

# A Bilevel Formalism for the Peer-Reviewing Problem

Gennaro Auricchio<sup>a,\*</sup>, Ruixiao Zhang<sup>b</sup>, Jie Zhang<sup>a</sup> and Xiaohao Cai<sup>b</sup>

<sup>a</sup>Department of Computer Science, University of Bath

<sup>b</sup>School of Electronic and Computer Science, University of Southampton

**Abstract.** Due to the large number of submissions that more and more conferences experience, finding an automatized way to well distribute the submitted papers among reviewers has become necessary. We model the peer-reviewing matching problem as a *bilevel programming (BP)* formulation. Our model consists of a lower-level problem describing the reviewers' perspective and an upper-level problem describing the editors'. Every reviewer is interested in minimizing their overall effort, while the editors are interested in finding an allocation that maximizes the quality of the reviews and follows the reviewers' preferences the most. To the best of our knowledge, the proposed model is the first one that formulates the peer-reviewing matching problem by considering two objective functions, one to describe the reviewers' viewpoint and the other to describe the editors' viewpoint. We demonstrate that both the upper-level and lower-level problems are feasible and that our BP model admits a solution under mild assumptions. After studying the properties of the solutions, we propose a heuristic to solve our model and compare its performance with the relevant state-of-the-art methods. Extensive numerical results show that our approach can find fairer solutions with competitive quality and less effort from the reviewers.<sup>1</sup>

## 1 Introduction

The peer-review process is the procedure followed by scientific journals to establish the novelty and correctness of research papers submitted for publication [26, 19]. The principle behind this process is simple. Once the editors of a journal receive a submission, they invite a sufficient number of reviewers to review it and, depending on the reviewers' reports, the editors decide whether the paper meets the criteria for publication or not. To make this process meaningful, there must be a concordance between the topic of the paper and the expertise of the reviewers, ensuring both the authors of the paper and the editors that the reviewers are competent on the topics discussed in the paper. Although the peer-review procedure has been proven to be reliable for scientific journals, it shows some flaws when applied to large conferences. In this case, all authors need to submit their papers by a specified deadline, after which all the submitted papers will be allocated to some reviewers by the editors. Since most conferences have only one submission deadline per year, all the papers are submitted in a few days; consequently, the editors need to handle a large number of papers at the same time. In particular, the editors have to determine the allocation hastily since this allocation process needs to respect the tight schedule of the conference.

Due to the large number of papers the editors usually receive and the lack of time, the allocation process becomes unaffordable through classic means. To meet this need for time efficiency, several automatized ways to determine a peer-reviewer allocation have been proposed. As we will detail in the related work below, the vast majority of the approaches assume that the best peer-reviewing matching is the minimum/maximum of a suitable objective function. This allows rephrasing the whole problem as an optimization problem whose objective encodes an appealing property that the peer-reviewing allocation should have. Albeit every model finds an allocation following its own criterion, all the models tacitly assume that the editors are the agents that actively decide the allocation. This is limiting for two reasons. One is that the efficiency of these methods is bound to how accurate the computation of the objective function is; the other is that reviewers are often allowed to interfere in the assignment procedure by refusing to review some papers or by bidding for them.

To represent the interplays between different agents (*i.e.*, editors and reviewers), we propose and study a model that phrases the peer-reviewing allocation problem through a *bilevel programming (BP)* formulation. Bilevel optimization has been proven to be a valuable expedient to describe several allocation problems in a multi-agent framework, such as the ones concerning resource allocation [31], traffic engineering [6], genomic problems [4], and Knapsack Problems [8]. The appeal of BP problems lies in the fact that relating the solutions of different optimization problems well captures how the choice of one agent affects the others. For a complete discussion on the BP problems, we refer the reader to [7].

**Related Work.** At its core, finding a peer-reviewing matching is a matching problem for indivisible goods. Matching problems were first introduced in the first half of the twenty century by Hitchcock [16] and Kantorovich [20]. Ever since they were introduced, this class of problems found several different applications, such as stable marriage [17, 25], workers and job matching [11, 28], resource allocation [13, 30], and biomedical data analysis [2].

It was observed that the first attempt at modeling the peer-reviewing process as a matching problem dates back to 2007; and then an integer linear programming (ILP) model was proposed, whose solution is an allocation minimizing the possible complaints of the reviewers. This was done with or without considering the reviewers' expertise [15]. In [14, 23], the authors improved this model by searching for a fair allocation, *i.e.*, an allocation that equalizes every reviewers' payoff as much as possible. In [14], this is accomplished by maximizing the minimum of the reviewers' payoffs, while in [23], the authors try to implement classic social inequality functions, such as the Gini's social evaluation function. Another similar approach is taken in [9], where the authors proposed a new cate-

\* Corresponding Author. Email: ga647@bath.ac.uk

<sup>1</sup> Our code website: <https://github.com/Galaxy-ZRX/Bilevel-Review>.

gory of matching problems, called *diverse matching problems*. Despite the differences of those models, they all share the idea that the optimal peer-reviewing allocation is characterizable as the minimum/maximum of an objective function over a suitable restricted set of matching. Moreover, all these routines are able to detect a solution without gathering data from the reviewers. Indeed, all the objective functions used by these procedures are assumed to be known *a priori* or to be computable through machine learning means using data from open-access sites, such as Google Scholar. These concepts are also at the base of the Toronto Paper Matching System (TPMS) [5] – one of the most credited and used peer-reviewing matching procedures.

**Paper Outline.** In Section 2, we briefly review the classic ILP model for the peer-reviewing matching problem. In Section 3, we introduce our BP model for the peer-reviewing matching problem. We describe and study the upper-level problem (ULp) and lower-level problem (LLp) of our formulation and give a set of conditions that ensure the existence of a solution. In Section 4, a heuristic solution is introduced to our BP model and its performance compared with the ones obtained by the relevant state-of-the-art. The results show that our model can find quasi-optimal solutions that are much fairer than the classic ones. Finally, we conclude in Section 5 and outline some future research directions. Due to space limit, we report the proofs of the proposed statements, the discussion on the Secondary Variational Problem, and additional numerical results in the Appendix.<sup>2</sup>

## 2 Maximum Edge-weighted Matching formulation

In this section, we recall the ILP model on which the TPMS is based [5]. Throughout the paper, we denote with  $\mathbb{B}_{nm}$  the set of  $n \times m$  binary matrices. Given  $A, B \in \mathbb{B}_{nm}$ , we say that  $A \leq B$  if  $a_{i,j} \leq b_{i,j}, \forall i \in [n], j \in [m]$ , where  $[\ell]$  denotes the set containing the first  $\ell$  natural numbers, *i.e.*,  $[\ell] := \{1, 2, 3, \dots, \ell\}$ .

In mathematical terms, peer-reviewing matching can be described as the solution of a maximum edge-weighted matching problem over a bipartite graph  $G$ . The two sides of the bipartite graph  $G$  are composed of the set of papers  $\mathfrak{P} = [n]$  and the set of reviewers  $\mathfrak{R} = [m]$ . The bipartite graph is therefore given by  $G = ([n] \cup [m], [n] \times [m])$  and any matrix  $X \in \mathbb{B}_{nm}$  describes a possible papers allocation. Following the classic conventions, the edge  $(i, j)$  is active, *i.e.*,  $X_{i,j} = 1$ , if and only if paper  $i$  is allocated to reviewer  $j$ .

Since the editors have to ensure a certain amount of reviews for every paper and no reviewer can be overburdened with papers to review, every feasible matching has to comply with some restriction that bounds the number of active edges connected to every vertex of the bipartite graph. The linear objective function of the problem is determined by the edge weight matrix  $W = \{w_{i,j}\}_{(i,j) \in G}$ . The weight  $w_{i,j}$  of the edge  $(i, j)$  represents the expected quality of the review that paper  $i$  receives from reviewer  $j$ , hence we refer to the matrix  $W$  as the *quality matrix*. In this framework, the optimal peer-reviewing matching is characterized as the solution of the maximum edge-weighted matching problem induced by the quality matrix  $W$  over the bipartite graph  $G$ , since it describes the matching maximizing the overall quality of the reviews. We below detail the constraints the editors must comply with and how the matrix  $W$  is determined.

**Problem Constraints.** We have two sets of constraints, one for each side of the bipartite graph. Their role is to regulate the number of reviews that every paper receives and to limit the number of reviews that every reviewer can be asked to perform. In detail, the con-

straints are as follows. (i) The set of constraints that is imposed over the set of reviewers is  $\sum_{i \in [n]} X_{i,j} \leq U_j, \forall j \in [m]$ . To do so, we restrict the number of papers the editors can assign to each reviewer. The value  $U_j$  is therefore the maximum number of papers that can be allocated to reviewer  $j$ . (ii) The set of constraints that is imposed over the set of papers is  $l_i \leq \sum_{j \in [m]} X_{i,j} \leq u_i, \forall i \in [n]$ . To do so, we bound the editors to assign every paper  $i$  to at least  $l_i$  and at most  $u_i$  reviewers. (iii) We assume every paper is indivisible, imposed by the restriction  $X \in \mathbb{B}_{nm}$ . Finally, the optimal peer-reviewing assignment is retrieved by solving the following ILP problem.

**Problem 1** Given  $n, m \in \mathbb{N}$ , and a quality matrix  $W$ , the classic ILP model for the peer-reviewing assignment is

$$\max_{X \in \mathbb{B}_{nm}} \langle W, X \rangle, \quad \text{s.t.} \quad l_i \leq \sum_{j \in [m]} X_{i,j} \leq u_i, \quad \sum_{i \in [n]} X_{i,j} \leq U_j,$$

where the constraints hold for every  $i \in [n]$  and  $j \in [m]$  and  $\langle \cdot, \cdot \rangle$  is the scalar product in the Euclidean space.

It is easy to see that, as long as  $l_i \leq u_i, \forall i \in [n]$  and  $\sum_{i \in [n]} l_i \leq \sum_{j \in [m]} U_j$ , Problem 1 is well-defined and admits a solution [12].

**Quality Matrix  $W$  and Its Role.** Determining a meaningful quality matrix  $W$  is of key importance for every allocation method based on the ILP model described in Problem 1. For this reason, retrieving a quality matrix  $W$  that reliably represents the expertise of the reviewers is still a topic of study and discussion.

One way to retrieve the matrix  $W$  is to associate every paper and reviewer to a set of labels that summarizes the contents and the fields of expertise, respectively. Since both the contents of the paper and the fields of expertise are drawn from the same set of topics, it is possible to represent these labels as vectors over an  $r$  dimensional space, where  $r$  is the total number of possible topics. If we denote with  $v_{p_i}$  the vector describing the topics of the paper  $i$  and with  $v_{r_j}$  the expertise vector of reviewer  $j$ , we can then compute the quality of the review performed by reviewer  $j$  on paper  $i$  as the scalar product between  $v_{p_i}$  and  $v_{r_j}$ , [27]. This metric does make sense since the more the vectors are aligned, the higher will be their scalar product. Moreover, if all the vectors have only positive entries, all the scalar products will return a positive value. Since there is no canonical way to represent papers and reviewers in an  $r$  dimensional space, this method highly depends on what embedding we use to translate papers and reviewers into vectors. Usually, the embedding is determined through neural network structures that use data reported from the reviewer or available from sites on the internet (such as Google Scholar [5]). Despite the remarkable results obtained by those procedures, it is a common belief that evaluating the quality of a review only from the publication records of the reviewer is limiting [5]. Indeed, these models do not consider other factors that are instead related to the effort the reviewer would put in reviewing the paper, such as personal interest, personal conflict, or even time at the disposal.

## 3 Bilevel Formulation of the Assignment Problem

In this section, we first introduce our BP formulation for the peer-reviewing matching problem, followed by studying the LLp and ULp and showing that our model admits a solution under mild conditions. We then highlight the relationship between our model and the classic literature. Throughout the paper, we assume that the editor in charge is just one person.

<sup>2</sup> See full version at <https://arxiv.org/abs/2307.12248>.

### 3.1 Bilevel Programming Model

Given a set of papers  $\mathfrak{P} = [m]$  and a set of reviewers  $\mathfrak{R} = [n]$ , we define our BP peer-reviewing matching problem as follows.

**Problem 2** Given  $l_i$ ,  $u_i$  and  $U_j$  as in Section 2, we consider the following problem:

$$\begin{aligned} \max_{Z, X \in \mathbb{B}_{n,m}} \quad & \langle W_E, X \rangle + \langle Y^*, X \rangle \\ \text{s.t.} \quad & l_i \leq \sum_{j \in [m]} X_{i,j} \leq u_i, \quad \sum_{i \in [n]} X_{i,j} \leq U_j, \\ & \sum_{i \in [n]} Z_{i,j} = U_j + \phi_j, \quad X \leq E - Z + Y^*, \quad (1) \\ & Y_j^* \in \underset{Y_j \in \mathbb{B}_n}{\operatorname{argmin}} \langle (W_R)_j, Y_j \rangle, \\ & \sum_{i \in [n]} Y_{i,j} = U_j, \quad \text{and} \quad 0 \leq Y_j \leq Z_j, \end{aligned}$$

where  $E$  is a  $n \times m$  matrix whose entries are all equal to 1. We denote with  $W_E$  the  $n \times m$  matrix describing the qualities of all the possible reviews from the editor's point of view, with  $W_R$  the  $n \times m$  matrix describing the reviewers' efforts, and with  $\phi_j \in \mathbb{N}$  the number of papers that the reviewers can refuse to review. Finally, since we assume that every reviewer bids independently, the LLp is a component-wise minimization, i.e.,  $Y^*$  is the matrix whose columns minimize the function  $Y_{:,j} \rightarrow \langle (W_R)_{:,j}, Y_{:,j} \rangle$ .

In Table 1, we report the variables of our model and their meaning.

**Remark 1** It is well-known that every BP problem describes a Stackelberg game [29, 24], and this is no exception. In our case, the leader is the editor who proposes  $Z$  to the followers, which are the reviewers. Afterwards, the reviewers report a vector  $Y$  according to their preferences. Every couple  $(Z, Y)$  uniquely defines a maximum edge-weighted matching problem, whose objective value corresponds to the payoff of the editor<sup>3</sup>. Notice that the BP formulation we introduced is also equivalent to the following three-phase procedure. First, the editor proposes a set of papers, described by the vector  $Z_j$ , to every reviewer  $j$ . Second, the reviewers select a set  $Y_j$  out of the set  $Z_j$ . The vector  $Y_j$  represents the papers that reviewer  $j$  would review if asked. Finally, the editor gathers the replies of all the reviewers and searches for a papers-reviewers allocation that maximizes the quality and matches the reports given by the reviewers as much as possible. For this reason, given any feasible triplet to problem (1), namely  $(X, Y, Z)$ , we refer to  $Z$  as the editor's proposal, to  $Y$  as the reviewers' bidding, and to  $X$  as the final assignment. A graphical description of this procedure is also given in Figure 1.

As in the previous model, we assume that the editor aims to maximize the quality of the peer-reviewing matching. In our case, however, both the objective function and the constraints of the problem depend on the reviewers' biddings, i.e., the solution of the LLp. The influence that the LLp has on both the objective function and the set of feasible matching of the ULp is the way in which we model the interaction between the reviewers and the editor.

**Remark 2 (The effort matrix  $W_R$  and its role)** To describe the reviewers' point of view, we introduce a matrix  $W_R$  that assigns an effort value to every reviewer-paper couple. When the number of papers is small, it is possible to retrieve  $W_R$  by directly asking the reviewers. As an alternative, we can build a proxy version of  $W_R$  using machine learning. Indeed, there are machine learning methods

**Table 1.** List of variables and parameters used in our BP model.

$l_i$	minimal number of reviews needed by paper $i$
$u_i$	maximal number of reviews needed by paper $i$
$U_j$	number of papers that reviewer $j$ bids for and maximum number of papers that $j$ may review
$Z_j$	vector describing the set of papers proposed to reviewer $j$
$Y_j$	vector describing the set of papers that reviewer $j$ bids for
$Y_j^*$	vector that minimizes reviewer $j$ 's effort
$X$	final allocation decided by the editor
$\phi_j$	degree of freedom of reviewer $j$
$W_E$	quality matrix from the editor's point of view
$W_R$	reviewers' effort matrix

able to find a complete preference order over a set of papers by only having a partial declaration of preferences [5]. In previous models these values were used to enhance the quality matrix  $W_E$  by refining the values that are less accurate due to a lack of information about the reviewer. Instead, in our model, we use the preference order to describe the criterion used by the reviewer during the bidding phase.

### 3.2 Lower-Level Problem

We now consider the LLp, which captures the reviewers' point of view. We first describe the objective function and the constraints, and then we show that the LLp has a solution and study its properties.

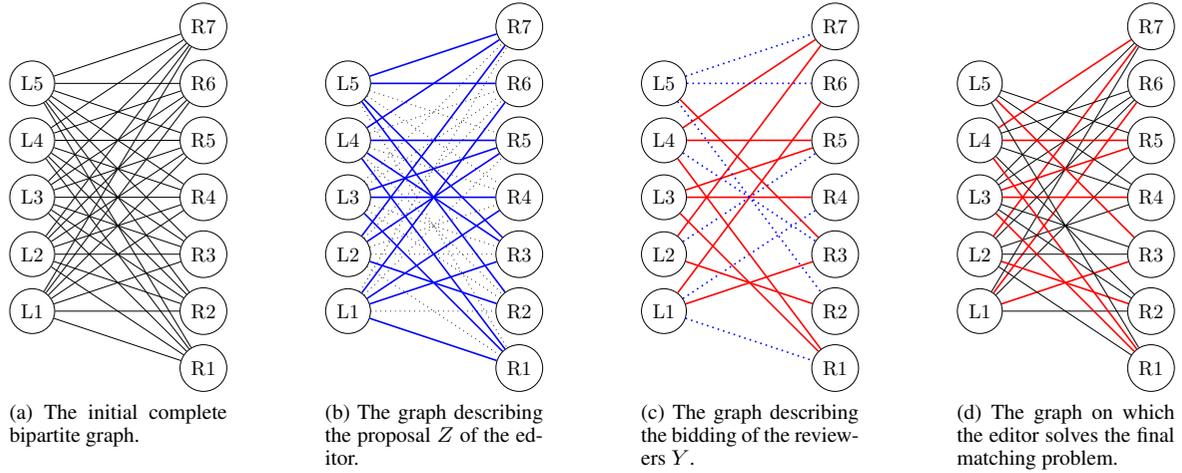
#### 3.2.1 Reviewer's Objective Function

Once the reviewers receive the editor's proposal  $Z$ , they have to bid for part of the papers proposed and report their bidding to the editor. In the following, we denote with  $Y_j$  the  $j$ -th column of the matrix  $Y$ , so that  $Y_j = Y_{:,j}$ . Similarly,  $(W_R)_j$  and  $Z_j$  denote the  $j$ -th column of the matrices  $W_R$  and  $Z$ , respectively. We model the interests of the reviewers through a linear objective function that evaluates the effort required by every reviewer  $j$  to review a set of papers  $Y_j$ . In this framework, the bidding criterion adopted by reviewer  $j$  is captured by the minimization of the objective function  $Y_j \rightarrow \langle (W_R)_j, Y_j \rangle$ . Every entry of the reviewers' effort vector  $((W_R)_j)_i$  describes how much effort reviewer  $j$  would spend to review paper  $i$ . To make the model meaningful, we assume that each entry of the vector  $(W_R)_j$  is positive, which means that reviewing is never effortless.

#### 3.2.2 Reviewer's Constraints

Every reviewer  $j$  has to comply with two sets of constraints. Since reviewers are anonymous to each other, we assume that the constraints imposed on one reviewer do not depend on the constraints imposed on other reviewers. Thus, we describe the constraints only for a fixed reviewer, namely  $j$ . To impose the indivisibility of the reviews, every  $Y_j$  is assumed to have binary entries. Finally, the set of constraints with which the reviewers have to comply is as follows. (i) The *consistency constraint*, that is  $Y_j \leq Z_j$ , where  $Z_j$  is the proposal advanced by the editor; this constraint imposes that every reviewer cannot bid for a paper that does not belong to the pool that the editor proposed. (ii) The *quantity constraint*, that is  $\sum_{i \in [n]} Y_{i,j} = U_j$ ; this constraint forces every reviewer to bid for  $U_j$  papers. Since the review effort is always non-negative, imposing the quantity constraint is equivalent to imposing the following constraint  $\sum_{i \in [n]} Y_{i,j} \geq U_j$ .

<sup>3</sup> If the couple  $(Z, Y)$  leads to an unfeasible problem, the payoff is  $-\infty$ . The reviewers' payoff is the total effort of the papers they bid for.



**Figure 1.** An example of the allocation procedure described by our BP problem. In (1a), we have the complete bipartite graph. The left side contains the papers, while the right side contains the reviewers. Reviewers  $R_6$ ,  $R_4$  and  $R_2$  have to review at most 1 paper, while the other reviewers have to review at most 2. All reviewers have a degree of freedom  $\phi_j$  equal to 1. In (1b), we report in blue the editor's proposal  $Z$  done by the editor, while the dotted lines represent the edges that have not been proposed by the editor. In (1c), we report in red the edges the reviewers' bidding  $Y$ , while the blue dotted lines represent the reviews that the reviewers have declined. Finally, in (1d), we report the final graph on which the editor has to solve a maximum edge-weighted matching problem.

**Uniqueness of the Solution and the Equivalence Result.** To conclude, we study the solution of the LLp. As we will see, the set of conditions needed to ensure the uniqueness of the solution and the equivalence with the relaxed problem is natural to assume.

**Lemma 1** *Given two integers  $0 < a < b$ , and let  $\{w_k\}_{k=1,\dots,b}$  be a set of positive and increasingly ordered real values, i.e.,  $0 \leq w_1 \leq \dots \leq w_a \leq w_{a+1} \leq \dots \leq w_b$ . If  $x_i \in [0, 1]$ ,  $\forall i \in [b]$  and  $\sum_{i=1}^b x_i = a$ , then it holds  $\sum_{i=1}^a w_i \leq \sum_{i=1}^b w_i x_i$ .*

We now prove that the ILP problem in the LLp is equivalent to the LP problem obtained by relaxing the constraint  $Y_j \in \{0, 1\}$ .

**Theorem 1** *For any given  $Z \in \mathbb{B}_{m,n}$ , the LP problem*

$$Y^* \in \operatorname{argmin}_{0 \leq Y \leq Z} \langle W, Y \rangle, \quad \text{s.t.} \quad \sum_{i \in [n]} Y_{i,j} = U_j, \quad (2)$$

and the ILP problem

$$\hat{Y} \in \operatorname{argmin}_{Y \in \mathbb{B}_{m,n}, Y \leq Z} \langle W, Y \rangle, \quad \text{s.t.} \quad \sum_{i \in [n]} Y_{i,j} = U_j, \quad (3)$$

attain the same optimal value, i.e.,  $\langle W, Y^* \rangle = \langle W, \hat{Y} \rangle$ . Moreover, if  $\forall j \in [m]$ , the set  $\{W_{i,j} : (i, j) \text{ s.t. } Z_{i,j} = 1\}$  does not contain two equal values, then the solution is unique and satisfies  $\hat{Y} \in \mathbb{B}_{m,n}$ , i.e., Problems (2) and (3) are equivalent.

Notice that the hypothesis of Theorem 3 is satisfied whenever every reviewer has a strict preference order over the set of papers. Finally, we show that minimizing the component-wise reviewers' effort is the same as minimizing the global reviewers' effort.

**Theorem 2** *For every given  $Z$ ,  $Y$  is a solution of the LLp if and only if  $Y$  minimizes the function  $Y \rightarrow \sum_{i \in [n]} \sum_{j \in [m]} (W_R)_{i,j} Y_{i,j}$  under the constraints  $\sum_{i \in [n]} Y_{i,j} - U_j \geq 0$  and  $Z_{i,j} - Y_{i,j} \geq 0$ .*

### 3.3 Upper-Level Problem

We now consider the ULp, which is the problem that captures the editor's point of view. We first describe the role of both the objective function and the constraints, and then we show that the solution to the ULp problem always exists under mild assumptions.

#### 3.3.1 The Editor's Objective Function

In our model, the editor needs to find an allocation that maximizes a linear objective function, described by the matrix  $W_E$  and the scalar product of the variables  $X$  and  $Y$ . While  $\langle W_E, X \rangle$  represents the quality of matching  $X$ , the scalar product between  $X$  and  $Y$  describes how much the final assignment  $X$  meets the reviewers' bidding  $Y$ . We, therefore, define the following quantity.

**Definition 1** *Given a solution  $(X, Y, Z)$  to the BP Problem 2, we define its Accordance Percentage as  $AC(X, Y, Z) := \frac{\langle X, Y \rangle}{\langle X, X \rangle}$ .*

By definition, we have  $AC(X, Y, Z) \in [0, 1]$  and  $AC(X, Y, Z) = 1$  if and only if  $X \leq Y$ , i.e. if and only if every reviewer receives only papers they bid for. If  $AC(X, Y, Z) = 1$  holds, we say that the solution  $(X, Y, Z)$  is *perfect*.

#### 3.3.2 The editor's constraints

In our model, the editor controls two variables,  $Z$  and  $X$ . Both  $Z$  and  $X$  are  $n \times m$  binary matrices that are subject to two different sets of constraints below.

**Constraints over  $Z$ .** We require the proposal  $Z$  to comply with only one set of constraints, which we call *freedom constraint*. The role of these constraints is to determine the number of papers the editor has to propose to every reviewer. In mathematical terms, this means that  $\sum_{i \in [n]} Z_{i,j} = U_j + \phi_j$ ,  $\forall j \in [m]$ , where the constants  $\phi_j \geq 0$  are the degrees of freedom. In fact, since the reviewer  $j$  has to bid for  $U_j$  papers, the parameter  $\phi_j$  corresponds to the number of papers that the reviewer can refuse to review.

**Constraints over  $X$ .** The final assignment  $X$  has to comply with the constraints as follows. (i) The *feasibility constraints*, that is,  $l_i \leq \sum_{j \in [m]} X_{i,j} \leq u_i, \forall i \in [n]$  and  $\sum_{i \in [n]} X_{i,j} \leq U_j, \forall j \in [m]$ . These constraints bound the editor to assign every paper  $i$  to a number of reviewers between  $l_i$  and  $u_i$  and to assign at most  $U_j$  reviews to every reviewer  $j$ . (ii) The *consistency constraint*, that is  $X \leq E - Z + Y^*$ , where  $Y^*$  is the solution to the LLp and  $Z$  is the editor's proposal. This constraint bounds the editor to assign a paper to a reviewer only if the reviewer has not refused it before.

**The Solutions of the ULp** In the previous section, we proved that the LLp admits a solution, regardless of what the editor proposes. As the next example shows, the same does not hold for the ULp.

**Example 1** *Let us consider the case in which we have three papers, namely  $p_1, p_2$ , and  $p_3$ , and two reviewers, namely  $r_1$  and  $r_2$ . We assume every paper needs exactly one review (i.e.,  $l_i = u_i = 1$ ) and that every reviewer can be asked to review at most 2 papers (i.e.,  $U_j = 2$ ). Finally, we assume that  $W_R$  is defined as  $W_R = \begin{pmatrix} 11 & 10 & 1 \\ 3 & 2 & 1 \end{pmatrix}$ . Both reviewers have the same preference order over the set of papers, i.e.,  $p_1 > p_2 > p_3$ . It is then easy to see that, if  $\phi_1, \phi_2 > 0$ , no reviewer will give its consensus to review paper  $p_3$ , making the ULp, and hence the whole BP problem, unfeasible.*

Example 1 points out that allowing the reviewers to review or refuse too many papers leads to an unsolvable problem in some pathological cases. In fact, if we ask the second reviewer to review only one paper, that is  $U_2 = 1$ , the BP problem described in Example 1 becomes feasible. In general, the editor can ensure the existence of a solution to Problem 2 by tuning the parameters of the problem.

**Theorem 3** *If, in addition to the assumptions that make Problem 2 feasible, it holds  $\max_{j \in [m]} \phi_j + 2 \max_{j \in [m]} U_j \leq n$ , then there exists a feasible triplet to Problem 1. In particular, under these assumptions, Problem 1 has a solution.*

Notice that whenever the number of papers  $n$  is large enough, the hypothesis of Theorem 3 is likely to be satisfied, since both  $U_j$  and  $\phi_j$  are parameters bounded to the human capacity. For example, if  $n = 30$ , the conditions will hold even if we require every reviewer to review at most 8 papers while allowing everyone to reject up to 10 papers. Finally, even when the number of papers is small, the editor can decrease  $\phi_j$  and  $U_j$  in order to make them suitable.

### 3.4 Fairness and the relation with the ILP model

In Example 1, we showed that when the reviewers share the same preference order over the set of papers, finding a solution to the BP formulation might be impossible. In the following, we show that this phenomenon is also due to an intrinsic fairness property that perfect solutions to the BP problem possess.

**Definition 2** *Let  $X$  be a matching over the bipartite graph  $G$  and let  $W_R$  be the reviewers' effort matrix. Given a vector  $\Phi := (\phi_1, \dots, \phi_m) \in \mathbb{N}^m$ , we say that  $X$  is  $\Phi$ -weakly fair if,  $\forall j \in [m]$ , the set  $BC_j(X) := \{i \in [n], \text{ s.t. } X_{i,j} = 0 \ \& \ (W_R)_{i,j} \geq \min_{X_{i,j}=1} (W_R)_{i,j}\}$ , contains at least  $\phi_j$  elements.*

Notice that the set  $BC_j(X)$  contains all the papers in  $[n]$  that require more effort than the highest effort-requiring paper assigned to

$j$  according to  $X$ . Therefore, a solution is  $\Phi$ -weakly fair if every reviewer  $j$  is not allocated with its  $\phi_j$  worst picks. Using the notion of a perfect solution, we are able to relate  $\Phi$ -weakly fair solutions of the classic ILP model to the perfect solutions of the BP model.

**Proposition 1** *There exists a perfect solution  $(X, Y, Z)$  that maximizes the quality if and only if the ILP problem admits a  $\Phi$ -weakly fair solution.*

To conclude, we show that our model extends the classic ILP formulation described in Problem 1. In particular, whenever we set  $\phi_j = 0, \forall j \in [m]$ , the solution of the LLp is actually determined by the editor. This allows us to find a bijection between the solutions of the classic ILP model and the solutions of our BP model.

**Proposition 2** *Set  $\phi_j = 0, \forall j \in [m]$ . Then, given a solution  $(X, Y, Z)$  of the BP problem, we have that  $X$  solves Problem 1. Vice-versa, if  $X$  is a solution to Problem 1, then all the triplets  $(X, Y, Z)$  such that  $X \subset Y = Z$  are perfect solutions to the BP problem.*

In Appendix B, we delve further into the relationships between our BP model and previous models. In particular, we focus on the class of diverse matching proposed in [9], show how this specific model is related to a Secondary Variational Problem induced by Problem 1, and adapt it to our BP model.

## 4 Numerical Experiments

In this section, we report the results of our numerical experiments. We first introduce the heuristic we use to solve the model. We then describe the experimental setting, the implementation details, and introduce the metrics used to compare the outcomes of our experiments. Finally, we comment on the results we find.

### 4.1 Greedy Heuristic

Due to the complexity of the BP model described in Problem 2, approaching the problem through Bilevel solvers is prohibitive. For this reason, we propose a greedy heuristic that finds a solution by fixing the editor's proposal  $Z$ . Indeed, once we fix  $Z$ , retrieving the reviewers' best response is simple since it is the minimum of  $m$  independent LP problems. Once we retrieve the best reply  $Y_Z$ , we obtain the final allocation by solving the maximum edge-weighted matching problem induced by  $Z$  and  $Y$ . For every  $Z$  and  $Y$ , we denote the related final allocation with  $X_{Z,Y}$ . Given a quality matrix  $W_E$ , we define the greedy proposal  $Z_g$  as the solution to the following ILP problem

$$Z_g = \operatorname{argmax}_{Z \in \mathbf{B}_{n,m}} \langle W_E, Z \rangle, \quad \sum_{i \in [n]} Z_{i,j} = U_j + \phi_j, \quad \forall j \in [m].$$

In particular, the greedy proposal  $Z_g$  is the editor's proposal that maximizes the total quality of the papers proposed. Following the previous notation, we define the greedy heuristic solution to Problem 2 as  $(X_g, Y_g, Z_g) := (X_{Z_g, Y_{Z_g}}, Y_{Z_g}, Z_g)$ . We notice that to build our heuristic solution, we need to solve at most three ILP problems. Therefore, complexity-wise, our approach is as costly as any method based on the resolution of an ILP problem. Moreover, when the set of reviewers is large and heterogeneous enough, the greedy heuristic always finds a feasible peer review matching.

**Theorem 4** *The greedy heuristic always finds a feasible peer review matching if, for every paper  $i \in [n]$ , there are at least  $L = \sum_{i \in [n]} l_i$  reviewers that do not place  $i$  among its top  $K := (\max_{j \in [m]} \phi_j + \max_{j \in [m]} U_j)$  picks.*

Finally, we show that if the accordance percentage of the heuristic solution is equal to 1, then the quality of the matching  $X_g$  lower bounds the quality of the optimal solution to the BP problem.

**Theorem 5** *Let  $(X_g, Y_g, Z_g)$  be the triplet found by the greedy heuristic. If  $AC(X_g, Y_g, Z_g) = 1$ , then, for every optimal solution  $(X^*, Y^*, Z^*)$  to the BP problem, we have  $\langle W_E, X_g \rangle \leq \langle W_E, X^* \rangle$ .*

## 4.2 Experimental Framework

We now detail the dataset, the parameters, and the solver we used during the implementation of our experiments.

**Dataset.** We use the multi-aspect review assignment evaluation dataset [22, 21], which is a classic benchmark dataset from UIUC. The dataset contains the information of 73 papers accepted by SIGIR 2007, as well as 189 prospective reviewers who had published in the main information retrieval conferences. Moreover, the dataset provides 25 major topics and, for each paper in the set, it provides a 25-dimensional label on that paper based on a set of defined topics. Similarly, for each of the 189 reviewers, a 25-dimensional expertise representation is provided. To the best of our knowledge, there is no available dataset from where we could retrieve the effort matrices needed for our tests. For this reason, we rely on two synthetically generated datasets as follows. (i) The *Aligned dataset*, in which there is a linear dependence between  $W_R$  and  $W_E$ . We assume that  $(W_R)_{i,j} = K - ((W_E)_{i,j} + \chi)$ , where  $K$  is a constant ensuring  $(W_R)_{i,j} > 0$  and  $\chi \sim \mathcal{N}(0, \sigma)$ . In this case, maximizing the quality of the assignment and minimizing the total reviewers' effort have the same objective function up to a Gaussian noise  $\chi$ , which describes the personal preferences the reviewers may have over papers from the same area. (ii) The *Random dataset*, in which every entry of  $W_R$  is defined as  $(W_R)_{i,j} = \chi$ , where  $\chi$  is a random variable. In our test, we consider two possible laws for  $\chi$ : the Uniform distribution  $\mathcal{U}[0, 1]$  and the Exponential distribution  $Exp(0.5)$ . In this scenario, the effort needed by any reviewer does not depend on the paper, but rather on external factors not related to the expected quality of the review, such as time at disposal and conflict of interest.

**Experiment Setting.** For every instance in our synthetic dataset, we use different methods to retrieve feasible peer-reviewer matchings. In particular, we consider: (i) The matching obtained by solving the ILP problem described in Problem 1. We call this solution *Pure-Quality solution* and denote it with  $X_{ILP}$ . (ii) The matching obtained by the greedy heuristic we presented above. We denote this matching with  $X_{BP}$ . (iii) The matching obtained by considering an ILP problem whose objective function tries to maximize the Quality and minimize the Reviewers' Effort at the same time. Given  $t$ , we say that  $X$  is a *t-tuned solution* if it maximizes  $\langle (1-t)W_E + t(-W_R), X \rangle$  under the constraints  $l_i \leq \sum_{j \in [m]} X_{i,j} \leq u_i$  and  $\sum_{i \in [n]} X_{i,j} \leq U_j$ .<sup>4</sup> We use  $X_{ILP}^{(t)}$  to denote the  $t$ -tuned solution.

Finally, we fix the parameters we use in our simulations. We set the values  $u_i$  and  $l_i$  respectively to be 3 and 5,  $\forall i \in [n]$ . We assume  $U_j$  does not depend on  $j \in [m]$  and let  $U_j$  range in  $\{6, 8\}$ . Similarly, we assume  $\phi_j$  does not depend on  $j \in [m]$  and let its value vary in a set that depends on the value of  $U_j$ . We assume the degree of freedom to be proportional to the value  $U_j$ . In particular, we consider  $\phi_j$  equals to the 50%, the 75%, and the 100% of  $U_j$ , and thus, for  $U_j = 8$ , we have  $\phi \in \{4, 6, 8\}$ . In the Aligned case, we let the variance of the Gaussian noise, namely  $\sigma$ , vary in the set  $\{0.1, 0.3\}$ , since we want

<sup>4</sup> We take  $-W_R$  instead of  $W_R$  because the LLp is a minimization problem.

$X$  to be only a small influence. Due to the fact that the main goal of the editor is to retrieve high quality solutions and since there does not exist a unique best parameter  $t$  to retrieve the  $t$ -tuned solutions, we let  $t$  range in  $\{0.05, 0.1, 0.15\}$ . In that regard, it is worth noticing that the best parameter  $t$  might change depending on the dimensions of the problem or on the reviewers' characteristics. Finally, since all the effort matrices are randomly generated, we run the experiment 250 times for every set of parameters and report the average of our results. Due to space limits, we report only the results for the case of  $U_j = 8$ . All the missing results are deferred in Appendix C.

**Implementation Details.** Our experiments are conducted in the Red Hat Linux Server with Intel Xeon Gold 6138 CPUs. We use Julia [3] as the main coding language and solve the ILP problems by using the JuMP package and HiGHS as an optimization solver [10].

## 4.3 Comparison Metrics

The metrics we use to evaluate our results are the *Quality Percentage (QP)*, the *Reviewers' Average Effort Ratio (RAER)*, and the *Fairness Ratio (FR)*, which we detail in the following.

**Quality Percentage.** Given a peer-reviewing allocation  $X$  and a quality matrix  $W_E$ , we define the quality of  $X$  as  $\mathcal{Q}(X) := \langle W_E, X \rangle$ . Given a solution  $X_{ILP}$  of the ILP problem and a matching  $X$  found, we define the *QP* as  $QP(X) := \frac{\mathcal{Q}(X)}{\mathcal{Q}(X_{ILP})}$ . Since  $X_{ILP}$  does maximize the quality over the set of feasible matchings, we have that  $QP(X) \in [0, 1]$  for every feasible matching.

**Reviewers' Average Effort Ratio.** Given a peer-reviewing allocation  $X$  and an effort matrix  $W_R$ , we define the effort required to reviewer  $j$  from  $X$  as  $\rho_j(X) := \langle (W_R)_{:,j}, X_{:,j} \rangle$  and say that  $j$  is *active* if  $\rho_j(X) > 0$ . We define the total reviewers' effort required by  $X$  as  $\mathcal{E}(X) := \langle W_R, X \rangle = \sum_{j \in [m]} \rho_j(X)$ . Finally, we define the average effort required by the active reviewers as  $\mathcal{E}_{avg}(X) := \frac{\mathcal{E}(X)}{n_{act}(X)}$ , where  $n_{act}(X)$  is the number of active reviewers according to  $X$ . Given two peer-reviewers allocations  $X$  and  $X'$ , we define the Reviewers' Average Effort Ratio as  $RAER(X, X') = \frac{\mathcal{E}_{avg}(X)}{\mathcal{E}_{avg}(X')}$ .

**Fairness Ratio.** Given an allocation matrix  $X$ , let  $\rho$  be the  $n_{act}(X)$ -dimensional vector containing the efforts of all the active reviewers. To measure how fairly spread the effort induced by  $X$  is, we use the variance of vector  $\rho$ , denoted as  $\theta(X)$  [18]. Given two peer-reviewers allocations  $X$  and  $X'$ , we define the Fair Ratio between  $X$  and  $X'$  as  $FR(X, X') = \frac{\theta(X)}{\theta(X')}$ .

## 4.4 Results

In this section, we compare the performance of the three matchings described in Section 4.2. First, we comment how  $X_{BP}$  compares with  $X_{ILP}$  and then compare  $X_{BP}$  with  $X_{ILP}^{(t)}$ .

**The Pure-Quality Solutions Against the Greedy Heuristic.** We first comment on the Aligned case. Notice that, in the Aligned case, maximizing the quality and minimizing the effort are similar tasks, thus this represents the best scenario for the ILP model. Nonetheless, we observe that our model is able to find solutions whose *QP* is consistently larger than 95%, such that (i) the *RAER* is 10% to 20% lower and (ii) the total effort is more evenly spread. Indeed, we observe that the *FR* ranges from 76% to 48%, so that the variance obtained by the ILP model is 1.3 times the variance obtained by our model in the best case and more than double in the worst case. We, therefore, conclude that the personal effort required from every reviewer according to the solution of the BP model is lower on average, spread amongst more reviewers, and, overall, more fairly distributed.

**Table 2.** Comparison between  $X_{BP}$  and  $X_{ILP}$ . Quantitative results for different values of  $\phi$  and differently generated effort matrices. Every column represents a different framework and is characterized by the effort matrix, the random variable used to generate the matrices, and the degree of freedom. For each framework, we report the averages over 250 instances of the Quality Percentage, the Reviewers' Average Effort Ratio, the Fairness Ratio, and the Accordance Percentage of the heuristic solution.

U=8	Aligned						Random					
	$\sigma = 0.1$			$\sigma = 0.3$			$X \sim \mathcal{U}[0, 1]$			$X \sim \text{Exp}(0.5)$		
	$\phi = 4$	$\phi = 6$	$\phi = 8$	$\phi = 4$	$\phi = 6$	$\phi = 8$	$\phi = 4$	$\phi = 6$	$\phi = 8$	$\phi = 4$	$\phi = 6$	$\phi = 8$
$QP(X_{BP})$	0.99	0.99	0.99	0.98	0.98	0.97	0.96	0.95	0.93	0.97	0.96	0.94
$RAER(X_{BP}, X_{ILP})$	0.89	0.89	0.86	0.87	0.87	0.80	0.63	0.53	0.47	0.45	0.35	0.30
$FR(X_{BP}, X_{ILP})$	0.65	0.75	0.57	0.61	0.61	0.48	0.35	0.24	0.19	0.16	0.09	0.07
$AC$	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99

**Table 3.** Comparison between  $X_{BP}$  and  $X_{ILP}^{(t)}$ . Quantitative results for different values of  $\phi$  and differently generated effort matrices. Every column represents a different framework and is characterized by the effort matrix, the random variable used to generate the matrices, and the degree of freedom. For each framework, we report the averages over 250 instances of the Quality Percentage, the Reviewers' Average Effort Ratio, and the Fairness Ratio.

U=8	Aligned						Random					
	$\sigma = 0.1$			$\sigma = 0.3$			$X \sim \mathcal{U}[0, 1]$			$X \sim \text{Exp}(0.5)$		
	$\phi = 4$	$\phi = 6$	$\phi = 8$	$\phi = 4$	$\phi = 6$	$\phi = 8$	$\phi = 4$	$\phi = 6$	$\phi = 8$	$\phi = 4$	$\phi = 6$	$\phi = 8$
$t = 0.05$	$QP(X_{ILP}^{(t)})$	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99
	$RAER(X_{BP}, X_{ILP}^{(t)})$	0.96	0.95	0.95	0.93	0.91	0.89	0.72	0.60	0.53	0.59	0.47
	$FR(X_{BP}, X_{ILP}^{(t)})$	0.77	0.77	0.76	0.70	0.65	0.62	0.43	0.29	0.23	0.26	0.16
$t = 0.1$	$QP(X_{ILP}^{(t)})$	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.98
	$RAER(X_{BP}, X_{ILP}^{(t)})$	0.95	0.95	0.95	0.95	0.93	0.91	0.78	0.66	0.58	0.71	0.57
	$FR(X_{BP}, X_{ILP}^{(t)})$	0.78	0.77	0.77	0.72	0.66	0.63	0.47	0.33	0.25	0.37	0.22
$t = 0.15$	$QP(X_{ILP}^{(t)})$	0.99	0.99	0.99	0.99	0.99	0.99	0.98	0.98	0.98	0.98	0.97
	$RAER(X_{BP}, X_{ILP}^{(t)})$	0.96	0.95	0.95	0.96	0.95	0.92	0.88	0.75	0.66	0.86	0.68
	$FR(X_{BP}, X_{ILP}^{(t)})$	0.79	0.77	0.77	0.74	0.69	0.65	0.55	0.38	0.30	0.53	0.32

We now comment on the Random case, *i.e.* the case in which there is no guaranteed relationship between the quality of a matching and the effort of the agents. For both the uniform and exponential distribution, we observe that the  $QP$  slightly drops between from 94% to 90%. However, the quality drop comes with a larger drop in the  $RAER$  and an even larger drop in the  $FR$ . Indeed, the allocation found by our heuristic halves the effort required by every reviewer in most cases, and in some instances, it achieves almost a third of the effort required by the ILP solution. Similarly, the  $FR$  drops between 0.35 and 0.07 hence the variance of the heuristic solution is up to 15 times lower than the one found by the classic ILP model. All the results of our experiments are reported in Table 2.

**The  $t$ -tuned Solutions Against the Greedy Heuristic.** Finally, we compare our heuristic to the  $t$ -tuned solutions for  $t = 0.05, 0.1, 0.15$ . As we observe from the results, taking a convex combination of the effort matrix and the quality matrix leads to solutions that better balance the quality and the effort. This is clear when we look at the performance on the aligned dataset: in this case, the  $QP$  of both  $X_{BP}$  and  $X_{ILP}^{(t)}$  are comparable for every  $t$ . Similarly, the  $RAER$  between  $X_{BP}$  and  $X_{ILP}^{(t)}$ , ranges between 90% and 95%, hence the  $X_{BP}$  still requires less effort to the reviewers, but only marginally. The real difference between the two solutions can be appreciated when look at their  $FR$ . Indeed, our heuristic is capable of finding a matching that spreads the effort in a more fair way amongst the reviewers. It is again worth stressing, that the aligned case represents the best set of conditions for any ILP model, since maximizing the quality and minimizing the effort do coincide in this specific framework. This consideration is confirmed by the results we get for the Random dataset. In this case, the performances of  $X_{ILP}^{(t)}$  is more in line with the performance of  $X_{ILP}$ : the loss in  $RAER$

is no longer negligible and the loss in  $FR$  increases. Finally, we observe that using our model has a further advantage over the  $t$ -tuned solutions: our method does not require to find a tuning parameter  $t$  to describe the interplay between the editor's and the reviewers' objective function. The experiment results are listed in the Table 3.

**Final Remarks.** In every setting considered, we observe that the variance of our solution is noticeably lower than the one found by any ILP based model. This is especially interesting since we are able to retrieve those fairer solutions without altering the objective function nor by determining a tuning parameter to mix the quality and the effort matrix. Finally, we notice that our accuracy results are acceptable since the percentages we found are well above the 83% obtained by other heuristic approaches, as the one proposed in [1].

## 5 Conclusion and Future Works

In this paper, we introduced an integer BP model for the peer review problem. To the best of our knowledge, this is the first model that finds a peer-reviewing matching using this formalism. We showed that the LLP and the ULp are well-defined problems and that solution has a neat burden-free property under mild assumptions. Finally, we defined a heuristic solution and validated it through numerical tests. From the results, we observe that the heuristic finds fairer solutions that require a lower average effort from the reviewers and achieves a competitive quality without requiring any parameter tuning.

For the future avenue, a first improvement to the model would be to relax the binary constraints of the allocation and search for a probabilistic allocation rather than a deterministic one. Another interesting aspect of the peer review problem is determining the expertise of the reviewers by comparing records from conferences.

## Acknowledgements

This project is partially supported by a Leverhulme Trust Research Project Grant (2021–2024). Jie Zhang is also supported by the EP-SRC grant (EP/W014912/1).

## References

- [1] Faez Ahmed, John P. Dickerson, and Mark Fuge, ‘Diverse Weighted Bipartite b-Matching’, in *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pp. 35–41, Melbourne, Australia, (2017). IJCAI.
- [2] Gennaro Auricchio, Federico Bassetti, Stefano Gualandi, and Marco Veneroni, ‘Computing kantorovich-wasserstein distances on d-dimensional histograms using (d+1)-partite graphs’, in *Advances in Neural Information Processing Systems*, eds., S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, volume 31, p. 5798–5808, Montreal, Canada, (2018). Curran Associates, Inc.
- [3] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah, ‘Julia: A fresh approach to numerical computing’, *SIAM Review*, **59**(1), 65–98, (2017).
- [4] Anthony P. Burgard, Priti Pharkya, and Costas D. Maranas, ‘Optknoack: a bilevel programming framework for identifying gene knockout strategies for microbial strain optimization’, *Biotechnology and bioengineering*, **84**(6), 647–657, (2003).
- [5] Laurent Charlin and Richard Zemel, ‘The toronto paper matching system: an automated paper-reviewer assignment system’, in *Proceedings of ICMML Workshop on Peer Reviewing and Publishing Models*, Atlanta, United States, (2013). ICML.
- [6] Stefano Coniglio, Luca G. Gianoli, Edoardo Amaldi, and Antonio Capone, ‘Elastic Traffic Engineering Subject to a Fair Bandwidth Allocation via Bilevel Programming’, *IEEE/ACM Transactions on Networking*, **28**(6), 2407–2420, (2020).
- [7] Stephan Dempe, *Foundations of bilevel programming*, Springer Science & Business Media, New York, United States, 2002.
- [8] Scott DeNegre, *Interdiction and discrete bilevel linear programming*, Ph.D. dissertation, Lehigh University, 2011.
- [9] John P. Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu, ‘Balancing Relevance and Diversity in Online Bipartite Matching via Submodularity’, in *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, pp. 1877–1884, Palo Alto, United States, (2019). AAAI press.
- [10] Iain Dunning, Joey Huchette, and Miles Lubin, ‘Jump: A modeling language for mathematical optimization’, *SIAM Review*, **59**(2), 295–320, (2017).
- [11] Thomas. E. Easterfield, ‘A Combinatorial Algorithm’, *Journal of The London Mathematical Society*, **21**, 219–226, (1946).
- [12] Jack Edmonds, ‘Maximum matching and a polyhedron with 0, 1-vertices’, *Journal of research of the National Bureau of Standards B*, **69**(125-130), 55–56, (1965).
- [13] Daquan Feng, Lu Lu, Yi Yuan-Wu, Geoffrey Ye Li, Gang Feng, and Shaoqian Li, ‘Device-to-device communications underlying cellular networks’, *IEEE Transactions on communications*, **61**(8), 3541–3551, (2013).
- [14] Naveen Garg, Telikepalli Kavitha, Amit Kumar, Kurt Mehlhorn, and Julián Mestre, ‘Assigning papers to referees’, *Algorithmica*, **58**(1), 119–136, (2010).
- [15] Judy Goldsmith and Robert H Sloan, ‘The AI conference paper assignment problem’, in *Proceedings AAAI Workshop on Preference Handling for Artificial Intelligence, Vancouver*, pp. 53–57, Vancouver, Canada, (2007). AAAI Workshop.
- [16] Frank L. Hitchcock, ‘The Distribution of a Product from Several Sources to Numerous Localities’, *Journal of Mathematics and Physics*, **20**(1-4), 224–230, (1941).
- [17] Robert W. Irving, ‘Stable marriage and indifference’, *Discrete Applied Mathematics*, **48**(3), 261–272, (1994).
- [18] Rajendra K Jain, Dah-Ming W Chiu, William R Hawe, et al., ‘A quantitative measure of fairness and discrimination’, *Eastern Research Laboratory*, **21**, 1–37, (1984).
- [19] Siladitya Jana, ‘A history and development of peer-review process’, *Annals of Library and Information Studies*, **66**(4), 152–162, (2019).
- [20] Leonid V. Kantorovitch, ‘On the Translocation of Masses’, *Management Science*, **5**(1), 1–4, (1958).
- [21] Maryam Karimzadehgan and ChengXiang Zhai, ‘Constrained multi-aspect expertise matching for committee review assignment’, in *Proceedings of the 18th ACM conference on Information and knowledge management*, pp. 1697–1700, New York, United States, (2009). Association for Computing Machinery.
- [22] Maryam Karimzadehgan, ChengXiang Zhai, and Geneva Belford, ‘Multi-aspect expertise matching for review assignment’, in *Proceedings of the 17th ACM conference on Information and knowledge management*, pp. 1113–1122, New York, United States, (2008). Association for Computing Machinery.
- [23] Julien Lesca and Patrice Perny, ‘LP Solvable Models for Multiagent Fair Allocation Problems.’, in *ECAI*, volume 2010, pp. 393–398, Lisboa, Portugal, (2010). IOS Press.
- [24] Zhi-Quan Luo, Jong-Shi Pang, and Daniel Ralph, *Mathematical programs with equilibrium constraints*, Cambridge University Press, 1996.
- [25] David G. McVitie and Leslie B Wilson, ‘The stable marriage problem’, *Communications of the ACM*, **14**(7), 486–490, (1971).
- [26] Nihar B Shah, ‘Challenges, experiments, and computational solutions in peer review’, *Communications of the ACM*, **65**(6), 76–87, (2022).
- [27] Kazunari Sugiyama and Min-Yen Kan, ‘Scholarly paper recommendation via user’s recent research interests’, in *Proceedings of the 10th annual joint conference on Digital libraries*, pp. 29–38, New York, United States, (2010). Association for Computing Machinery.
- [28] Robert L. Thorndike, ‘The problem of classification of personnel’, *Psychometrika*, **15**, 215–235, (1950).
- [29] Heinrich Von Stackelberg, *The theory of the market economy*, William Hodge, 1952.
- [30] Li Wang, Huaqing Wu, Wei Wang, and Kwang-Cheng Chen, ‘Socially enabled wireless networks: resource allocation via bipartite graph matching’, *IEEE Communications Magazine*, **53**(10), 128–135, (2015).
- [31] Jiuping Xu, Yan Tu, and Ziqiang Zeng, ‘Bilevel optimization of regional water resources allocation problem under fuzzy random environment’, *Journal of Water Resources Planning and Management*, **139**(3), 246–264, (2013).