

Supplementary Material for “Group Fairness in Peer Review”

A Execution of CoBRA in a Toy Example

		Rounds					Reviewers		1		1		2	
		1	2	3	4	5	1	2	1	2	1	2		
Rounds	1 :	p_3	p_2	p_4			1 :	p_3, p_2, p_4	p_3, p_6	p_2, p_4	p_6	p_2, p_4	p_4, p_5	
Reviewers	2 :	p_1	p_3		p_6		2 :	p_1, p_3, p_6	p_1, p_3, p_6			p_1, p_3, p_6		
	3 :	p_2	p_1			p_4	3 :	p_2, p_1, p_4	p_2, p_1, p_4			p_2, p_1, p_4		
	4 :			p_1		p_5	4 :	p_1, p_5, p_6	p_1, p_5, p_6			p_1, p_5, p_6		
	5 :				p_2	p_3	5 :	p_2, p_3	p_2, p_3			p_2, p_3, p_4		
	6 :				p_5		6 :	p_5, p_4	p_5, p_4, p_3			p_5, p_4, p_3		
							Phase 1		Phase 2					

(a) Execution of PRA-TTC

(b) Execution of Filling-Gaps

Figure 1: Execution of CoBRA when $n = 6$, $k_p = k_a = 3$, $\sigma_1 = 2 \succ 3 \succ 4 \succ \dots$, $\sigma_2 = 3 \succ 1 \succ 5 \succ \dots$, $\sigma_3 = 1 \succ 2 \succ 5 \succ \dots$, $\sigma_4 = 1 \succ 3 \succ 5 \succ \dots$, $\sigma_5 = 6 \succ 4 \succ \dots$ and $\sigma_6 = 2 \succ \dots$. On the left table, we see the assignments that are established in each round of PRA-TTC by eliminating cycles. After the execution of PRA-TTC, three papers, p_4, p_5, p_6 are not completely assigned. Thus, $U = \{4, 5, 6\}$ and $L = \{3\}$. On the right table, we see the execution of Filling-Gaps. There is a cycle in the greedy graph which is eliminated at the first round of Phase 1. In Phase 2, where $\vec{p} = (6, 5)$, at the first round, since p_3 is authored by an agent in $U \cup L \setminus \{6\}$, is not reviewed by 6 and is completely assigned, p_3 is assigned to 6 while it is removed from 1 in which p_6 is now assigned. At the second round, since p_4 is authored by an agent in $U \cup L \setminus \{5\}$, is not reviewed by 5 and is completely assigned, p_4 is assigned to 5 while it is removed from 1 in which p_5 is now assigned.

Figure 1 shows an execution of CoBRA at a small instance.

Let us briefly describe how we get this assignment. In the preference graph, we see that 1 has an outgoing edge to agent 2, agent 2 has an outgoing edge to agent 3 and agent 3 has an outgoing edge to agent 1. By eliminating this cycle we assign p_3 to 1, p_2 to 3 and p_1 to 2. By continuing to detect cycles in the preference graph given the preferences, we conclude on the partial assignment shown in the table on the left. For the table on the right, at the first phase of Filling-Gaps, when we create the greedy graph, we see that it consists of 3 nodes, since $U = \{4, 5, 6\}$, 4 has outgoing edges to 5 and 6 (since none of them review p_4), 6 has outgoing edges to 4 and 5, and 5 has no outgoing edge. Hence there is only one cycle and that is why p_6 is assigned to 4 and p_4 is assigned to 6. Then, 4 is moved to L , since p_4 becomes completely assigned. Then, there is no cycle in the greedy graph as while 6 has an outgoing edge to 5, 5 has no outgoing edges. Hence, we proceed to the second phase where there are two non-completely assigned papers, p_5 and p_6 . $\vec{p} = (5, 6)$ as agent 6 reviews p_5 , but 5 does not review p_6 . Then, the algorithm makes sure that p_6 and then p_5 become completely assigned as it is described in the caption of the figure.

B Proof of Lemma 1

Proof. We start by showing that during the execution of PRA-TTC and the execution the first phase of Filling-Gaps, it holds the following lemma.

Lemma 2. *During the execution of PRA-TTC and the execution of the first phase of Filling-Gaps, for each $i \in N$ with $\bar{P}_i \neq \emptyset$, it holds that*

$$|R_i^a| = \sum_{j \in [m^*]} |R_{p_i, j}^p|.$$

Proof. Note that during the execution of PRA-TTC, if an agent i with $\bar{P}_i \neq \emptyset$ is assigned one submission due to the elimination of a cycle, then we know that one of her submissions that is

incompletely assigned is also assigned to an agent that does not review it already. Hence, we see that until \bar{P}_i becomes empty, we have that $|R_i^a| = \sum_{j \in [m^*]} |R_{p_{i,j}}^p|$.

Next, we focus on the execution of Filling-Gaps. We know that any $i \in N$ with $\bar{P}_i \neq \emptyset$ is included in U . In the first phase, the algorithm eliminates cycles in the greedy graph. With similar arguments as in the elimination of cycles in the preference graph, we get that during and after the first phase of Filling-Gaps, it is still true that $|R_i^a| = \sum_{j \in [m^*]} |R_{p_{i,j}}^p|$ for any $i \in U$ with $\bar{P}_i \neq \emptyset$. \square

Next, we need to show that Line 8 of PRA-TTC is valid which is true if $|U| \leq k_p$.

Lemma 3. *PRA-TTC returns $|U| \leq k_p$.*

Proof. For each $i \in U$, from Lemma 2, we know that

$$|R_i^a| = \sum_{j \in [m^*]} |R_{p_{i,j}}^p| < m^* \cdot k_p \leq k_a$$

where the first inequality follows since there exists at least one submission of i that is assigned to less than k_p agents. Hence, we get that each $i \in U$ can review more submissions, when PRA-TTC terminates. Now, suppose for contradiction that at the last iteration of PRA-TTC, each agent $i \in U$ has an outgoing edge in the preference graph. In this case, we claim that there exists a directed cycle in the preference graph which is a contradiction since PRA-TTC would not have been terminated yet. To see that, note that each outgoing edge of an agent $i \in U$ either goes to another agent $i' \in U$, since i' can review more submissions, or goes to an agent $i' \notin U$ whose all submissions are completely assigned. In the latter case, i' has an outgoing edge to an agent in U by the definition of the preference graph. Thus, starting from any agent in U , we conclude in an agent in U and eventually we would find a cycle. Therefore, we have that there exists an agent $i^* \in U$ that at the last iteration of the algorithm arbitrarily picks her incomplete submission p_{i^*, ℓ^*} and does not have any outgoing edge to any other agent. This means that all the agents that can review more submissions, already review p_{i^*, ℓ^*} . Since all the agents in $U \setminus \{i^*\}$ can review more submissions, we get that all of them are assigned p_{i^*, ℓ^*} . But since p_{i^*, ℓ^*} is not completely assigned, we conclude that $|U \setminus \{i^*\}| < k_p$, which means that $|U| \leq k_p$. \square

We proceed by showing that Lines 11- 12 of Filling-Gaps are valid.

Lemma 4. *When Fillings-Gaps enter the second phase with a non empty U , for each $t \in [|U|]$ and for each $p_{\rho(t), \ell} \in \bar{P}_{\rho(t)}$, it exists a completely assigned submission $p_{i', \ell'}$ with $i' \in U \cup L \setminus \{\rho(t)\}$ that is not reviewed by $\rho(t)$, and it also exists an $i'' \in N$ that reviews $p_{i', \ell'}$, but she does not review $p_{\rho(t), \ell}$.*

Proof. When U is non empty and no more cycles exists in the greedy graph, the algorithm constructs the topological order of the greedy graph, denoted by $\vec{\rho}$.

First, we show the following proposition.

Proposition 1. *For each $t \in [|U|]$, $\rho(t)$ reviews all the incompletely assigned submissions of each $i \in U \setminus \{\rho(1), \dots, \rho(t-1), \rho(t)\}$.*

Proof. Since $\vec{\rho}$ is the topological ordering of the greedy graph, we have that no $i \in U \setminus \{\rho(1), \dots, \rho(t-1), \rho(t)\}$ has an outgoing edge to $\rho(t)$. But from Lemma 2, we get that $\rho(t)$ can review more submissions, since $\rho(t)$ has submissions that are incompletely assigned which means that

$$\sum_{\ell \in [m^*]} |R_{p_{\rho(1), \ell}}^p| < k_p \cdot m^* \leq k_a.$$

Therefore, from the definition of the greedy graph, we get that $\rho(t)$ reviews all the incompletely assigned submissions of each i in $U \setminus \{\rho(1), \dots, \rho(t-1), \rho(t)\}$. \square

Next, we show by induction that for each $t \in [|U|]$ as long as $\bar{P}_{\rho(t)}$ is non-empty and it holds that $|R_{\rho(t)}^a| = \sum_{\ell \in [m^*]} |R_{p_{\rho(t),\ell}}^p|$, there exists a completely assigned submission $p_{i',\ell'}$ of an agent $i' \in U \cup L \setminus \{\rho(t)\}$ that is not reviewed by $\rho(t)$, and $i'' \in N$ that reviews $p_{i',\ell'}$ and does not review $p_{\rho(t),\ell} \in \bar{P}_{\rho(t)}$.

We start with $t = 1$. First, suppose for contradiction that $\rho(1)$ reviews all the submissions of all the agents in $U \cup L \setminus \{\rho(1)\}$. Then, we would have that

$$|R_{\rho(1)}^a| = \sum_{\ell \in [m^*]} |R_{p_{\rho(1),\ell}}^p| = k_p \cdot m^*,$$

where the first equality follows from the assumption that $|R_{\rho(1)}^a| = \sum_{\ell \in [m^*]} |R_{p_{\rho(1),\ell}}^p|$ and the second inequality follows from the facts that $|U \cup L \setminus \{\rho(1)\}| = k_p$ and each agent has m^* submissions. But then we would conclude that all the submissions of $\rho(1)$ are completely assigned since $\rho(1)$ has m^* submissions and each of them should be assigned to k_p reviewers which is a contradiction. Moreover, from Proposition 1 we know that $\rho(1)$ reviews all the incompletely assigned submissions that belongs to some agent $i \in U \setminus \{\rho(1)\}$. Hence, we get that since $\rho(1)$ reviews all the incompletely assigned submissions but she cannot review all the submissions of all agents in $i \in U \cup L \setminus \{\rho(1)\}$, there exists a completely assigned submission $p_{i',\ell'}$ that belongs to some $i' \in U \cup L \setminus \{\rho(1)\}$ and is not reviewed by $\rho(1)$. In addition, since $p_{i',\ell'}$ is reviewed by k_p agents and not from $\rho(1)$, while $p_{\rho(1),\ell} \in \bar{P}_{\rho(1)}$ is reviewed by strictly less than k_p agents, it exists an agent i'' that reviews the former submission but not the latter. It remains to show that during the execution of the 1-th iteration of the for loop in the second phase of Filling-Gaps, it holds that $|R_{\rho(1)}^a| = \sum_{\ell \in [m^*]} |R_{p_{\rho(1),\ell}}^p|$. Note that if every time that the algorithm enters the second while loop of the algorithm, this property is satisfied, then the property remains true at the end of this execution, since as we show above, in this case there are $p_{i',\ell'}$ and i'' with the desired properties, and therefore one incompletely assigned submission of $\rho(1)$ is assigned to a new reviewer and concurrently $\rho(1)$ is assigned a new submission to review. We get that $|R_{\rho(1)}^a| = \sum_{\ell \in [m^*]} |R_{p_{\rho(1),\ell}}^p|$ is true during the execution of the 1-st iteration of the for loop by noticing that from Lemma 2, we know that this is true when we first enter the while loop of the second phase.

Suppose that the hypothesis holds for $t - 1$. Note that from the base case and the hypothesis, at iteration t , all the submissions of each agent in $i' \in L \cup \{\rho(1), \dots, \rho(t-1)\}$ are completely assigned. Thus, any incompletely assigned submission, that does not belong to $\rho(t)$, belongs to some agent $i \in U \setminus \{\rho(1), \dots, \rho(t-1), \rho(t)\}$. But, from Proposition 1 we already know that $\rho(t)$ reviews any such submission. Moreover, we note that $\rho(t)$ cannot review all the submissions of all the agents in $U \cup L \setminus \{\rho(t)\}$. Indeed, if we assume for contradiction that $\rho(t)$ reviews all the submissions of all the agents in $U \cup L \setminus \{\rho(t)\}$, then we have that

$$|R_{\rho(t)}^a| = \sum_{\ell \in [m^*]} |R_{p_{\rho(t),\ell}}^p| = k_p \cdot m^*,$$

where the first inequality follows from the assumption that $|R_{\rho(t)}^a| = \sum_{\ell \in [m^*]} |R_{p_{\rho(t),\ell}}^p|$ and the second follows from the facts that $|U \cup L \setminus \{\rho(t)\}| = k_p$ and each of them has m^* submissions, which would imply that all the submissions of $\rho(t)$ are completely assigned. Hence, we get that since $\rho(t)$ reviews all the incompletely assigned submissions but cannot review all the submissions of all agents in $i \in U \cup L \setminus \{\rho(t)\}$, there exists a completely assigned submission that belongs to some $i' \in U \cup L \setminus \{\rho(t)\}$ and is not reviewed by $\rho(t)$. Moreover, we show that there exists i'' that reviews $p_{i',\ell'}$, but does not review $p_{\rho(t),\ell} \in \bar{P}_{\rho(t)}$. Indeed, since $p_{i',\ell'}$ is reviewed by k_p agents and not from $\rho(t)$, while $p_{\rho(t),\ell}$ is reviewed by strictly less than k_p agents, it exists an agent that reviews the former submission but not the latter. It remains to show that during the execution of the t -th iteration of the for loop in the second phase of Filling-Gaps, it holds that $|R_{\rho(t)}^a| = \sum_{\ell \in [m^*]} |R_{p_{\rho(t),\ell}}^p|$. Note that if when we first enter the while loop of the t -th iteration it is indeed true that $|R_{\rho(t)}^a| = \sum_{\ell \in [m^*]} |R_{p_{\rho(t),\ell}}^p|$, then during the whole execution of the while loop this property remains true, since as we show above, in this case there are $p_{i',\ell'}$ and i'' with the desired properties. From Lemma 2, we know that this is true when we enter the second phase of Filling-Gaps. Before, round t , if $\rho(t)$ is assigned a new submission to review, she is removed one of the old assigned submissions, while none of her incompletely assigned submissions is assigned to any agent. Hence, indeed we have the desired property, when we first enter the t -th round. \square

Dataset	Alg	USW	ESW
CVPR '17	CoBRA	1.644	0.000
	TPMS	1.970	0.000
	PR4A	1.919	0.384
CVPR '18	CoBRA	1.208	0.000
	TPMS	1.586	0.004
	PR4A	1.560	0.731
ICLR '18	CoBRA	0.251	0.015
	TPMS	0.284	0.038
	PR4A	0.278	0.086

Table 2: USW and ESW without subsampling on CVPR 2017 and 2018, and ICLR 2018.

We partition the agents in N , into two parts N_1 and N_2 , where N_1 contains all the agents that do not belong in U that PRA-TTC returns and $N_2 = N \setminus N_1$. We proceed by separately showing that the assignment that CoBRA returns is valid over the agents in N_1 and over the agent in N_2 .

Valid Assignment over the agents in N_1 . Note that in PRA-TTC, each submission is assigned to at most k_p reviewers and therefore, during the execution of PRA-TTC, for each $i \in N$, it holds that $\sum_{j \in [m^*]} |R_{p_i, j}^p| \leq k_p \cdot m^*$. From Lemma 2, since $k_p \cdot m \leq k_a$, we get that in PRA-TTC, each $i \in N_1$ is not assigned more than k_a papers to review until the point where all of her submissions become completely assigned. After that point, an agent may still participate in a cycle as long as she reviews strictly less than k_a submissions. Therefore, when we exit PRA-TTC, each agent in N_1 does not review more than k_a submissions and all her submissions are completely assigned. In Filling-Gaps, from Lemma 4, we get that if an agent in N_1 is assigned a new submission to review, she is removed one of the submissions that she already reviews. Moreover, the assignments of submissions that belong to agents in N_1 do not change. Hence, we can conclude that the assignment that CoBRA returns is valid with respect to the agents in N_1 .

Valid Assignment over the agents in N_2 . From Lemma 2, we get that during the execution of PRA-TTC each agent i that is included in U that PRA-TTC returns reviews less than k_a submissions, since some of her submissions are not completely assigned (which means that $\sum_{\ell \in [m^*]} |R_{p_i, \ell}^p| < k_p \cdot m^*$). From the same lemma, we have that after the execution of the first phase of Filling-Gaps, it holds that $|R_i^a| = \sum_{\ell \in [m^*]} |R_{p_i, \ell}^p|$. Next, we show that this property remains true after the second phase of Filling-Gaps. Indeed, from Lemma 4, we have that during the second phase of Filling-Gaps, if $i \in U$ is assigned a new submission to review without one of her incompletely assigned submissions is assigned to a new reviewer, she is removed one of her assigned submissions; on the other hand, if one of her incompletely assigned submissions is assigned to a new reviewer, she is also assigned to review a new submission. Lastly, from Lemma 4, we conclude that in the second phase of Filling-Gaps, all the submissions become eventually completely assigned since in each iteration of the while loop, an incompletely assigned submission is assigned to one more reviewer. Therefore in the assignment that Filling-Gaps returns, no agent in N_2 reviews more than k_a submissions and all the submissions of the agents in N_2 are completely assigned, which means that the assignment is valid with respect to the agents in N_2 as well. \square

C Supplementary Experiments

In Section 4, we show that TPMS and PR4A often motivate group of authors to deviate and redistribute their submissions among themselves. The size of a deviating groups is also an interesting measure, for evaluating if such a group indeed consists a distinct subcommunity of researchers that has incentives to build its own conference rather than an extremely tiny group of authors that could locally benefit by exchanging their papers for reviewing. In Table 3, we can see the maximum size of a successfully deviating coalition, averaged across 100 runs, together with the standard error. As before, each run is a subsampled dataset of size 100, so these can be interpreted as percentages. It seems that under both TPMS and PR4A across all three datasets, the largest deviating communities are 6-15% of the conference size, which we can indeed reflect the sizes of some of the largest subcommunities at

Dataset	Alg	Largest Deviating Group
CVPR '17	TPMS	6.64 ± 0.77
	PR4A	7.50 ± 0.77
CVPR '18	TPMS	10.54 ± 1.29
	PR4A	11.49 ± 1.49
ICLR '18	TPMS	11.25 ± 1.76
	PR4A	15.01 ± 1.76

Table 3: Largest Size of Deviating Group on CVPR 2017 and 2018, and ICLR 2018.

CVPR and ICLR. Of course, there are deviating groups with smaller size as well which can consist smaller communities.