find a match within a given (random) period of time, and (ii) the matching policy is of the max-weight class, with possible errors in the assessment of the system state.

Specifically, we consider that the policy decisions may be subject to noise that can lead to errors in the estimate of the queue lengths, which is necessary to implement matching policies of the max-weight type.

Indeed, as is well known from the vast literature on the load balancing of parallel service systems, implementing policies that are sensitive to the queue sizes, such as the max-weight type policies considered in the present paper, becomes algorithmically prohibitive as the number of queues grows large. Efficiently sorting the queue sizes may easily lead to inaccuracies, entailing routing errors.

Specifically, in many practical cases the queue sizes of the various servers are not always known: they are either estimated by the incoming customers (as in physical queues in supermarkets, for instance), or by a central dispatcher (e.g., in passport-checking stations). Likewise, information regarding the queues to each of the servers in large web-server farms or call centers require constant communication between the job dispatchers and the servers, slowing down the response time. This information is thus not always available in real time (Lu et al. 2011), and in such cases, it can then be preferable to *estimate* the queue sizes, based on past or partial information, for instance. Observe that this source of inaccuracies is the main motivation for introducing “Power of d ”-type policies to reduce the complexity of routing and the probability of errors by routing the requests only to a subset of the available servers having a more tractable size (see Mitzenmacher (1996) and related references on this matter).

In the context of matching systems, in most of the aforementioned applications, the requests of the various classes are physically gathered in different queues; that is, each node of the graph has its own buffer for storing the corresponding items. In many cases, the size of the graph, that is, the number of classes, is large, leading to routing problems that are similar to the context of parallel service systems. For instance, liver transplants are subject to compatibility constraints that lie far beyond the blood types of givers and receivers, namely, age, the nature of the pathology of the patient (cancer or cirrhosis), and immunological criteria. This leads to an increase of the number of classes of givers and receivers, and in turn, to an increase of the complexity of implementing max-weight type policies and to the potentiality of allocation errors or inaccuracies.

To account for this difficulty, in the spirit of Moyal and Perry (2021) for parallel service systems, we introduce an error parameter in the estimate of the queue sizes in the context of matching systems. In this context, it is then reasonable to consider that these errors are independent of one another but follow distributions that may depend on the specificity of the considered nodes. For instance, an arriving item to node j chooses a match from the neighboring node i , making an error of estimate of the queue length of node i , that may depend on i and j ; see our later assumptions.

All in all, by “generalized max-weight policies,” we mean policies that are sensitive to the queue sizes, that incorporate rewards to give specific weights to the various matchings, and that take into account the possibility of errors in assessing the queue sizes of the various nodes of the graph.

1.2. Main Contributions

As we developed previously, the extension of the general stochastic matching model introduced in Mairesse and Moyal (2016), to the paradigm of reneging systems, is prevalent in practice. As will be specified later, adding an impatience parameter (possibly to some, but not all of the item classes) makes the analysis more intricate than in the original model. This constitutes the first main contribution of this work.

Using the specific structure of the max-weight policy, our second main contribution, Theorem 1, is to show that a simple quadratic function is a Lyapunov function whenever the traffic condition “ $\lambda \in \text{NCOND}(G, \gamma)$ ” (see (4)) that can be understood as a generalization of the classical condition “ $\lambda \in \text{NCOND}(G)$ ” to the case of reneging, is satisfied. This result is interesting and original for its own sake, even in the no-reneging case. This allows us directly to characterize the stability region and, as a corollary, to conclude that the max-weight policy is maximum stable. By doing so, we widely generalize results in Mairesse and Moyal (2016) to the case of generalized max-weight policies (of which the match the longest (ML) policy introduced in Bušić et al. (2013) and then Mairesse and Moyal (2016) is a particular case, see the later discussion) and to matching models with reneging. This quadratic function was already identified in Mairesse and Moyal (2016) and Bušić et al. (2013) as a Lyapunov function by an indirect proof: First, by stating that if the Foster-Lyapunov criterion applies to a quadratic Lyapunov function for some matching policy, then it must be also the case for the policy ML, and second, by showing that the criterion does apply to an ad hoc randomized policy. However, this argument cannot be generalized to the much wider case of max-weight policies, whence the main interest of the direct proof that is presented here.

We say that a couple graph/reneging vector (G, γ) is *stabilizable* whenever there exists a matching policy and an arrival vector λ making the corresponding matching model stable. As a consequence of Theorem 1, (G, γ) is stabilizable if and only if the set $\text{NCOND}(G, \gamma)$ is nonempty. It is well known (theorem 1 in Mairesse and Moyal 2016) that for a general stochastic matching model without reneging, the graph G is stabilizable if and only if it is *not* bipartite, because this condition is equivalent to $\text{NCOND}(G, \gamma)$ being nonempty. Our third main contribution is a quite remarkable result: In Proposition 2, we show a very simple sufficient condition for the stabilizability of

the system: it is sufficient that G be nonbipartite, or that only *one* class of elements renege. In particular, for a bipartite G (as is the case for organ transplants model, for instance), the system is stabilizable if, and only if at least one class of elements renege. Said differently, whenever G is bipartite and arrivals are made by single items, it is enough that one node receives impatient elements to make the system stabilizable, rather than artificially assuming that arrivals are made by pairs (as is the case in the (extended) bipartite matching models studied in the literature).

Showing that a given (even simple) function is a Lyapunov function is in general a challenging issue, and our case is not an exception, especially in such general a context. Our proof is quite intricate and relies on carefully exploiting this generalized NCOND condition. Combining these ideas, as our fourth main contribution, we are also able to prove exponential convergence speed toward the stationary measure, and additionally, by combining with stochastic comparisons, to construct bounds for the stationary mean of the largest queue size of customers in the system.

We finally illustrate some of those results using simulations on toy examples.

1.3. Outline of the Paper

In Section 2, we give the formal definition of the model we study. Section 3 is dedicated to the theoretical results. Section 3.1 focuses on stability: we characterize the stability region of the models in Theorem 1; this theorem is based on Proposition 1, a key result in which we control the drift of a quadratic function. Last, Proposition 2 describes precisely whether a given model is stabilizable. In Section 3.2, we then investigate the speed of convergence to equilibrium: a geometric Lyapunov bound is given in Proposition 3, and the exponential convergence to equilibrium follows as a consequence in Theorem 2. In Section 3.3, some bounds are given for moments of the size of the largest queue of the system in equilibrium in Theorem 3. Section 4 is dedicated to simulations: we compare the max-weight policy with a priority policy in Section 4.1, and an analysis of the theoretical bounds obtained in the previous section is given in Section 4.2. All proofs are provided in Section 5.

2. Model

2.1. Notation

The graph $G = (V, E)$ is assumed in all the article to be finite, simple, undirected, and connected. Because the set of nodes V represents customers classes, we use the words node and class in an indistinguishable way during all the article. The nonconnected case can be addressed by decomposing into connected components. Furthermore, to avoid trivialities, we assume $|V| > 1$. We write $\{i, j\} \in E$ to denote that there is an edge connecting the nodes $i, j \in V$ in the graph G . We say that a subset $I \subset V$ is an independent set if $i, j \in I$ implies $\{i, j\} \notin E$. For any subset $W \subset V$, we let

$$\mathbb{I}(W) = \{I \subset W : I \text{ is an independent set}\} \setminus \{\emptyset\}$$

be the family of nonempty independent subsets of W , in a way that $\mathbb{I} := \mathbb{I}(V)$ gathers all nonempty independent sets of the graph G . Let also $E(W) = \{i \in V : \{i, j\} \in E \text{ for some } j \in W\}$. For singletons, we write $E(i)$ for $E(\{i\})$. Observe that $W \cap E(W)$ is the union of the connected components of W that are not singletons; in particular, $W \cap E(W) = \emptyset$ if and only if W is an independent set. Let $\mathbb{R}, \mathbb{R}_+, \mathbb{R}_+^*, \mathbb{Z}, \mathbb{N}_0$ and \mathbb{N} , respectively, denote the sets of real numbers, non-negative real numbers, positive real numbers, integers, nonnegative integers, and positive integers. For $p, q \in \mathbb{Z}$, we let $\llbracket p, q \rrbracket := [p, q] \cap \mathbb{Z}$ be the set of all integers in the closed interval $[p, q]$. Let also, for $z \in \mathbb{R}^V$ and $W \subset V$, $\|z\|_W = \max_{i \in W} |z(i)|$, writing simply $\|z\|$ in the case $W = V$ (this is the supremum norm), and letting $\|z\|_\emptyset = 0$.

2.2. Stochastic Matching Model with Reneging

The (continuous-time) stochastic matching model with renegeing (G, MW, λ, γ) is formally defined as follows: we give ourselves a set of $|V|$ independent Poisson processes $N(i)$, $i \in V$, of respective intensities $\lambda(i) > 0$, $i \in V$, and denote by $\lambda := (\lambda(i), i \in V)$ the *arrival intensity vector* of the model. For any $i \in V$, we denote by $\{T_n^i : n \in \mathbb{N}\}$ the sequence of successive arrival times in the process $N(i)$. For any subset $W \subset V$, denote by $\lambda(W) = \sum_{i \in W} \lambda(i)$ the intensity of the superposed Poisson processes having indexes in W . The overall superposition of the Poisson process of the $N(i)$ s, $i \in V$, is denoted by N . It is thus a Poisson process of intensity $\lambda(V)$, and its points are denoted by $\{T_n : n \in \mathbb{N}\} := \bigcup_{i \in V} \{T_n^i : n \in \mathbb{N}\}$. For all i , on each point T_n^i , $n \in \mathbb{N}$, an *item of class i* (or *i -item* for short) enters the system. Thus, the points of N are the overall arrival times in the system. We say that two items of respective classes i and j in V are *compatible* if and only if $\{i, j\} \in E$. To each class $i \in V$ is also associated a parameter $\gamma(i) \geq 0$. Denote $\gamma := (\gamma(i), i \in V)$, the *reneging rate vector* of the model. Unlike the vector λ , we allow γ to have null entries. We then associate to each i -item an exponential random variable (r.v.) of distribution $\mathcal{E}(\gamma(i))$, interpreted as its *patience time*. An r.v. of distribution $\mathcal{E}(0)$ is interpreted as almost surely (a.s.) infinite, and we denote by

$$R := \{i \in V : \gamma(i) > 0\} \quad \text{and} \quad R^c := \{i \in V : \gamma(i) = 0\}$$

the subsets of nodes whose items do (respectively, do not) renege from the system. The corresponding item then reneges from the system at the end of this time period provided that it has not left before because of the matching mechanism we describe later. We assume that the patience times of all items are independent of one another and of the other involved r.v.s.

The dynamics of the model (G, MW, λ, γ) is then defined recursively at the successive points of N , as follows:

- We start at time 0 with a state in which all present items are incompatible. In other words, in the system at time 0, for any $\{i, j\} \in E$, there are either no i items, or no j items, or neither. Such a buffer content is called *admissible*. We associate a patience to each of these initial items that is drawn, independently of everything else, from the distribution corresponding to its class, for example, from $\mathcal{E}(\gamma(j))$ for a j -item. As previously mentioned, if $j \in R^c$, that is, $\gamma(j) = 0$, the patience time of the corresponding item is set as infinite.

- Let $n \in \mathbb{N}$, and suppose that the buffer content of the system was admissible at time T_{n-1} . We first update the buffer content at each reneging time in $[T_{n-1}, T_n)$, by discarding all items present in the system at time T_{n-1} and whose patience time has expired during that time interval. Observe that the buffer content remains admissible after this procedure. Then, suppose that T_n is a point of the j th arrival process, that is, $T_n = T_k^j$ for some $k \in \llbracket 1, n \rrbracket$, or in other words, a j -item enters the system at T_n . This happens with probability

$$\mu_\lambda(j) := \frac{\lambda(j)}{\lambda(V)}. \quad (1)$$

We then face the following alternatives:

(a) Either the buffer contains no items compatible with the incoming j -item, that is, there are no stored items of classes in $E(j)$, in which case the latter item is stored in the buffer. We then draw the (possibly infinite) patience time of this item from the distribution $\mathcal{E}(\gamma(j))$, independently of everything else.

(b) Or there is at least one remaining item having class in $E(j)$ in the buffer. Then the incoming j -item must be matched with one of these compatible items, and it is the role of the matching policy to determine its match. We assume that the matching policy is of the *generalized max-weight policy* type with parameters w and F ($MW_{w,F}$, or simply MW , for short). To define this, let $w := \{w_{(j,i)} : (j,i) \in V^2 \text{ s.t. } \{i,j\} \in E\}$ be a family of nonnegative numbers and $\{F_{(j,i)} : (j,i) \in V^2 \text{ s.t. } \{i,j\} \in E\}$ be a family of cumulative distribution functions (c.d.f.s). We index with ordered pairs (j,i) instead of the unordered ones because the involved parameters are not assumed to be symmetric. The first entry of the pair j will typically represent the class of the incoming item, and the second one i , the class of the item j performs a match with. Denote, for all $i \in V$, by $x(i)$ the number of stored i -items at this time. Then the match of the incoming j -item is chosen uniformly at random from the set

$$\arg \max_{i \in E(j) : x(i) > 0} \left([x(i) + U_{(j,i),n}]^+ + w_{(j,i)} \right), \quad (2)$$

where $U_{(j,i),n}$ is an r.v. drawn at time n independently of everything else from the c.d.f. $F_{(j,i)}$, interpreted as the measurement error of the queue length of node i , made on an arrival of a j element at time T_n , and $w_{(j,i)}$ is interpreted as the specific reward of matching an incoming j -item with an i -item. Observe that in (2), we make the simplifying assumption that measurement errors are made only regarding the length of *nonempty* queues; that is, empty queues are identified correctly. This makes sense in practice, first, because the overestimation of the size of an empty queue might lead to a matching involving an item that is not actually present in the system. Second, it is usually simpler to determine whether a given buffer is empty or not (by sending a ping to the corresponding server) than to precisely measure the size of a (potentially much longer) nonempty queue.

Then, once the matching is done, the incoming j -item leaves the system right away, together with its match. We emphasize that, unlike the notion of compatibility given by the undirected set of edges E , symmetry between i and j is a priori not assumed in the definitions of the measurement errors and the matching rewards.

It is then easily seen that, in any case, the buffer content at time T_n is also admissible.

- We continue the construction inductively.

Remark 1. Particular cases of generalized max-weight policies are as follows:

1. The classical max-weight policies, as defined in Tassiulas and Ephremides (1992), recovered by taking constant weights and, for every $(j,i) \in V^2$ such that $\{i,j\} \in E$, the c.d.f. $F_{(j,i)} := \mathbb{1}_{\mathbb{R}_+}$ as the one corresponding to the Dirac δ_0 distribution, that is, no error.

2. The match the longest policy introduced in Mairesse and Moyal (2016), recovered by taking $F_{(j,i)} := \mathbb{1}_{\mathbb{R}_+}$ and $w_{(j,i)} = 0$ for all $(j,i) \in V^2$ such that $\{i,j\} \in E$, that is, no measurement errors and no rewards.

Remark 2. The error measurement model we consider is additive, namely, a scale-free error is made on all queue lengths, and we prevent absurd measurement of negative queue lengths by taking the positive part of the latter (see (2)). Another interesting model would be to assume multiplicative error measurements; namely, (2) would read

$$\arg \max_{i \in E(j) : x(i) > 0} (x(i)U_{(j,i),n} + w_{(j,i)}).$$

We strongly believe that all results to come also hold in that case but leave the proof of this conjecture to the reader.

2.3. Markov Representation

Let, for all $t \geq 0$, $X_t := (X_t(i), i \in V)$ be the queue length vector at time t , namely, for any $i \in V$, $X_t(i)$ is the number of i -items in the buffer at time t . It is easily seen that the process $X = \{X_t : t \geq 0\}$ is a regular continuous time Markov chain (CTMC) on the (countable) commutative state-space

$$\mathcal{X} = \{x \in \mathbb{N}_0^V : x(i)x(j) = 0 \text{ for any } \{i, j\} \in E\}.$$

For any $x \in \mathcal{X}$, let $S_x = \{i \in V : x(i) > 0\}$ be its support. Observe that S_x is an independent set for every $x \in \mathcal{X}$. For all $i \in V$, denote by δ_i the i th canonical vector: $\delta_i(j) = \mathbf{1}\{i = j\}$. The infinitesimal generator \mathcal{L} of X is then simply given by

$$\mathcal{L}f(x) = \sum_{i \notin E(S_x)} [f(x + \delta_i) - f(x)]\lambda(i) + \sum_{i \in S_x} [f(x - \delta_i) - f(x)] \left[\gamma(i)x(i) + \sum_{j \in E(i)} \lambda(j)v_{x,j}(i) \right] \quad (3)$$

for any $x \in \mathcal{X}$ and any function $f : \mathcal{X} \rightarrow \mathbb{R}$, where, for all $i \in V$ and $j \in E(i)$, $v_{x,j}(i)$ is the probability of performing the match $\{i, j\}$, given that a j -item enters the system in state x . From the very definition of MW, the latter is fully determined by the c.d.f.s $F_{(j,i)}$, $i \in S_x \cap E(j)$.

3. Results

3.1. Stability

The state $x \in \mathcal{X}$ is said to be *positive recurrent* if the expected amount of time until the CTMC revisits x given that $X_0 = x$, is finite. The CTMC is said to be positive recurrent if every state $x \in \mathcal{X}$ is. Because the CTMC is irreducible, this occurs if and only if some state is positive recurrent. Positive recurrence is equivalent to the existence and uniqueness of an invariant distribution. We say that the matching model $(G, \text{MW}, \lambda, \gamma)$ is stable if the CTMC $(X_t)_{t \geq 0}$ is positive recurrent. The next result provides a necessary and sufficient condition for stability.

Theorem 1. For any connected graph G and any $\gamma \in \mathbb{R}_+^V$, the model $(G, \text{MW}_{w,F}, \lambda, \gamma)$ is stable if and only if its intensity vector λ belongs to the set

$$\text{NCOND}(G, \gamma) = \left\{ \lambda \in (\mathbb{R}_+^*)^V : \lambda(I) < \lambda(E(I)) \text{ for any } I \in \mathbb{I}(R^c) \right\}. \quad (4)$$

Observe that if we define

$$\eta = \eta(\lambda) := \begin{cases} \min\{\lambda(E(I)) - \lambda(I) : I \in \mathbb{I}(R^c)\} & \text{if } R^c \neq \emptyset \\ \infty & \text{else,} \end{cases}$$

then $\lambda \in \text{NCOND}(G, \gamma)$ if and only if $\eta(\lambda) > 0$.

For $(j, i) \in V^2$ such that $\{i, j\} \in E$, let $U_{(j,i)}$ be an r.v. with c.d.f. given by $F_{(j,i)}$; that is, $U_{(j,i)}$ has the same distribution than $U_{(j,i),n}$ for every n . For $\kappa \in (0, \lambda(V))$, define

$$u_{\kappa,(j,i)} = \min \left\{ u \in \mathbb{R}_+ : \mathbb{P}(|U_{(j,i)}| \leq u) \geq \left(1 - \frac{\kappa}{\lambda(V)}\right)^{\frac{1}{|E|}} \right\},$$

and let

$$u_\kappa = \max\{u_{\kappa,(j,i)} : (j, i) \in V^2 \text{ s.t. } \{i, j\} \in E\}.$$

The quantity u_κ is tailor-made for the event

$$\mathcal{B}_{\kappa,n} := \bigcap_{(j,i) \in V^2 : \{i,j\} \in E} \{|U_{(j,i),n}| \leq u_\kappa\}$$

to satisfy

$$\mathbb{P}(\mathcal{B}_{\kappa,n}) \geq 1 - \frac{\kappa}{\lambda(V)} \quad \text{for every } n \in \mathbb{N}. \quad (5)$$

Theorem 1 is a corollary of the following key generator bound for a quadratic function.

Proposition 1. Assume $\lambda \in \text{NCOND}(G, \gamma)$ and let $\kappa \in (0, \lambda(V))$. Let $f_2: \mathcal{X} \rightarrow \mathbb{R}$ be the quadratic function defined by $f_2(x) = \sum_{i \in V} x(i)^2$, and let $\check{w} = \max_{(j,i) \in V^2: \{i,j\} \in E} w_{(j,i)}$.

1. If $R = \emptyset$, then

$$\mathcal{L}f_2(x) \leq \lambda(V) + 2\lambda(V)(2u_\kappa + \check{w})|V| + 2(\kappa - \eta)\|x\| \quad (6)$$

for every $x \in \mathcal{X}$.

2. Otherwise, if $R \neq \emptyset$, we have

$$\begin{aligned} \mathcal{L}f_2(x) &\leq \lambda(V) + \sum_{i \in R} \left[-2\gamma(i)x(i)^2 + (\gamma(i) + 2\lambda(i))x(i) \right] \\ &\quad + 2[\|x\|_R + (2u_\kappa + \check{w})|R^c|]\lambda(R^c) \\ &\quad + 2[\lambda(V)(2u_\kappa + \check{w})|R^c| + (\kappa - \eta)\|x\|_{R^c}] \mathbb{1}\{\|x\|_{R^c} > \|x\|_R + (2u_\kappa + \check{w})|R^c|\} \end{aligned} \quad (7)$$

for every $x \in \mathcal{X}$.

As can be easily seen in the proof of the necessity statement in Theorem 1, for a connected graph G and a reneging rate vector γ , λ being an element of $\text{NCOND}(G, \gamma)$ is necessary for the existence of a matching policy that makes the system stable, whatever the matching policy is. The sufficiency statement means that such models always exist whenever λ do belong to $\text{NCOND}(G, \gamma)$ —It suffices to implement a matching policy in the max-weight class. Therefore, it is natural to say that the couple (G, γ) is *stabilizable* if there exists an intensity vector λ such that the model with reneging $(G, \text{MW}, \lambda, \gamma)$ is stable, that is, if $\text{NCOND}(G, \gamma)$ is nonempty. The following proposition makes precise the class of stabilizable graphs.

Proposition 2. Let $G = (V, E)$ be a connected graph and $\gamma \in \mathbb{R}_+^V$ be a reneging rate vector. Then $\text{NCOND}(G, \gamma)$ is nonempty if and only if G is nonbipartite or $R \neq \emptyset$.

In other words, if G is bipartite, it is enough that one class of elements renege to make the system stabilizable. As will be made clearer in the proof of Proposition 2, the reason behind this (apparently surprising) result is as follows: if $G = (V, E)$ is bipartite, say of bipartition $V = V_1 \cup V_2$, and no elements renege, then the *only* independent set I that cannot satisfy $\lambda(I) < \lambda(E(I))$ for some intensity vector λ is either $I = V_1$ or $I = V_2$ because, as G is bipartite and connected, $\{V_1, V_2\}$ is the only pair of sets such that $V_2 = E(V_1)$ and $V_1 = E(V_2)$, and we cannot have $\lambda(V_1) < \lambda(V_2)$ together with $\lambda(V_2) < \lambda(V_1)$. Then, adding reneging to a single node of V_1 (if $I = V_1$), or to V_2 (if $I = V_1$) allows to simply break this paradoxical symmetrical condition. See the details in Section 5.3.

To summarize, gathering Theorem 1 and Proposition 2, we obtain the following panorama regarding the stability of the model:

1. Any couple (G, γ) is stabilizable provided that G is nonbipartite or that some classes of elements renege, that is, $R \neq \emptyset$.
2. Under these conditions, any model $(G, \text{MW}, \lambda, \gamma)$ with $\lambda \in \text{NCOND}(G, \gamma)$, is stable.

Remark 3 (Full Reneging Case). In the case where $R^c = \emptyset$, that is, all items may renege, Condition (4) is empty, and we retrieve the good-sense result that any such matching model is stable, whatever the intensity vector λ .

Remark 4 (No-Reneging Case). In the case where $R = \emptyset$ (no reneging), the model amounts to a continuous-time version of that introduced in Mairesse and Moyal (2016). From Theorem 1, any model $(G, \text{MW}, \lambda, \gamma)$ is stable whenever λ belongs to the set

$$\text{NCOND}(G) = \left\{ \lambda \in (\mathbb{R}_+^*)^V : \lambda(I) < \lambda(E(I)), \text{ for all } I \in \mathbb{I} \right\}. \quad (8)$$

In other words, any policy of the MW type is then maximal stable, in the sense that it has the largest stability region. This result generalizes those of Mairesse and Moyal (2016). Indeed, setting $F_{(j,i)} := \mathbb{1}_{\mathbb{R}_+}$ for all (j, i) (no errors) and $w_{(j,i)} := w_j$ for all $j \in V$ and all $i \in E(j)$ (equal rewards for all matches involving an incoming j -item, $j \in V$), the matching policy amounts to match the longest (priority is given to the node having the largest queue), and we retrieve assertion (15) in theorem 2 of Mairesse and Moyal (2016). Theorem 1 is thus a wide generalization of the latter result, to the case of weighted MW policies with measurement errors.

3.2. Speed of Convergence to Equilibrium

In this and the following section, we assume $R = \emptyset$ for simplicity. The involved techniques can be used to obtain analogous results in the general case (see Remark 5).

Building on the previous results and following classical approaches for Markov processes with bounded jumps, we now show how to define a geometric Lyapunov function that is very useful to quantify the speed of convergence to the steady state. We first have the following result.

Proposition 3. *Let $(G, MW_{w,F}, \lambda, \gamma)$ be a max-weight stochastic matching model with $\gamma \equiv 0$ and λ belonging to the set $\text{NCOND}(G, \gamma)$ defined by (4). Then, if $\alpha > 0$ is small enough, there exist $c, k > 0$ such that*

$$\mathcal{L}e^{\alpha\|x\|_2} \leq -ce^{\alpha\|x\|_2}$$

for any $x \in \mathcal{X}$ such that $\|x\|_2 \geq k$, being $\|x\|_2 = \left(\sum_{i \in V} x(i)^2\right)^{1/2}$ and \mathcal{L} is the generator of the process X defined by (3).

The previous proposition implies that the distribution of the Markov process X converges to its stationary distribution exponentially fast. More precisely, the following result follows immediately from Proposition 3, using, for example, the results in Hairer (2014).

Theorem 2. *As in Proposition 3, let $(G, MW_{w,F}, \lambda, \gamma)$ be a max-weight stochastic matching model with $\gamma \equiv 0$ and λ belonging to $\text{NCOND}(G, \gamma)$, and let π be the unique invariant distribution of the process X (granted by Theorem 1). For any $\alpha > 0$ there exist constants $c > 0$ and $0 < \rho < 1$ such that*

$$\|\mathbb{P}_x(X_t = \cdot) - \pi(\cdot)\|_{\text{TV}} \leq c\rho^t e^{\alpha\|x\|}$$

for all $x \in \mathcal{X}$. Here \mathbb{P}_x is the law of the process conditioned to $[X_0 = x]$, and $\|\cdot\|_{\text{TV}}$ denotes the total variation distance.

3.3. Moments Bounds at Equilibrium

We control below the first moment of the maximum queue size under the steady state of the system.

Theorem 3. *Assume that we are under hypotheses of Proposition 3, namely $R = \emptyset$ and $\lambda \in \text{NCOND}(G, \gamma)$, and let π be the unique invariant distribution. Let $U_{(j,i)}$ have distribution $F_{(j,i)}$ (this r.v. has been defined in Section 3.1), and assume that there exists $B \geq 0$ such that $\mathbb{P}[-B \leq U_{(j,i)} \leq B] = 1$ for every $(j, i) \in V^2$ such that $\{i, j\} \in E$. Then we have that*

$$\max_{i \in V} \frac{\rho_i}{1 - \rho_i} \leq \int \|x\|_{\infty} \pi(dx) \leq \frac{\lambda(V)}{\eta} \left(\frac{1}{2} + (\tilde{w} + 2B)|V| \right), \quad (9)$$

where $\rho_i = \lambda(i)/\lambda(E(i))$ for all $i \in V$.

Remark 5. *With exactly the same procedures, the content of Sections 3.2 and 3.3 can be extended in the following two directions: on the one hand, one may weaken hypothesis $R = \emptyset$; on the other, bounds for higher order moments under π may be obtained. We prefer to avoid developing this direction since the computations become rapidly cumbersome.*

4. Simulations

In this section, we present a few simulation results. In the first set of simulations, we illustrate the gain of performance of max-weight policies with respect to other, more static, policies, for systems with reneging. To do so, as an illustrative example, in Section 4.1, a max-weight policy is compared with policies of the *priority* type, namely, policies that only take into consideration the rewards of the various matches but not the queue sizes. Then, in Section 4.2, we compare the theoretical bounds obtained by Lyapunov techniques in Section 3.3 with simulated values. These bounds are not meant to be very tight, and we leave their optimization for future research.

4.1. Comparison with Priority Policies

In this section, we assume $F_{(j,i)} = \mathbb{1}_{\mathbb{R}_+}$ for every $i, j \in V$ such that $\{i, j\} \in E$, namely no measurement errors are considered. A policy is said to be of the *priority* type if, whenever the system is in state $x \in \mathcal{X}$ and a j -item enters the system, the match of the incoming j -item (if any) is chosen uniformly at random from the set

$$\arg \max_{i \in E(j): x(i) > 0} w_{(j,i)}.$$

This amounts to saying that each j -item has a fixed order of priorities between the neighboring nodes of i , for choosing its match. (Observe that a priority policy can be seen as a max-weight policy where we set the errors to $U_{(j,i),n} = -\infty$ a.s. for all $n, (j, i)$.) As it is illustrated in the two examples in section 5 of Mairesse and Moyal (2016), priority policies can be maximal stable or not. Conversely, as is shown in theorem 3 of Moyal and Perry (2017), for any graph G outside a particular class of graphs, there always exists a nonmaximal priority policy.

To compare the asymptotic behavior of max-weight and priority policies for a variety of graphs, arrival and departure rates, and rewards, we randomize all these parameters. We perform 100 simulations. In each one, we first sample the graph G , the arrival rates λ , departure rates γ , and rewards w , as follows:

1. The graph $G = (V, E)$ is an Erdős-Rényi one with parameters $|V| = 30$ and $p = 0.1$;
 2. The arrival rates $(\lambda(i))_{i \in V}$ are i.i.d. with common distribution $\text{Unif}[0, 10]$;
 3. The reneging rates are i.i.d. with common distribution $\text{Ber}(0.5)$ (on average, half of the nodes do not renege);
- and
4. We consider symmetric rewards, namely $w_{(j,i)} = w_{(i,j)}$ for every i, j such that $\{i, j\} \in E$, and the family $\{w_{(j,i)} : \{i, j\} \in E\}$ is i.i.d. with common law $\text{Unif}[0, 10]$.

In each case, the condition NCOND is forced; that is, if λ is not an element of $\text{NCOND}(G)$, the simulation is rerun. Then, for each simulation we compare a system ruled by max-weight to the same system ruled by the corresponding priority policy, for the same sampled parameters. In other words, we perform paired replications in the sense that a random network is generated by items 1–4, and the same network is used for comparing the corresponding policies. Two level of randomness are involved: the first one to define the network (formed by the graph, and the values of the rates and rewards), and the second one to draw, conditioned to the already sampled network, the effective dynamic of the process (consisting on the arrival and departure episodes, and the randomness to decide which item class to pick in case of a tie in the value of the weights). Observe that, from Theorem 1, the max-weight system is necessarily stable, whereas in view of the previous observations, the system ruled by the priority policy is not necessarily so.

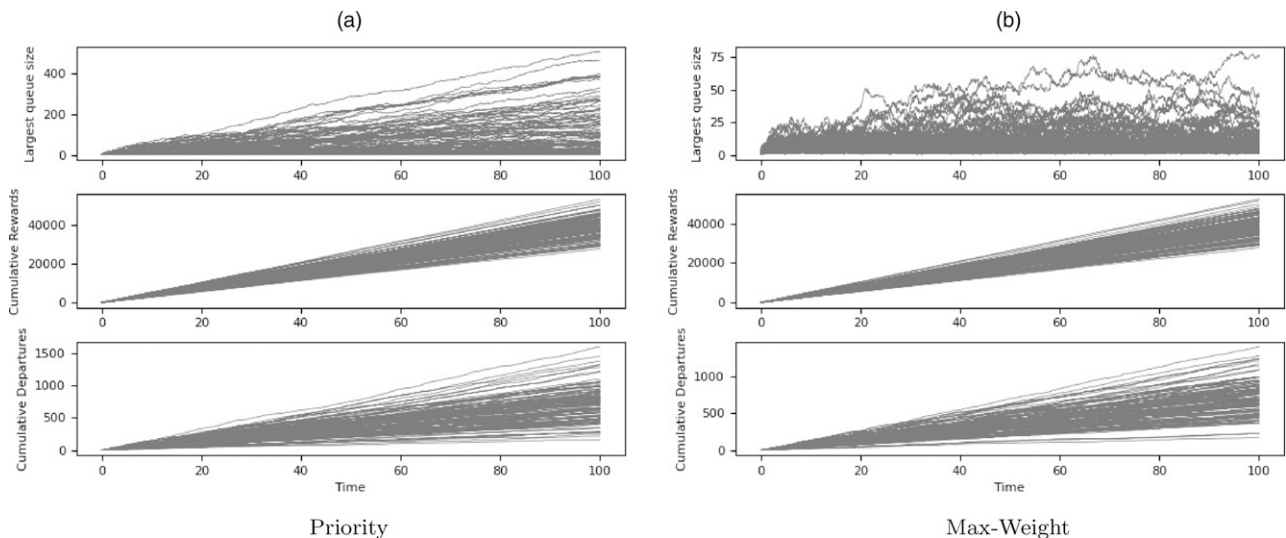
Figure 1 represents, in vertical order, the size of the largest queue, the cumulative reward (each time a (j, i) -match occurs, a reward $w_{j,i}$ accumulates) and the cumulative number of departures, the three quantities versus time, in the time interval $[0, 100]$. Figure 2 is the boxplot for the size of the largest queue at time $t = 100$.

These results put in evidence an important and striking qualitative difference between the two policies, in the behavior of the buffer size process. Specifically, max-weight keeps far fewer items in the system, whereas both policies have essentially the same performance in terms of cumulative reward and departures. In many cases, the priority policy seems to fail in stabilizing the system, whereas the corresponding max-weight system is always stable. Finally, it is worth mentioning that the same picture was obtained in unpublished simulations ran in much larger time intervals.

4.2. Bounds in the Stationary Regime

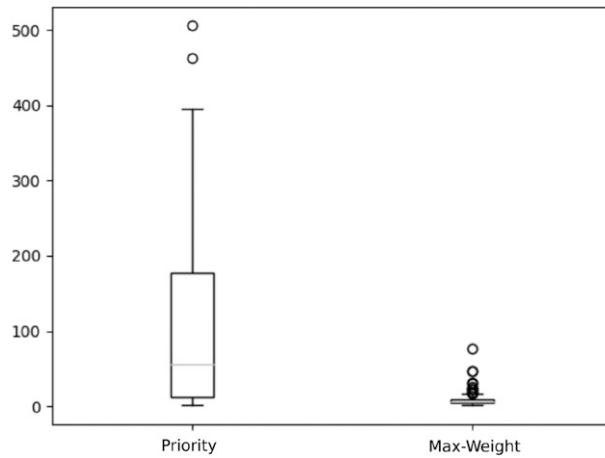
Unlike Section 4.1, in this section, impatience is excluded. We first restrict our attention to the classical match the longest policy; that is, we take $w_{j,i} = 0$ for every $i, j \in V$ such that $\{i, j\} \in E$, and we work without measurement errors. An important difference with respect to the previous subsection is that here the parameters (G, λ) are

Figure 1. From Top to Bottom, Largest Queue Size, Cumulative Reward, and Cumulative Number of Departures, in the Time Interval $[0, 100]$



Notes. (a) Priority. (b) Max-weight.

Figure 2. Boxplots for 100 Iterations at Time $t = 100$ of the Largest Queue Size



drawn only once at the beginning. We clarify that no comparison with a priority policy is made here. The distribution of the pair (G, λ) is maintained: Erdős-Rényi distribution with parameters $|V| = 30$ and $p = 0.1$ for G , and common distribution $\text{Unif}[0, 10]$ for the i.i.d. family $(\lambda(i))_{i \in V}$. Figure 3 represents a histogram at time $t = 100$ of 500 simulations. The vertical lines represent the lower and upper bounds for the expected size of the largest queue given by Theorem 3.

Finally, we study the case with measurement errors and rewards. The random variables $\{U_{(j,i)} : \{i, j\} \in E\}$ giving place to the noises are chosen to be i.i.d. with common distribution $\text{Unif}[-1, 1]$, and the rewards are again symmetric and i.i.d. with common distribution $\text{Unif}[0, 10]$. The distributions of G and λ do not change, and the tuple (G, λ, U, w) is drawn only once at the beginning. Figure 4 is a histogram also for 500 iterations at time $t = 100$, and the vertical line represents the lower bound for the expectation. The numerical value for the upper bound for the expectation is 26,867.96, and it is not included in the graphic.

For match the longest, we observe tight bounds for the expected value. All the upper bounds become very loose when we include the presence of rewards or noises. From unpublished simulations, we observe that this is a general picture and not only a particular situation for this choice of the parameters. In the presence of asymmetric rewards and noise, the bounds are not informative. We leave for future work to obtain better tailored made bounds in this case. A typical histogram is very close to the lower bound for the expectation, and the upper bound is more accurate if larger values of p in the distribution of G are considered.

Figure 3. Histogram for 500 Simulations at Time $t = 100$ of the Largest Queue Size for Match the Longest (i.e., MW with $w_{ij} = 0$)





Figure 4. Histogram for 500 Simulations at Time $t = 100$ of the Largest Queue Size for Max-Weight with Rewards and Noise

5. Proofs

5.1. Proof of Proposition 1

Fix $x \in \mathcal{X}$ and $\kappa \in (0, \lambda(V))$. To shorten notation, we write

$$\tau_i(x) = \sum_{j \in E(i)} \lambda(j) v_{x,j}(i).$$

We assume during the proof that we are in the $(n-1)$ th step, namely, the law of $v_{x,j}$ is ruled by the family of r.v.s $\{U_{(j,i),n} : \{i,j\} \in E\}$; of course, by homogeneity, this can be done without loss of generality. For the quadratic function f_2 , identity

$$f_2(x \pm \delta_i) - f_2(x) = 1 \pm 2x(i)$$

holds as long as $x \pm \delta_i \in \mathcal{X}$. Hence, in view of the expression defining the generator (3), we get

$$\mathcal{L}f_2(x) = \sum_{i \notin E(S_x)} [1 + 2x(i)]\lambda(i) + \sum_{i \in S_x} [1 - 2x(i)][\gamma(i)x(i) + \tau_i(x)]. \quad (10)$$

Because $x(j) = 0$ for all $j \in E(S_x)$, the first sum equals

$$\lambda(E(S_x)^c) + 2 \sum_{i \in V} \lambda(i)x(i).$$

Observe that

$$\begin{aligned} \sum_{i \in S_x} \tau_i(x) &= \sum_{i \in S_x} \sum_{j \in E(i)} \lambda(j) v_{x,j}(i) \\ &= \sum_{i \in V} \sum_{j \in V} \mathbb{1}\{i \in S_x\} \mathbb{1}\{\{i,j\} \in E\} \lambda(j) v_{x,j}(i) \\ &= \sum_{i \in V} \sum_{j \in V} \mathbb{1}\{i \in S_x\} \mathbb{1}\{j \in E(S_x)\} \lambda(j) v_{x,j}(i) \\ &= \sum_{i \in V} \sum_{j \in V} \mathbb{1}\{j \in E(S_x)\} \lambda(j) v_{x,j}(i) \\ &= \sum_{j \in E(S_x)} \lambda(j) \sum_{i \in V} v_{x,j}(i) \\ &= \sum_{j \in E(S_x)} \lambda(j) \\ &= \lambda(E(S_x)), \end{aligned} \quad (11)$$

the third and the fourth identities following because $v_{x,j}(i)$ is zero if $\{i,j\} \notin E$ or $i \notin S_x$. Hence, the second sum in (10) can be rewritten as

$$\sum_{i \in V} \gamma(i)x(i) + \lambda(E(S_x)) - 2 \sum_{i \in V} \gamma(i)x(i)^2 - 2 \sum_{i \in V} x(i)\tau_i(x).$$

Putting these observations together, we get

$$\mathcal{L}f_2(x) = \lambda(V) + \sum_{i \in V} [-2\gamma(i)x(i)^2 + x(i)\gamma(i)] + 2 \sum_{i \in V} x(i) [\lambda(i) - \tau_i(x)]. \quad (12)$$

From now on the proof is devoted to controlling the last term; an objective that will be carried out through the use of the following technical result.

Lemma 1. Fix $x \in \mathcal{X}$ and $\kappa \in (0, \eta]$, and assume that the nonempty subset $W \subset S_x$ satisfies the following hypotheses:

1. For every $j \in E(W)$, $\sum_{i \in W \cup x_j} \mathbb{1}_{\mathcal{B}_{\kappa, n}}(i) = 1$.
2. For every $A \subset W$, $\lambda(A) < \lambda(E(A))$.

Then

$$\sum_{i \in W} x(i) [\lambda(i) - \tau_i(x)] \leq \lambda(E(W))(2u_\kappa + \check{w})|W| + (\kappa - \eta)\|x\|_W. \quad (13)$$

In the assumptions of Lemma 1, hypothesis 1 states that, conditioned to $\mathcal{B}_{\kappa, n}$, an arriving j -item necessarily matches with an item whose class belongs to W . On another hand, hypothesis 2 states that the inequalities defining NCOND are satisfied inside the set W . The connection between the two hypotheses can be informally understood as follows: the second one gives stability inside W , and the first one tells you that this stability is not being taken advantage of by classes outside $W \cup E(W)$.

Before providing the proof of Lemma 1, we show how to conclude from it. To avoid trivialities, we assume $S_x \neq \emptyset$ from now on.

1. Assume $R = \emptyset$ (no reneging). It is easy to see that the hypotheses of Lemma 1 are satisfied for $W = S_x$. Hence, from (13) it follows that

$$\sum_{i \in V} x(i) [\lambda(i) - \tau_i(x)] \leq \lambda(V)(2u_\kappa + \check{w})|V| + (\kappa - \eta)\|x\|. \quad (14)$$

Then simply replace in (12) to obtain the desired inequality (6).

2. We now turn to the case where $R \neq \emptyset$. First, the result in the subcase $R^c = \emptyset$ follows immediately from (12). Therefore, let us finally assume that $R \neq \emptyset$ and $R^c \neq \emptyset$. Coming back to (12), we have

$$\begin{aligned} \mathcal{L}f_2(x) &\leq \lambda(V) + \sum_{i \in V} [-2\gamma(i)x(i)^2 + x(i)\gamma(i)] + 2 \sum_{i \in R} x(i)\lambda(i) + 2 \sum_{i \in R^c} x(i) [\lambda(i) - \tau_i(x)] \\ &= \lambda(V) + \sum_{i \in R} [-2\gamma(i)x(i)^2 + x(i)(\gamma(i) + 2\lambda(i))] + 2 \sum_{i \in R^c} x(i) [\lambda(i) - \tau_i(x)]. \end{aligned} \quad (15)$$

If $\|x\|_{R^c} \leq \|x\|_R + |R^c|(2u_\kappa + \check{w})$, we can easily conclude because

$$\sum_{i \in R^c} x(i) [\lambda(i) - \tau_i(x)] \leq \sum_{i \in R^c} x(i)\lambda(i) \leq [\|x\|_R + |R^c|(2u_\kappa + \check{w})]\lambda(R^c). \quad (16)$$

From now on we assume that $\|x\|_{R^c} > \|x\|_R + |R^c|(2u_\kappa + \check{w})$, and we recall that we are in the situation in which $R \neq \emptyset$ and $R^c \neq \emptyset$. One is tempted to apply Lemma 1 with $W = R^c \cap S_x$, especially because, for this particular choice of W , hypothesis 2 of this lemma is guaranteed by the hypothesis NCOND(G, γ). Unfortunately, hypothesis 1 of Lemma 1 fails. The rest of the proof is devoted to finding a proper subset $W \subset R^c \cap S_x$ fitting this missing requirement.

Here we introduce a crucial notion of connectivity that depends on the number of stored items, notion that is also used in the proof of Lemma 1. We define an auxiliary graph (V', E') by setting $V' = S_x \cap R^c$ and $\{i, j\} \in E' \iff i \neq j$ and $|x(i) - x(j)| \leq 2u_\kappa + \check{w}$. Let \mathcal{C} be the partition of V' induced by the connected components of the graph (V', E') : for every $C \in \mathcal{C}$, C is a subset of V' , and every pair $i, j \in C$ is connected by a path of edges in E' . The set

$$W = \bigcup \{C \in \mathcal{C} : x(i) > \|x\|_R + 2u_\kappa + \check{w} \text{ for every } i \in C\}$$

is the one that does the trick as we show here:

• We first see that $W \neq \emptyset$. Let $i_0 \in V'$ be such that $\|x\|_{R^c} = x(i_0)$, and let $C_0 \in \mathcal{C}$ be the subset to which i_0 belongs. We will prove that $C_0 \subset W$. We have to see that $x(i) > \|x\|_R + 2u_\kappa + \check{w}$ for every $i \in C_0$. If $i = i_0$, this follows because

$$x(i_0) = \|x\|_{R^c} > \|x\|_R + |R^c|(2u_\kappa + \check{w}) \geq \|x\|_R + 2u_\kappa + \check{w}. \quad (17)$$

If $i \neq i_0$, consider a path in (V', E') connecting i_0 with i , namely, take $i_1, \dots, i_L \in C_0$ (L is simply the length of the path) such that $i_L = i$, $\{i_{l-1}, i_l\} \in E'$ for every $1 \leq l \leq L$, and $i_l \neq i_{l'}$ for every $l, l' \in \{0, \dots, L\}$, $l \neq l'$. The definition of

C_0 implies that

$$x(i_0) - x(i) \leq \sum_{l=1}^L |x(i_{l-1}) - x(i_l)| \leq L(2u_\kappa + \check{w}). \quad (18)$$

Using that $L + 1 \leq |R^c|$, we have

$$\begin{aligned} x(i) &\geq x(i_0) - L(2u_\kappa + \check{w}) \\ &> \|x\|_R + |R^c|(2u_\kappa + \check{w}) - L(2u_\kappa + \check{w}) \geq \|x\|_R + 2u_\kappa + \check{w} \end{aligned}$$

as desired.

• We prove now that hypothesis 1 of Lemma 1 holds. For $j \in E(W)$ fixed, we have to see that, conditioned to $\mathcal{B}_{\kappa,n}$, the policy necessarily performs a match with an item class in W . This follows if we prove that

$$[x(i) + U_{(j,i),n}]^+ + w_{(j,i)} > [x(i') + U_{n,(j,i')}]^+ + w_{(j,i')} \quad (19)$$

for every $i \in W$ and $i' \in S_x \setminus W$. We split into the following two subcases: $i' \in S_x \cap R$ and $i' \in (S_x \cap R^c) \setminus W$.

(a) If $i' \in S_x \cap R$, we have

$$\begin{aligned} [x(i) + U_{(j,i),n}]^+ + w_{(j,i)} &> \|x\|_R + 2u_\kappa + \check{w} + U_{(j,i),n} + w_{(j,i)} \\ &\geq x(i') + u_\kappa + \check{w} \\ &\geq [x(i') + U_{n,(j,i')}]^+ + w_{(j,i')}. \end{aligned}$$

(b) For the subcase $i' \in (S_x \cap R^c) \setminus W$, we will use the following obvious property: if $i_1, i_2 \in V'$ both belong to the same set of connected nodes $C \in \mathcal{C}$, and $i_3 \in V'$ is such that $x(i_1) \wedge x(i_2) - (2u_\kappa + \check{w}) \leq x(i_3) \leq x(i_1) \vee x(i_2) + 2u_\kappa + \check{w}$, then also i_3 belongs to C . We next show that

$$x(i) > x(i') + 2u_\kappa + \check{w}. \quad (20)$$

Suppose by contradiction that the inequality is false. Let $C' \in \mathcal{C}$ be the subset to which i' belongs, and let $i'' \in C'$ be such that $x(i'') \leq \|x\|_R + 2u_\kappa + \check{w}$ (such an i'' exists because otherwise i' would belong to W). In particular,

$$x(i') \wedge x(i'') - (2u_\kappa + \check{w}) \leq \|x\|_R + 2u_\kappa + \check{w} < x(i) \leq x(i') \vee x(i'') + 2u_\kappa + \check{w}.$$

The mentioned property implies that $i \in C'$, contradictorily implying that $i \notin W$. Once we have Inequality (20), we can conclude because

$$\begin{aligned} [x(i) + U_{(j,i),n}]^+ + w_{(j,i)} &> x(i') + 2u_\kappa + \check{w} + U_{(j,i),n} + w_{(j,i)} \\ &\geq [x(i') + U_{n,(j,i')}]^+ + w_{(j,i')}. \end{aligned} \quad (21)$$

• As mentioned before, hypothesis 2 of Lemma 1 holds for V' because of the $\text{NCOND}(G, \gamma)$ assumption, and it extends to W simply for being a subset.

We can now apply Lemma 1 to W , use that $x(i) \leq \|x\|_R + (2u_\kappa + \check{w})|R^c|$ for every $i \in R^c \setminus W$, and use that $\|x\|_W = \|x\|_{R^c}$, to obtain

$$\begin{aligned} \sum_{i \in R^c} x(i)[\lambda(i) - \tau_i(x)] &\leq \sum_{i \in R^c \setminus W} x(i)\lambda(i) + \sum_{i \in W} x(i)[\lambda(i) - \tau_i(x)] \\ &\leq [\|x\|_R + |R^c|(2u_\kappa + \check{w})]\lambda(R^c) + \lambda(W^c)(2u_\kappa + \check{w})|W| + (\kappa - \eta)\|x\|_W \\ &\leq [\|x\|_R + |R^c|(2u_\kappa + \check{w})]\lambda(R^c) + \lambda(V)(2u_\kappa + \check{w})|R^c| + (\kappa - \eta)\|x\|_{R^c}. \end{aligned}$$

Replace in (15) to conclude.

Proof of Lemma 1. We divide the proof into three steps.

Step 1: Separation into connected components. We define an auxiliary graph (W, \tilde{E}) by $\{i, j\} \in \tilde{E} \Leftrightarrow |x(i) - x(j)| \leq 2u_\kappa + \check{w}$, and we denote by $\{C_1, \dots, C_L\}$ the partition of W into connected subsets of nodes. We suppose that they are ordered in a decreasing way with respect to the number of stored items: if $l < l'$, and if $i \in C_l$ and $i' \in C_{l'}$, then $x(i) > x(i') + 2u_\kappa + \check{w}$. More precisely, if $i_1, \dots, i_{|W|}$ is an enumeration of the nodes in W such that $x(i_k) \leq x(i_{k+1})$ for every $k \in \{1, \dots, |W| - 1\}$, there are exactly $L - 1$ of them $i_{k_1}, \dots, i_{k_{L-1}}$ such that $x(i_{k_1+1}) > x(i_{k_1}) + 2u_\kappa + \check{w}$ for every $l \in \{1, \dots, L - 1\}$, and they satisfy $x(i_{k_l}) = \max\{x(i) : i \in C_{l+1}\}$ and $x(i_{k_{l+1}}) = \min\{x(i) : i \in C_l\}$. The idea behind these definitions is to distinguish nodes whose weights are too far apart. In particular, we will use the following

property:

$$\text{for every } l \in \llbracket 1, L \rrbracket \text{ and every } j \in E(C_1 \cup \dots \cup C_l), \quad \sum_{i \in C_1 \cup \dots \cup C_l} v_{x,j}(i | \mathcal{B}_{\kappa,n}) = 1. \quad (22)$$

In other words, conditioned to $\mathcal{B}_{\kappa,n}$, an incoming item whose class belongs to $E(C_1 \cup \dots \cup C_l)$ necessarily matches with an item whose class belongs to $C_1 \cup \dots, C_l$. To see why this is true, take $j \in E(C_1 \cup \dots \cup C_l)$. In particular, we have $j \in E(W)$, so $\sum_{i \in W} v_{x,j}(i | \mathcal{B}_{\kappa,n}) = 1$ by hypothesis; hence, we are done if $l = L$. The case $l < L$ follows because, for every $i \in C_1 \cup \dots \cup C_l$ and every $i' \in C_{l+1} \cup \dots \cup C_L$, we have

$$\begin{aligned} [x(i) + U_{(j,i)}]^+ + w_{(j,i)} &= x(i) + U_{(j,i)} + w_{(j,i)} > x(i') + U_{(j,i)} + 2u_\kappa + \check{w} + w_{(j,i)} \\ &\geq [x(i') + U_{(j,i')}]^+ + w_{(j,i')}. \end{aligned} \quad (23)$$

We will also use that

$$\|x\|_{C_l} - (2u_\kappa + \check{w})|W| \leq x(i) \leq \|x\|_{C_l} \quad (24)$$

for any $l \in \llbracket 1, L \rrbracket$ and any $i \in C_l$.

Step 2: Control of the number of stored items in each C_l . In this step, we reduce the problem by controlling the maximum and minimum number of stored items belonging to the same connected subsets. First observe that

$$\sum_{i \in W} x(i)[\lambda(i) - \tau_i(x)] \leq \sum_{i \in W} x(i) \left[\lambda(i) - \sum_{j \in E(i)} \lambda(j)v_{x,j}(i | \mathcal{B}_{\kappa,n})\mathbb{P}(\mathcal{B}_{\kappa,n}) \right] \quad (25)$$

$$\begin{aligned} &= \sum_{i \in W} x(i)h_i(x) \quad (26) \\ &= \sum_{l=1}^L \sum_{i \in C_l} x(i)h_i(x), \end{aligned}$$

the first identity being simply a definition of $h_i(x)$ as the expression in square brackets. Using (24), we have

$$\begin{aligned} \sum_{i \in C_l} x(i)h_i(x) &= \sum_{i \in C_l}^+ x(i)h_i(x) + \sum_{i \in C_l}^- x(i)h_i(x) \leq \|x\|_{C_l} \sum_{i \in C_l}^+ h_i(x) + [\|x\|_{C_l} - (2u_\kappa + \check{w})|W|] \sum_{i \in C_l}^- h_i(x) \\ &= \|x\|_{C_l} \sum_{i \in C_l} h_i(x) - (2u_\kappa + \check{w})|W| \sum_{i \in C_l}^- h_i(x), \end{aligned} \quad (27)$$

where the superscripts $+$ (respectively, $-$) in the sums mean that we are summing over the indexes i s for which $h_i(x) \geq 0$ (respectively, $h_i(x) < 0$). Hence, calling $H_l(x) = \sum_{i \in C_l} h_i(x)$, we have

$$\sum_{l=1}^L \sum_{i \in C_l} x(i)h_i(x) \leq \sum_{l=1}^L \|x\|_{C_l} H_l(x) - (2u_\kappa + \check{w})|W| \sum_{l=1}^L \sum_{i \in C_l}^- h_i(x). \quad (28)$$

However, observe that

$$-h_i(x) = \sum_{j \in E(i)} \lambda(j)v_{x,j}(i | \mathcal{B}_{\kappa,n})\mathbb{P}(\mathcal{B}_{\kappa,n}) - \lambda(i) \leq \sum_{j \in E(i)} \lambda(j)v_{x,j}(i | \mathcal{B}_{\kappa,n}),$$

implying that

$$\begin{aligned} -\sum_{l=1}^L \sum_{i \in C_l}^- h_i(x) &\leq \sum_{l=1}^L \sum_{i \in C_l} \sum_{j \in E(i)} \lambda(j)v_{x,j}(i | \mathcal{B}_{\kappa,n}) \\ &\leq \sum_{l=1}^L \sum_{i \in C_l} \sum_{j \in E(i)} \lambda(j)v_{x,j}(i | \mathcal{B}_{\kappa,n}) \\ &= \sum_{i \in W} \sum_{j \in E(i)} \lambda(j)v_{x,j}(i | \mathcal{B}_{\kappa,n}) \\ &= \sum_{j \in E(W)} \lambda(j) \sum_{i \in W \cap E(j)} v_{x,j}(i | \mathcal{B}_{\kappa,n}) \\ &= \sum_{j \in E(W)} \lambda(j) \\ &\leq \lambda(E(W)), \end{aligned}$$

the second identity following as in (11), the last identity following from hypothesis 1, and the last inequality from the fact that $W \subseteq S_x$ is an independent set. An application of this bound to the right-hand side (r.h.s.) of (28) reduces the proof to proving that

$$\sum_{l=1}^L \|x\|_{C_l} H_l(x) \leq (\kappa - \eta) \|x\|_W,$$

and for this we use an iterative procedure.

Step 3: Iterative procedure. First, recalling (5), observe that

$$\begin{aligned} H_1(x) &= \sum_{i \in C_1} \lambda(i) - \mathbb{P}(\mathcal{B}_{\kappa,n}) \sum_{j \in E(C_1)} \lambda(j) \sum_{i \in C_1} v_{x,j}(i | \mathcal{B}_{\kappa,n}) \\ &= \lambda(C_1) - \mathbb{P}(\mathcal{B}_{\kappa,n}) \sum_{j \in E(C_1)} \lambda(j) \leq \lambda(C_1) - \left(1 - \frac{\kappa}{\lambda(V)}\right) \lambda(E(C_1)) \leq \kappa - \eta. \end{aligned} \quad (29)$$

In the last inequality, we used hypothesis 2; in the second identity, we use Property (22). We are in the following alternative: if it holds that

$$H_l(x) \leq 0 \quad \forall l \in \{2, \dots, L\}, \quad (30)$$

we are done because, in view of (29),

$$\sum_{l=1}^L \|x\|_{C_l} H_l(x) \leq \|x\|_{C_1} H_1(x) \leq (\kappa - \eta) \|x\|_W$$

(observe that $\|x\|_{C_1} = \|x\|_W$). If (30) does not hold, we define

$$L' = \max\{l \in \llbracket 1, L \rrbracket : H_l(x) > 0\},$$

which, from (29), is necessarily strictly larger than one. Then we have that

$$\sum_{l=1}^L \|x\|_{C_l} H_l(x) \leq \sum_{l=1}^{L'} \|x\|_{C_l} H_l(x) \leq \sum_{l=1}^{L'-2} \|x\|_{C_l} H_l(x) + \|x\|_{C_{L'-1}} [H_{L'-1}(x) + H_{L'}(x)], \quad (31)$$

where sums over empty sets are set to zero. If $H_{L'-1}(x) + H_{L'}(x) \leq 0$, we bound the last quantity by $\sum_{l=1}^{L'-2} \|x\|_{C_l} H_l(x)$ and restart the procedure; that is, we look for

$$\max\{l \in \llbracket 1, L' - 2 \rrbracket : H_l(x) > 0\},$$

and so on. If $H_{L'-1}(x) + H_{L'}(x) > 0$, we bound the r.h.s. of (31) by

$$\sum_{l=1}^{L'-3} \|x\|_{C_l} H_l(x) + \|x\|_{C_{L'-2}} [H_{L'-2}(x) + H_{L'-1}(x) + H_{L'}(x)],$$

separate between the cases

$$H_{L'-2}(x) + H_{L'-1}(x) + H_{L'}(x) \leq 0 \text{ and } H_{L'-2}(x) + H_{L'-1}(x) + H_{L'}(x) > 0,$$

and proceed as previously shown. It only remains to address the case where all partial sums are positive; that is, $\sum_{l=l_0}^{L'} H_l(x) > 0$ for any $l_0 \in \llbracket 2, L' - 2 \rrbracket$. In this case, the final step of the iteration gives

$$\sum_{l=1}^{L'} \|x\|_{C_l} H_l(x) \leq \|x\|_{C_1} \sum_{l=1}^{L'} H_l(x). \quad (32)$$

However, as in (29), we have that

$$\begin{aligned} \sum_{l=1}^{L'} H_l(x) &= \sum_{i \in C_1 \cup \dots \cup C_{L'}} \lambda(i) - \mathbb{P}(\mathcal{B}_{\kappa,n}) \sum_{j \in E(C_1 \cup \dots \cup C_{L'})} \lambda(j) \sum_{i \in C_1 \cup \dots \cup C_{L'}} v_{x,j}(i | \mathcal{B}_{\kappa,n}) \\ &= \lambda(C_1 \cup \dots \cup C_{L'}) - \mathbb{P}(\mathcal{B}_{\kappa,n}) \lambda[E(C_1 \cup \dots \cup C_{L'})] \leq \kappa - \eta, \end{aligned}$$

where we are again using Property (22). Plugging this into (32), we get that

$$\sum_{l=1}^{L'} \|x\|_{C_l} H_l(x) \leq (\kappa - \eta) \|x\|_W,$$

which concludes the proof. \square

5.2. Proof of Theorem 1

Necessity of the condition NCOND can be shown similarly to proposition 2 in Mairesse and Moyal (2016): For any $I \in \mathbb{I}(R^c)$ and any $t \geq 0$, all items of class in I that have departed the system before t have necessarily been matched with an element of $E(I)$ arrived before t ; thus, we have that

$$\begin{aligned} X_t(I) &:= \sum_{i \in I} X_t(i) \geq \sum_{i \in I} X_0(i) - \sum_{i \in E(I)} X_0(i) + \sum_{i \in I} N(i)_t - \sum_{i \in E(I)} N(i)_t \\ &\geq \sum_{i \in I} X_0(i) + M_t^I + t(\lambda(I) - \lambda(E(I))), \end{aligned}$$

where for all i , $N(i)_t$ denotes the number of points of the process $N(i)$ up to t , and M^I is a martingale with respect to the natural filtration of N . Therefore, it is routine to check that the above process tends a.s. to infinity as $t \rightarrow \infty$ if $\lambda(I) - \lambda(E(I)) > 0$, and is at best null recurrent if $\lambda(I) - \lambda(E(I)) = 0$.

We now prove sufficiency. By the continuous-time version of the Foster-Lyapunov theorem (Meyn and Tweedie 1993), it suffices to show that

$$\mathcal{L}f_2(x) \leq -1$$

for every x outside a finite set. Take $\kappa = \eta/2$. If $R = \emptyset$, Inequality (6) becomes

$$\mathcal{L}f_2(x) \leq c_1 - \eta \|x\| \tag{33}$$

for a constant c_1 that does not depend on x ; this clearly implies the assertion. If $R \neq \emptyset$, Inequality (7) becomes

$$-2\hat{\gamma} \|x\|_R^2 + c_2 \|x\|_R + c_3 - \eta \|x\|_{R^c} \mathbb{1}\{\|x\|_{R^c} > \|x\|_R + c_4\}, \tag{34}$$

being $\hat{\gamma} = \min_{i \in R} \gamma(i)$. The r.h.s. of (34) is larger than -1 only when $\|x\|_R$ and $\|x\|_{R^c}$ are simultaneously small enough, a fact that occurs only for finitely many x s. \square

5.3. Proof of Proposition 2

To expedite the proof, first observe that, for any G and $\gamma, \lambda \in (\mathbb{R}_+^*)^V$ is an element of $\text{NCOND}(G, \gamma)$ defined by (4) if and only if μ_λ defined by (1) is an element of

$$\text{NCOND}'(G, \gamma) = \{\mu \in \mathcal{P}_V^* : \mu(I) < \mu(E(I)) \text{ for any } I \in \mathbb{I}(R^c)\},$$

being \mathcal{P}_V^* the set of strictly positive probabilities defined on V ; similarly, λ belongs to $\text{NCOND}(G)$ defined by (8) if and only if μ_λ belongs to

$$\text{NCOND}'(G) = \{\mu \in \mathcal{P}_V^* : \mu(I) < \mu(E(I)) \text{ for any } I \in \mathbb{I}\}.$$

First, if G is bipartite and $R = \emptyset$, then theorem 1 of Mairesse and Moyal (2016) shows that the set $\text{NCOND}'(G, \gamma) = \text{NCOND}'(G)$ is empty. From the previous remark, so is the set $\text{NCOND}(G, \gamma) = \text{NCOND}(G)$.

Regarding the converse, let us first assume that G is nonbipartite. Then, again from theorem 1 in Mairesse and Moyal (2016), the set $\text{NCOND}'(G)$ is nonempty. However, plainly, for any γ , we have $\text{NCOND}'(G) \subset \text{NCOND}'(G, \gamma)$ because any element of $\mathbb{I}(R)$ (if any) is also an element of \mathbb{I} . This shows that $\text{NCOND}'(G, \gamma)$ is nonempty and therefore is $\text{NCOND}(G, \gamma)$.

It only remains to consider the case where G is bipartite (in which case $\text{NCOND}'(G)$ is empty again by theorem 1 of Mairesse and Moyal (2016)), and γ is such that $R \neq \emptyset$. Let $V_1 \cup V_2$ be a bipartition of V into two independent sets. Then, it follows from theorem 4.2 in Bušić et al. (2013) that there exists a probability measure $\mu \in \mathcal{P}_V^*$ such that

$$\mu(V_1) = \mu(V_2) = \frac{1}{2} \quad \text{and} \quad \mu(I) < \mu(E(I)) \quad \text{for all } I \in \mathbb{I} \setminus \{V_1, V_2\}. \tag{35}$$

Suppose that $R \cap V_1 \neq \emptyset$ (the case where $R \cap V_2 \neq \emptyset$ is symmetric). Our idea is to modify the measure μ by transferring a small enough amount of mass from a node of V_2 to a node of $R \cap V_1$, so as to belong to $\mathbb{I}(R^c)$. For this, we first set a node $i \in V_2$ and a positive real number ε as follows:

(i) If $V_1 \subset R$, i is arbitrary in V_2 and

$$\varepsilon = \frac{1}{2} \left(\mu(i) \wedge \min\{1 - \mu(j) : j \in R \cap V_1\} \right).$$

(ii) Else, i is an arbitrary element of $E(R^c \cap V_1)$ and we set

$$\varepsilon = \frac{1}{3} \left(\min\{\mu(E(I)) - \mu(I) : I \in \mathbb{I}(R^c \cap V_1)\} \wedge \mu(i) \wedge \min\{1 - \mu(j) : j \in R \cap V_1\} \right), \tag{36}$$

which is strictly positive because V_1 is not an element of $\mathbb{I}(R^c \cap V_1)$.

We then construct the measure $\hat{\mu}$, as follows:

$$\begin{cases} \hat{\mu}(i) = \mu(i) - \varepsilon; \\ \hat{\mu}(j) = \mu(j) + \varepsilon \text{ for some arbitrary } j \in R \cap V_1; \\ \hat{\mu}(k) = \mu(k) \text{ for all other } k \in V. \end{cases}$$

Then, $\hat{\mu}$ clearly is a probability measure on V , such that

$$\mu(A) \leq \hat{\mu}(A) \leq \mu(A) + \varepsilon \text{ and } \mu(B) - \varepsilon \leq \hat{\mu}(B) \leq \mu(B), \text{ for all } A \subset V_1 \text{ and } B \subset V_2. \quad (37)$$

We aim at showing that

$$\hat{\mu}(I) < \hat{E}(I), \text{ for all } I \in \mathbb{I}(R^c). \quad (38)$$

For this, first observe that because V_1 intersects with R it is not an independent set of $\mathbb{I}(R^c)$. Second, if $V_2 \in \mathbb{I}(R^c)$ (which is true if and only if $R \subset V_1$), we get

$$\hat{\mu}(V_2) = \mu(V_2) - \varepsilon = \mu(V_1) - \varepsilon = \hat{\mu}(V_1) - 2\varepsilon = \hat{\mu}(E(V_2)) - 2\varepsilon < \hat{\mu}(E(V_2)).$$

Now take an independent set $I \in \mathbb{I}(R^c) \setminus \{V_2\}$ such that $I \cap V_1 = \emptyset$. Then we have that

$$\hat{\mu}(E(I)) - \hat{\mu}(I) \geq \mu(E(I)) - \mu(I) > 0, \quad (39)$$

where the first inequality follows from (37) together with the fact that $E(I) \subset V_1, I \subset V_2$, and the second, from the facts that $\mathbb{I}(R^c) \subset \mathbb{I}, I \notin \{V_1, V_2\}$ and (35).

Last, take an independent set $I \in \mathbb{I}(R^c) \setminus \{V_2\}$ such that $I \cap V_1 \neq \emptyset$. (Such an element exists only if V_1 is not included in R , so we are in case (ii).) Then we have

$$\begin{aligned} \hat{\mu}(E(I)) - \hat{\mu}(I) &= \hat{\mu}(E(I) \cap V_1) + \hat{\mu}(E(I) \cap V_2) - \hat{\mu}(I \cap V_2) - \hat{\mu}(I \cap V_1) \\ &= \hat{\mu}(E(I \cap V_2)) - \hat{\mu}(I \cap V_2) + \hat{\mu}(E(I \cap V_1)) - \hat{\mu}(I \cap V_1) \\ &\geq \hat{\mu}(E(I \cap V_1)) - \hat{\mu}(I \cap V_1) \\ &\geq \mu(E(I \cap V_1)) - \varepsilon - (\mu(I \cap V_1) + \varepsilon) > 0, \end{aligned}$$

where in the second equality, we use the fact that the graph is bipartite, implying that $E(I) \cap V_1 = E(I \cap V_2)$ and $E(I) \cap V_2 = E(I \cap V_1)$; the first inequality follows from (39) in the case where $I \cap V_2 \neq \emptyset$, or else is an equality; and where we use (37) in the second inequality and (36) is the third one.

We have thus proven that $\hat{\mu}(I) < \hat{E}(I)$ for any $I \in \mathbb{I}(R^c)$. This implies that $\hat{\mu}$ is an element of $\text{NCOND}'(G, \gamma)$ and thereby an element of $\text{NCOND}(G, \gamma)$ that let us conclude. \square

5.4. Proof of Proposition 3

Neglecting the reneging part, we get for all x ,

$$\mathcal{L}e^{\alpha\|x\|_2} = \sum_{i \notin E(S_x)} \lambda(i)(e^{\alpha\|x+\delta_i\|_2} - e^{\alpha\|x\|_2}) + \sum_{i \in S_x} \tau_i(x)(e^{\alpha\|x-\delta_i\|_2} - e^{\alpha\|x\|_2})$$

(recall that $\tau_i(x) = \sum_{j \in E(i)} \lambda(j)v_{x,j}(i)$). Taking common factor $e^{\alpha\|x\|_2}$, using that $e^a - 1 \leq a + a^2$ if $|a|$ is small enough, and that $(\|x \pm \delta_i\|_2 - \|x\|_2)^2 \leq 1$ (this is simply a triangular inequality), last expression can be bounded by

$$\begin{aligned} &e^{\alpha\|x\|_2} \left[\sum_{i \notin E(S_x)} \lambda(i)[\alpha(\|x + \delta_i\|_2 - \|x\|_2) + \alpha^2] + \sum_{i \in S_x} (\alpha(\|x - \delta_i\|_2 - \|x\|_2) + \alpha^2)\tau(i) \right] \\ &= e^{\alpha\|x\|_2} \alpha \left[\sum_{i \notin E(S_x)} \lambda(i)(\|x + \delta_i\|_2 - \|x\|_2) + \sum_{i \in S_x} (\|x - \delta_i\|_2 - \|x\|_2)\tau(i) + \alpha\lambda(V) \right]. \end{aligned}$$

From a Taylor development,

$$\|x \pm \delta_i\|_2 - \|x\|_2 \leq \frac{\pm x(i)}{\|x\|_2} + \frac{1}{2\|x\|_2}.$$

Using this inequality, we get

$$\begin{aligned} & \sum_{i \in E(S_x)} \lambda(i)(\|x + \delta_i\|_2 - \|x\|_2) + \sum_{i \in S_x} (\|x - \delta_i\|_2 - \|x\|_2)\tau(i) \\ & \leq \frac{1}{\|x\|_2} \left(\sum_{i \in S_x} x(i)[\lambda(i) - \tau(i)] + \frac{1}{2} \lambda(V) \right) \\ & \leq \frac{1}{\|x\|_2} \left(\lambda(V)(2u_\kappa + \check{w})|V| + (\kappa - \eta)\|x\|_\infty + \frac{1}{2} \lambda(V) \right) \\ & \leq \frac{\lambda(V)(2u_\kappa + \check{w})|V|}{\|x\|_2} + (\kappa - \eta)\sqrt{|V|} + \frac{\lambda(V)}{2\|x\|_2} \end{aligned}$$

for every $\kappa \in (0, \lambda(V))$. In the second inequality, we used (6). We have obtained

$$\mathcal{L}e^{\alpha\|x\|_2} \leq e^{\alpha\|x\|_2} \left(\frac{\alpha\lambda(V)(2u_\kappa + \check{w})|V|}{\|x\|_2} + \frac{\alpha\lambda(V)}{2\|x\|_2} + \alpha(\kappa - \eta)\sqrt{|V|} + \alpha^2\lambda(V) \right).$$

We can conclude by choosing $\kappa < \eta$ and α such that $\alpha(\kappa - \eta)\sqrt{|V|} + \alpha^2\lambda(V) < 0$. \square

5.5. Proof of Theorem 3

Under our assumptions, Inequality (6) takes the simpler form

$$\mathcal{L}f_2(x) \leq \lambda(V) + 2\lambda(V)(2B + \check{w})|V| - 2\eta\|x\|.$$

The upper bound in (9) follows after integrating with respect to π and using that

$$\int (\mathcal{L}f)(x)\pi(dx) = 0 \quad \text{for every } \pi\text{-integrable } f : \mathcal{X} \rightarrow \mathbb{R}.$$

We now turn to the lower bound of (9). For this, we introduce the following stochastic lower bound for X : Consider a fictional matching model such that, for any $i \in V$, as soon as an item of a class $j \in E(i)$ enters the system, one class- i item leaves the system. Let \tilde{X} be the $\mathbb{N}^{|V|}$ -valued process counting the number of items of each class in the system at any time. It is easy to prove by an immediate coupling argument, that $\tilde{X}_t(i) \leq_{st} X_t(i)$ for all $i \in V$, where “ \leq_{st} ” stands for the strong stochastic ordering. Moreover, it is also immediate to observe that the i th marginal of \tilde{X} is just a birth and death process with up rate $\lambda(i)$ and down rate $\lambda(E(i))$ in any state. Hence, the marginal of the stationary measure $\tilde{\pi}$ of \tilde{X} satisfies

$$\int x(i)\tilde{\pi}(dx) = \frac{\rho_i}{1 - \rho_i}, \quad \text{for } \rho_i = \frac{\lambda(i)}{\lambda(E(i))}.$$

Then, we get that

$$\frac{\rho_i}{1 - \rho_i} \leq \int x(i)\pi(dx) \leq \int \|x\|_\infty\pi(dx),$$

which completes the proof. \square

6. Conclusion

By proving that a quadratic function and its exponential are Lyapunov functions under the generalized N-COND condition, we showed that the generalized max-weight policy has very desirable properties: it is maximal stable even in the presence of noisy observations and has exponential convergence toward its stationary measure. In an ongoing work, we show using systematic and extensive simulations that it also compares very favorably to state-of-the-art matching policies both in terms of long-term variance and convergence.

Acknowledgments

We are grateful to the anonymous referees for several comments that helped us give a clear focus to the article.

References

- Adan I, Weiss G (2012) Exact FCFS matching rates for two infinite multitype sequences. *Oper. Res.* 60(2):475–489.
- Adan I, Weiss G (2014) A skill based parallel service system under FCFS-ALIS—Steady state, overloads, and abandonments. *Stochastic Systems* 4(1):250–299.

- Adan I, Busic A, Mairesse J, Weiss G (2017) Reversibility and further properties of FCFS infinite bipartite matching. *Math. Oper. Res.* 43(2): 598–621.
- Adan I, Kleiner I, Richter R, Weiss G (2018) FCFS parallel service systems and matching models. *Performance Evaluation* 127–128:253–272.
- Boxma OJ, David I, Perry D, Stadjé W (2011) A new look at organ transplantation models and double matching queues. *Probability Engrg. Inform. Sci.* 25(2):135–155.
- Büke B, Chen H (2015) Stabilizing policies for probabilistic matching systems. *Queueing Systems* 80(1–2):35–69.
- Büke B, Chen H (2017) Fluid and diffusion approximations of probabilistic matching systems. *Queueing Systems* 86(1–2):1–33.
- Bušić A, Meyn S (2016) Approximate optimality with bounded regret in dynamic matching models. Preprint, submitted June 25, <https://arxiv.org/abs/1411.1044>.
- Bušić A, Gupta V, Mairesse J (2013) Stability of the bipartite matching model. *Adv. Appl. Probability* 45(2):351–378.
- Caldentey R, Kaplan E, Weiss G (2009) FCFS infinite bipartite matching of servers and customers. *Adv. Appl. Probability* 41(3):695–730.
- Gurvich I, Ward A (2014) On the dynamic control of matching queues. *Stochastic Systems* 4(2):479–523.
- Hairer M (2014) *Convergence of Markov Processes*. Lecture Notes available in the web-page of the author, <https://www.hairer.org/notes/Convergence.pdf>.
- Lu Y, Xie Q, Kliot G, Geller A, Larus JR, Greenberg A (2011) Join-Idle-Queue: A novel load balancing algorithm for dynamically scalable web services. *Performance Evaluation* 68(11):1056–1071.
- Mairesse J, Moyal P (2016) Stability of the stochastic matching model. *J. Appl. Probability* 53(4):1064–1077.
- Meyn SP, Tweedie RL (1993) Stability of Markovian processes. III. Foster-Lyapunov criteria for continuous-time processes. *Adv. Appl. Probability* 3(25):518–548.
- Mitzenmacher M (1996) The power of two choices in randomized load balancing. PhD thesis, University of California, Berkeley, CA.
- Moyal P, Perry O (2017) On the instability of matching queues. *Ann. Appl. Probability* 27(6):3385–3434.
- Moyal P, Perry O (2021) Stability of parallel server systems. *Oper. Res.*, ePub ahead of print August 16, <https://doi.org/10.1287/opre.2021.2125>.
- Moyal P, Busic A, Mairesse J (2018) Loynes construction for the extended bipartite matching. Preprint, submitted March 7, <https://arxiv.org/abs/1803.02788>.
- Moyal P, Busic A, Mairesse J (2020) A product form for the general stochastic matching model. Preprint, submitted February 5, <https://arxiv.org/abs/1711.02620>.
- Nazari M, Stolyar AL (2019) A stochastic matching model on hypergraphs. *Queueing Systems* 91(1–2):143–170.
- Rahmé Y, Moyal P (2019) Reward maximization in general dynamic matching systems. Preprint, submitted July 30, <https://arxiv.org/abs/1907.12711>.
- Tassiulas L, Ephremides A (1992) Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automated Control* 37(12):1936–1948.