

On Good and Fair Paper-Reviewer Assignment

Cheng Long, Raymond Chi-Wing Wong, Yu Peng, Liangliang Ye
 The Hong Kong University of Science and Technology
 Kowloon, Hong Kong
 {clong,raywong,gracepy,llye}@cse.ust.hk

Abstract—Peer review has become the most common practice for judging papers submitted to a conference for decades. An extremely important task involved in peer review is to assign submitted papers to reviewers with appropriate expertise which is referred to as *paper-reviewer assignment*. In this paper, we study the paper-reviewer assignment problem from both the goodness aspect and the fairness aspect. For the goodness aspect, we propose to maximize the *topic coverage* of the paper-reviewer assignment. This objective is new and the problem based on this objective is shown to be NP-hard. To solve this problem efficiently, we design an approximate algorithm which gives a $\frac{1}{3}$ -approximation. For the fairness aspect, we perform a detailed study on conflict-of-interest (COI) types and discuss several issues related to using COI, which, we hope, can raise some open discussions among researchers on the COI study. Finally, we conducted experiments on real datasets which verified the effectiveness of our algorithm and also revealed some interesting results of COI.

I. INTRODUCTION

Peer review has become the most common practice for judging papers submitted to a conference for decades [1]. One important task involved in peer review is to assign reviewers from the program committee to submitted papers. This task is referred to as *paper-reviewer assignment* (PRA). Since the number of submitted papers and the size of the program committee in a conference is large (e.g., in ICDM 2012, there were 756 submissions and 234 PCs¹), *manual* PRA is not feasible. Instead, people resort to *automatic* PRA methods.

The first work about automatic PRA was due to Dumais and Nielsen [2], where PRA is regarded as an *information retrieval* problem. Specifically, a paper is used as a *query* and each reviewer is represented by a text document (e.g., the expertise statement provided by the reviewer or simply the publications of the reviewer). The problem is to retrieve a certain number of reviewers who are the most *relevant* to the paper. Following the same idea of [2], some other methods for automatic PRA have been proposed [3, 1, 4-7]. These retrieval-based methods differ from each other by using different information retrieval techniques: Latent Semantic Indexing (LSI) [2, 3], Vector Space Model [1, 4], Topic Model [5], Mixture Language Model [7] and other models [6].

A common drawback of these retrieval-based methods is that the retrieval process has to be done for each paper *independently* such that an assignment between papers and reviewers can be constructed, which, however, introduces several problems. First, it might happen that some popular

reviewers are assigned with excessive papers while some other reviewers with few or even no papers. Second, even though they can avoid the first problem by incorporating a hard constraint on the reviewers' workload, they still suffer a lot. One example is that the constructed assignment is sensitive to the order of papers processed since the papers processed earlier can be assigned with relevant reviewers but the papers processed later may be assigned with irrelevant reviewers due to the hard constraint introduced. Another example is that the process is heuristic-based without any optimized objectives.

To avoid the above drawback of the retrieval-based methods, more recent studies regard PRA as a *matching* problem between a paper set and a reviewer set [8-15]. In this way, all papers *as a batch* are assigned to the reviewers. In general, matching-based methods has the following two-phase framework. In the first phase, a *weighted* bipartite graph between a paper set and a reviewer set is constructed. The weight of an edge between a paper and a reviewer is usually set to a value denoting the *relevance* between the paper and the reviewer. In the second phase, based on the constructed bipartite graph, a matching, which is used to construct the final paper-reviewer assignment, is found such that several constraints (e.g., each paper is assigned to a certain number of reviewers and each reviewer is not assigned with excessive papers) are satisfied and one appropriate objective function is optimized. Under this framework, most of the matching-based methods [16, 17, 1, 14] share the following ideas. Firstly, the relevance used in the first phase is defined based on *topics*. Specifically, each paper (and also each reviewer) is associated with a set of *topics* from a topic domain (pre-specified or learned). The relevance between a paper and a reviewer is defined based on the number of *common* topics shared by the paper and the reviewer. Secondly, the matching described in the second phase usually corresponds to the *maximum weight matching* (or its variant) based on the bipartite graph. With these ideas, these methods assign papers to reviewers so the *total weight* of a matching is maximized. This idea is intuitive, but it does not directly maximize the *coverage* of the topics of the papers *covered* by the assigned reviewers, which could be quite problematic in some cases. To illustrate, let us work through a toy example in detail.

Example 1 (Motivating Example): In Figure 1(a), we have 1 paper p_1 , a topic domain containing 5 topics, namely t_1, \dots, t_5 , and 4 reviewers, namely r_1, \dots, r_4 . A link between a paper and a topic denotes that the paper covers the topic, and a link between a reviewer and a topic denotes that the

¹<http://www.cs.uvm.edu/~icdm/Slides/ICDM12-CommunityMtnng.pdf>




Fig. 1. A motivating example

reviewer has his/her expertise on the topic. Assume that we want to assign 2 reviewers to paper p_1 . Existing methods such as [17] would first construct a bipartite graph as shown in Figure 1(b). The weight of edge (p_1, r_1) in the graph is equal to 3 since paper p_1 and reviewer r_1 share 3 topics (i.e., t_1, t_2 and t_3). Similarly, we have the weights of other edges as shown in the graph. Then, the assignment $A_1 = \{(p_1, r_1), (p_1, r_2)\}$ (which means that paper p_1 is assigned to reviewers r_1 and r_2) is computed since it corresponds to the *maximum weight matching* subject to the constraint that p_1 is assigned to 2 reviewers. Unfortunately, assignment A_1 is not good enough since all the assigned reviewers (i.e., r_1 and r_2) have their expertise on only 3 *distinct* topics of p_1 (i.e., t_1, t_2 and t_3) while there are still 2 topics of p_1 (i.e., t_4 and t_5) left uncovered. This might result in a fairly poor judgement on p_1 . A better choice is to assign r_1 (or r_2) and r_4 to p_1 since they together have their expertise in more distinct topics of p_1 (because the number of distinct topics covered is 5) and thus the resulting judgement is probably more qualitative. \square

Motivated by the above observation, in this paper, we propose a new problem called *Maximum Topic Coverage Paper-Reviewer Assignment* (MaxTC-PRA), which assigns papers to reviewers such that the total number of *distinct* topics of papers that are covered by the assigned reviewers is maximized and the following three constraints are satisfied: (1) *Paper Demand Constraint*: Each paper is reviewed by a certain number of reviewers, (2) *Reviewer Workload Constraint*: Each reviewer reviews at most a certain number of papers, and (3) *COI Constraint*: There exists no conflict-of-interest (COI) between the authors of each paper and the assigned reviewers.

Compared with the existing matching-based methods such as those in [14, 16, 17], our MaxTC-PRA gives a *broader* coverage of the topics of the papers by only counting the *distinct* topics covered by the assigned papers. This makes a big difference since it is desirable that the topics covered by the assigned reviewers together should be as *broad* as possible so that multiple aspects of the paper can be judged qualitatively.

The new objective of MaxTC-PRA is more natural and desirable than those of the existing methods but it makes the problem harder. Firstly, the existing matching-based methods cannot be adapted to our problem. This is because our MaxTC-PRA problem cannot be fit in the framework of matching-based methods since the weight between a paper and a reviewer is correlated with the weight between the same paper and another reviewer (in terms of the topic coverage). Secondly, we prove that MaxTC-PRA is NP-hard. Although solving MaxTC-PRA optimally is difficult, we propose a greedy algorithm which iteratively assigns a paper to a reviewer such that the *marginal gain* of the objective is the

greatest. Interestingly, this greedy algorithm provides a $\frac{1}{3}$ -factor approximation for MaxTC-PRA.

In addition, we discuss in detail three issues of using COI in PRA, namely, (1) what types of COI (e.g., the co-author relationship is a type of COI) should be used in PRA; (2) whether it is reliable to use the COIs specified by the authors and/or the reviewers only; and (3) whether it is always good to use as many COIs as possible. We hope these issues could open some more thorough and useful discussions among researchers on the COI study.

We summarize our main contributions as follows.

- To the best of our knowledge, we are the *first* to propose the MaxTC-PRA problem, which assigns reviewers to papers for maximizing the total number of *distinct* topics of papers covered by the assigned reviewers. MaxTC-PRA is superior over existing methods since it gives a *broader* coverage of the topics of the papers.
- We prove that MaxTC-PRA is NP-hard and design an approximate algorithm for MaxTC-PRA, which provides a $\frac{1}{3}$ -factor approximation.
- We discuss three issues related to using COI in PRA, which we hope, could serve as the initiative efforts on the COI study.
- We conducted a comprehensive empirical study which verifies the superiority of our MaxTC-PRA and the effectiveness of our approximate algorithm. Besides, we observed some interesting findings of the effects of different COI types on the paper-reviewer assignment.

The remainder of this paper is organized as follows. We define and solve the MaxTC-PRA problem in Section II and in Section III, respectively. Then, we discuss three issues related to COI in Section IV and provide the empirical study in Section V. We review the related work in Section VI and conclude the paper in Section VII.

II. PROBLEM DEFINITION

Let $P = \{p_1, p_2, \dots, p_n\}$ be a set of n papers and $R = \{r_1, r_2, \dots, r_m\}$ be a set of m reviewers. Let $T = \{t_1, t_2, \dots, t_k\}$ be the topic domain whose collection will be described in Section V-A. Each paper p_i (reviewer r_j) is captured by a set of topics from the topic domain T , denoted by $T(p_i)$ ($T(r_j)$) for each $i \in [1, n]$ ($j \in [1, m]$). For example, in Example 1, $P = \{p_1\}$, $R = \{r_1, r_2, r_3, r_4\}$, $T = \{t_1, t_2, \dots, t_5\}$, $T(p_1) = \{t_1, t_2, \dots, t_5\}$ and $T(r_1) = \{t_1, t_2, t_3\}$.

Given a paper $p_i \in P$ and a reviewer $r_j \in R$, (p_i, r_j) is said to be a *match* if p_i is assigned to r_j . Given a match $M = (p_i, r_j)$, we denote the paper (reviewer) involved in M by $M.p$ ($M.r$) which is p_i (r_j).

Any possible set of matches corresponds to an *assignment* between the paper set P and the reviewer set R . Given an assignment $A \subseteq P \times R$, we denote by $A(p_i)$ the set of matches in A that involve paper $p_i \in P$, i.e., $A(p_i) = \{M | M \in A, M.p = p_i\}$. Similarly, we define $A(r_j) = \{M | M \in A, M.r = r_j\}$ for each $r_j \in R$. For example, in Example 1, $A = \{(p_1, r_1), (p_1, r_4)\}$ denotes an

assignment. Besides, we have $A(p_1) = \{(p_1, r_1), (p_1, r_4)\}$ and $A(r_1) = \{(p_1, r_1)\}$.

In a typical PRA scenario, each paper should be assigned to a certain number τ_p of reviewers for cross-verification consideration (e.g., in ICDM, each paper is assigned to 3 reviewers), which we call the *Paper Demand Constraint*, and each reviewer should not be assigned with more than a certain number τ_r of papers for workload consideration, which we call the *Reviewer Workload Constraint*. Here, τ_p and τ_r are two positive integers at least 1.

Besides, we have to take COIs into consideration in PRA, which we call the *COI Constraint*. Specifically, a paper should not be assigned to a reviewer if the authors of the paper have COI with the reviewer. We can always capture the COIs by a set \mathcal{C} which contains all pairs of papers and reviewers in the form of (p_i, r_j) such that p_i cannot be assigned to r_j . We will discuss the issues related to COI in Section IV in detail.

Let A be an assignment between the paper set P and the reviewer set R . Let $\sigma_i(A)$ denote the number of *distinct* topics of paper p_i , which are covered by the assigned reviewers wrt A . For instance, in Example 1, consider the assignment $A_1 = \{(p_1, r_1), (p_1, r_2)\}$. We have $\sigma_1(A_1) = 3$ since 3 distinct topics of p_1 are covered by the assigned reviewers in A_1 . Similarly, for the assignment $A_2 = \{(p_1, r_1), (p_1, r_4)\}$, we have $\sigma_1(A_2) = 5$. Then, we define $\sigma(A)$ to be the total number of *distinct* topics of all papers covered by the assigned reviewers wrt A , i.e., $\sigma(A) = \sum_{i=1}^n \sigma_i(A)$.

Problem Statement. Formally, we present the MaxTC-PRA problem as follows.

Problem 1: Given a set P of n papers, namely p_1, p_2, \dots, p_n , a set R of m reviewers, namely r_1, r_2, \dots, r_m , two integers τ_p and τ_r , and a COI set \mathcal{C} , the **MaxTC-PRA problem** is to find the assignment A that satisfies

- **Paper Demand Constraint:** Each paper is assigned to τ_p reviewers, i.e., for each $p_i \in P$, $|A(p_i)| = \tau_p$,
- **Reviewer Workload Constraint:** Each reviewer reviews at most τ_r papers, i.e., for each $r_j \in R$, $|A(r_j)| \leq \tau_r$,
- **COI Constraint:** $A \cap \mathcal{C} = \emptyset$,

and maximizes $\sigma(A)$. \square

Consider Example 1. Suppose $\tau_p = 2$, $\tau_r = 1$ and $\mathcal{C} = \{(p_1, r_1)\}$. Then, the solution of the MaxTC-PRA problem is $A = \{(p_1, r_2), (p_1, r_4)\}$ with the greatest value of $\sigma(A) = 5$.

Note that the definition of MaxTC-PRA in Definition 1 does not guarantee that there always exists a solution. For example, when $n \cdot \tau_p > m \cdot \tau_r$, the Paper Demand Constraint and the Reviewer Workload Constraint cannot be satisfied simultaneously. Even when $n \cdot \tau_p \leq m \cdot \tau_r$, if the COI set is large (an extreme case is that all pairs of papers and reviewers cannot be made into matches), it is possible that there exist no assignments that satisfy all the three constraints. Motivated by these considerations, we replace the Paper Demand Constraint with a *relaxed* one: we only require that each paper is reviewed at most τ_p reviewers. The MaxTC-PRA problem adopting the relaxed Paper Demand Constraint enjoys two benefits. First, it always guarantee an existence of a solution. Second, in the

case that there are enough reviewers and/or the COI set is not that large (which is the common case in practice), the optimal solution of the MaxTC-PRA problem with the relaxed constraint could be made exactly the same as that of the original version of MaxTC-PRA either *automatically* (because of the objective function $\sigma(\cdot)$ since including more matches in the assignment usually increases the objective function) or *manually* by post-assigning those papers that have not been assigned to exactly τ_p reviewers to the reviewers who are assigned with fewer than τ_r papers. Thus, in the following, we focus on the MaxTC-PRA problem with the relaxed Paper Demand Constraint only.

Intractability. Unfortunately, it turns out that the MaxTC-PRA problem is NP-hard.

Lemma 1: The MaxTC-PRA problem is NP-hard. \square

Proof: The proof could be found in the full version of this paper [18]. \blacksquare

III. SOLUTION OF MAXTC-PRA

In this section, we develop an approximate algorithm called *Greedy* for MaxTC-PRA.

Let A be the assignment to be returned by the algorithm. *Greedy* initializes A to be \emptyset and thus all the three constraints of the MaxTC-PRA problem are satisfied at the right beginning (note that we consider the relaxed Paper Demand Constraint in the problem). Let U be a set of all possible matches to be used in the algorithm. It initializes U to $P \times R - \mathcal{C}$, which excludes all pairs of papers and reviewers in \mathcal{C} for the COI Constraint. During the execution, it maintains U such that for each match $M \in U$, the relaxed Paper Demand Constraint and the Reviewer Workload Constraint are satisfied by $A \cup \{M\}$ (where A is being updated during the execution) by removing all those matches from U that violate this condition. Note that any match in U that violates this condition cannot be included in A (or any super-set of A) subject to the constraints of MaxTC-PRA. Besides, note that there is no need to do any removal operation on U initially since the initial content of U (i.e., $P \times R - \mathcal{C}$) contains all possible matches to be used in the algorithm. Then, it iteratively augments A with the match $M \in U$ such that the marginal gain (in terms of the objective function) of adding M into A is maximized. Note that all the three constraints are satisfied if one of the matches in U is inserted into A . The process stops when U is empty. Since A satisfies all the constraints at the beginning and augmenting A in each iteration does not break these constraints, we know that the final assignment A satisfies all the constraints of MaxTC-PRA. We present the *Greedy* algorithm in Algorithm 1.

Approximation Quality Analysis. We prove that our *Greedy* algorithm gives a $\frac{1}{3}$ -factor approximation for the MaxTC-PRA problem. The details could be found in [18].

IV. CONFLICT-OF-INTEREST

A *conflict-of-interest* (COI) between an author of a paper and a reviewer means that the reviewer has an a-priori bias for/against the paper. Therefore, avoiding different types of

Algorithm 1 The *Greedy* algorithm

Input: A paper set $P = \{p_1, p_2, \dots, p_n\}$; A reviewer set $R = \{r_1, r_2, \dots, r_k\}$; Parameters τ_p , τ_r and a COI set \mathcal{C}

Output: An approximate solution of the MaxTC-PRA problem

- 1: $A \leftarrow \emptyset$
 - 2: $U \leftarrow P \times R - \mathcal{C}$
 - 3: **while** U is non-empty **do**
 - 4: pick the match $M \in U$ such that $\sigma(A \cup \{M\}) - \sigma(A)$ is maximized
 - 5: $A \leftarrow A \cup \{M\}$; $U \leftarrow U - \{M\}$
 - 6: remove each match $M' \in U$ such that the relaxed Paper Demand Constraint or the Reviewer Workload Constraint is not satisfied by $A \cup \{M'\}$
 - 7: **return** A
-

COI when assigning papers to reviewers is critical for unbiased judgement of papers. In this part, we discuss 3 issues related to the problem of using COI for a PRA task.

Issue 1: What types of author-reviewer relationship should be considered as COI types?

We do a literature study first. Some author-reviewer relationships that are well-recognized as COI types by popular conference management systems (e.g., Microsoft Research’s CMT tool (<http://msrcmt.research.microsoft.com/cmt/>), EasyChair (<http://www.easychair.org/>) and CyberChair (<http://www.borbala.com/cyberchair/>)) include the *co-author relationship* (the author and the reviewer have co-authored some papers), the *colleague relationship* (the author and the reviewer have worked/collaborated at the same affiliations), and the *advisor-advisee* relationship (the author was/is the thesis advisor of the reviewer or vice versa).

We note that all these COI types share a common feature that a reviewer who has a COI with an author will be biased for the paper of the author. In other words, they are used to avoid *false positives* of paper judgements. In contrast, no COI types have been proposed for the case that a reviewer is biased against the paper of the author. In this paper, for the first time, we identify such a COI type called the *competitor relationship*.

An author who submitted a paper and a reviewer are said to be in a *competitor relationship* with each other if the reviewer has also submitted a paper in the conference and the fraction of the common topics of these two papers is above a pre-specified threshold δ . Note that it is a common case that the program committee members (i.e., reviewers) of a conference would usually submit papers to the conference as well. As a result, it is not uncommon that a reviewer is assigned with a paper which has quite similar topics as his/her own paper submitted to the same conference (e.g., the two papers appear in the same area track or even study the same problem). In this case, it is hard to assume that the reviewer would judge the paper objectively since s/he has his/her own paper which is competitive with the paper being considered. Thus, in order to avoid this potential un-biased factor for safety, it is reasonable to consider the competitor relationship as a COI type. Clearly, the competitor relationship could be used to avoid *false negatives* of paper judgements.

In this paper, we study 4 types of author-reviewer relation-

ships as COI types, namely, the co-author relationship, the colleague relationship, the advisor-advisee relationship and the competitor relationship. The first three are existing COI types and the last one is newly-proposed in this paper.

Issue 2: Is it reliable to use the COI information specified by the authors and/or the reviewers only?

With the growing of the number of paper submissions, the number of reviewers (i.e., program committee members) in a conference could simply be several hundreds (e.g., ICDM 2012 involves 234 reviewers), in which case, the task of specifying COIs manually by authors and/or reviewers becomes time-consuming and it could happen often that some COIs are left un-specified simply because of the tedious aspect of this task. In addition, what should not be ignored is that there might exist some bad-citizens who would cheat the system by either leaving some COIs un-specified or specifying some fake COIs intentionally. The above two considerations trigger us to come up with the idea of specifying the COIs *automatically* which are then used as a complementary source.

In this paper, we mine multiple sources from the Web for detecting the aforementioned 4 types of COIs automatically. More details will be discussed in Section V.

Issue 3: Is it always good to use as many COIs as possible?

To answer this question, we should understand the effects of COIs on the PRA task. Without doubt, COIs make the PRA task fairer. But, we should also notice that a COI might prevent a paper being assigned to the reviewer who have much expertise on the topics of the paper and this might degrade the *goodness* (i.e., the topic coverage) of the resulting assignment. That is, there is a trade-off between the goodness and the fairness of the paper-reviewer assignment.

In this paper, we study the effect of each of the above COI types on the goodness of the assignment returned by our *Greedy* algorithm to see to what extent this COI type degrades the goodness of the assignment. Details will be discussed in Section V-C. This result can be served as a reference to researchers to understand the significance of each COI type on the paper-reviewer assignment in order to determine whether some of the COI types should be adopted finally in the field.

V. EXPERIMENT

A. Experimental Setup

Datasets. We collected all the papers published in KDD from year 2006 to year 2010 as the submitted papers. We have 496 papers in total. We collected all the program committees (PCs) of ICDM 2010 and KDD 2010 as the reviewers (since ICDM and KDD correspond to the two major prestigious conferences in the data mining community and we expect that the program committees of these two conferences together cover the topics of data mining well). We have 550 reviewers in total. We collected all the subject areas specified in KDD 2011 for paper submission as the topic domain. We have 49 topics in total.

Topic Extraction. Following the convention [4, 3, 6], we retrieve the topics covered by a paper (reviewer) by using the




Fig. 2. Topic coverage: *Greedy* vs. ILP

TF-IDF weighted vector space model [19]. The details could be found in [18].

COI Collection. Details are included in [18].

Algorithms. We considered two algorithms in our experiments, namely, *Greedy* and ILP. *Greedy* is the approximate algorithm proposed in this paper and ILP is the art-of-the-state [17] among all algorithms that consider *topic coverage* for paper-reviewer assignment. ILP is a matching-based method, which sets the weight of the edge between a paper and a reviewer as the number of common topics shared by the paper and the reviewer in Phase 1 and computes the maximum weight matching with an *integer linear program* in Phase 2.

Configurations. 4 factors are studied in our experiments, namely, the number of papers (i.e., $|P|$), the number of reviewers (i.e., $|R|$), the parameter of the (relaxed) Paper Demand Constraint (i.e., τ_p) and the parameter of the Reviewer Workload Constraint (i.e., τ_r). The values used for $|P|$ are 100, 200, 300, 400 and **496**, where the value in bold font is used by default. The values used for $|R|$ are 150, 250, **350**, 450 and 550. The values used for τ_p are 1, 2 and **3** and those used for τ_r are **5**, 6, 7 and 8.

B. Verification of MaxTC-PRA & *Greedy*

We used a measure called *topic coverage*, which is defined to be the average percentage of the distinct topics of papers covered by the assigned reviewers. Clearly, a better algorithm gives a higher topic coverage. We used all COI types in this set of experiments. The results when varying $|P|$, $|R|$ and τ_p are shown in Figure 2(a), (b) and (c), respectively.

Varying $|P|$. We have the following observations. First, the topic coverage of the assignment returned by *Greedy* is consistently higher than that returned by ILP. In fact, the topic coverage of *Greedy* is 15% larger than that of ILP. Second, the topic coverage of *Greedy* is more than 90% in all settings, which also implies that the empirical approximation factor of *Greedy* (which is at least 0.90) is much better the theoretical one (which is $\frac{1}{3}$). Third, the topic coverage of both *Greedy* and ILP decreases slightly when $|P|$ increases. This could be explained as follows. When the number of papers increases, the total number of distinct topics of the papers is larger. Since the number of reviewers is kept unchanged, it is likely that the topic coverage would decrease.

Varying $|R|$. Again, there is a clear gap between the topic coverage of *Greedy* and that of ILP.

Varying τ_p . Still, we note that the topic coverage of *Greedy* is consistently higher than that of ILP except for the case of $\tau_p = 1$. Specifically, when $\tau_p = 1$ (i.e., each paper is reviewed

by only 1 reviewer), the topic coverage of both algorithms is around 65%, which is not high enough for obtaining a qualitative paper judgement. This implies that setting $\tau_p = 1$ is not a good choice in real applications. In contrast, when $\tau_p = 3$ which is a common practice in real applications, the topic coverage of *Greedy* is more than 90%, which is about 15% higher than that of ILP.

Varying τ_r . The trends are similar to those of varying $|R|$. This is because increasing τ_r is essentially “adding” (more or less similar) reviewers into the reviewer set (with the original Reviewer Workload Constraint). Bearing this in mind, in the following, we omit the experimental results of varying τ_r .

C. Studies of the Effect of COI

1) *On the Fairness of PRA:* We evaluate the effect of each COI type on the fairness of PRA by comparing the resulting assignments with and without using this COI type. Specifically, we collected the percentage of matches in the assignment that violate the COI type being considered if this COI type is not used for PRA. The results of varying $|P|$, $|R|$ and τ_p are shown in Figure 3(a), (b) and (c), respectively.

Varying $|P|$. We have the following observations. First, surprisingly yet interestingly, we find that COI4 (i.e., the competitor relationship) has the strongest effect on PRA. Specifically, around 15-17% of the matches in the resulting assignment violate COI4 if it is not used as a COI type. This might be explained by the fact that the reviewers of a conference usually have their own papers submitted in the conference and at the same time, the papers (by others) that share some common topics with these papers would probably be assigned to these reviewers since the reviewers usually have much expertise on these common topics. This results in a situation where many author-reviewer pairs violate COI4. Second, the (decreasing) ordering of the remaining COI types by their effects on the assignment is COI1 (i.e., the co-author relationship), COI2 (i.e., the colleague relationship), COI3 (i.e., the advisor-advisee relationship). Third, the percentage of matches that violate a COI (any COI type) is more than 22%, which provides a clear argument that the assignment made without using any COI types could be fairly unfair.

Varying $|R|$ and τ_p . The results provide us similar clues and thus the discussion on these results are omitted here.

2) *On the Goodness of PRA:* We evaluate the effects of the COI types on the quality (i.e., topic coverage) of the assignment returned by our *Greedy* algorithm. The results of varying $|P|$, $|R|$ and τ_p are shown in Figure 4(a), (b) and (c), respectively.




Fig. 3. Effects of COI on the fairness of the assignment




Fig. 4. Effects of COI on the goodness of the assignment

Varying $|P|$. In Figure 4(a) where we show the effect of the whole set of 4 COI types. We observe that using the whole set of COI types degrades the topic coverage by 2%-3% only compared to the case where no COIs are used. This is a good news since it implies that the side-effect (in terms of the quality) of using all COI types studied in this paper is negligible. The results of the effect corresponding to each COI type only are not shown here since they are quite minor (e.g., ranging from 0.1% to 1%). Instead, we show the results of the overall effect of the whole set of 4 COI types.

The decreasing trends of the curves when $|P|$ increases could be explained in the way as we did for Figure 2(a).

Varying $|R|$ and τ_P . The results are similar to the case of varying $|P|$ except for the trends of the curves in Figure 4(b) and (c) which could be explained similarly as we did for Figure 2(b) and (c), respectively.

VI. RELATED WORK

In the past two decades, a bulk of studies studied the problem of (automatic) paper-reviewer assignment (PRA), which we categorize into two branches, namely the *retrieval-based* methods [2, 6, 4, 1, 5, 7, 20, 21, 3] which regard each paper as a query to retrieve the *relevant* experts (i.e., reviewers) and the *matching-based* methods [8-17] which compute a *matching* based on the bipartite graph between the paper set and the reviewer set. Refer to [18] for more details.

VII. CONCLUSION

In this paper, we proposed a new problem called MaxTC-PRA, which favors the diversification of the topics of the papers covered by the assigned reviewers. We showed that this problem is NP-hard and proposed a $\frac{1}{3}$ -approximate algorithm for this problem. We also discussed several issues related to using COI in PRA, which we hope, could serve as an initiative effort on COI study. We conducted experiments on real datasets which verified our algorithm and revealed some interesting results of COI.

Acknowledgements: We are grateful to the anonymous reviewers for their constructive comments on this paper. The research is supported by grant FSGRF12EG50.

REFERENCES

- [1] H. K. Biswas and M. Hasan, "Using publications and domain knowledge to build research profiles: An application in automatic reviewer assignment," in *ICICT*, 2007.
- [2] S. T. Dumais and J. Nielsen, "Automating the assignment of submitted manuscripts to reviewers," in *SIGIR*. ACM, 1992.
- [3] J. Zablocki and R. Lee, "Auto-assign: An implementation to assign reviewers using topic comparison in start."
- [4] S. Hettich and M. J. Pazzani, "Mining for proposal reviewers: lessons learned at the national science foundation," in *SIGKDD*. ACM, 2006.
- [5] D. Mimno and A. McCallum, "Expertise modeling for matching papers with reviewers," in *SIGKDD*. ACM, 2007.
- [6] C. Basu, H. Hirsh, W. Cohen, and C. Nevill-Manning, "Recommending papers by mining the web," in *Proceedings of the IJCAI99 Workshop on Learning about Users*, 1999.
- [7] M. Karimzadehgan, C. Zhai, and G. Belford, "Multi-aspect expertise matching for review assignment," in *CIKM*. ACM, 2008.
- [8] D. Hartvigsen, J. C. Wei, and R. Czuchlewski, "The conference paper-reviewer assignment problem*," *Decision Sciences*, vol. 30, no. 3, pp. 865-876, 1999.
- [9] S. Benferhat and J. Lang, "Conference paper assignment," *International Journal of Intelligent Systems*, 2001.
- [10] J. Merelo-Guervós and P. Castillo-Valdivieso, "Conference paper assignment using a combined greedy/evolutionary algorithm," in *Parallel Problem Solving from Nature-PPSN VIII*. Springer, 2004, pp. 602-611.
- [11] C. J. Taylor, "On the optimal assignment of conference papers to reviewers," Computer and Information Science Department, University of Pennsylvania, Tech. Rep. MS-CIS-08-30, 2008.
- [12] D. Conry, Y. Koren, and N. Ramakrishnan, "Recommender systems for the conference paper assignment problem," in *RecSys*. ACM, 2009, pp. 357-360.
- [13] N. Garg, T. Kavitha, A. Kumar, K. Mehlhorn, and J. Mestre, "Assigning papers to referees," *Algorithmica*, 2010.
- [14] W. Tang, J. Tang, and C. Tan, "Expertise matching via constraint-based optimization," in *WI-IAT*, 2010.
- [15] L. Charlin, R. S. Zemel, and C. Boutilier, "A framework for optimizing paper matching," *arXiv preprint:1202.3706*, 2012.
- [16] S. Ferilli, N. D. Mauro, T. Basile, F. Esposito, and M. Biba, "Automatic topics identification for reviewer assignment," *AAAI*, pp. 721-730, 2006.
- [17] M. Karimzadehgan and C. Zhai, "Constrained multi-aspect expertise matching for committee review assignment," in *CIKM*. ACM, 2009, pp. 1697-1700.
- [18] C. Long, R. C.-W. Wong, Y. Peng, and L. Ye, "On good and fair paper-reviewer assignment (technical report)," in <http://www.cse.ust.hk/~raywong/paper/paperAssign-technical.pdf>, 2013.
- [19] G. Salton and M. J. McGill, *Introduction to modern information retrieval*. Facet publishing, 1986.
- [20] M. A. Rodriguez and J. Bollen, "An algorithm to determine peer-reviewers," in *CIKM*. ACM, 2008, pp. 319-328.
- [21] P. Serdyukov, H. Rode, and D. Hiemstra, "Modeling multi-step relevance propagation for expert finding," in *CIKM*. ACM, 2008, pp. 1133-1142.