Supplementary Materials of "On Strategyproof Conference Peer Review"

A Extension to Review Scores

We define our setup when we have review scores instead of rankings here. In this setting, every reviewer gives his ranking as a vector $\pi^{(i)} = (s_{ij})_{j=1}^{n}$, where (s_{ij}) is the score of paper p_j given by reviewer r_i if $p_j \in \mathcal{P}_i$, and $s_{ij} = \perp$ (placeholder) otherwise. We use the definition of strategyproof in this setting, and define group unanimity as below:

Definition 5 (Group Unanimity (GU) for review scores)**.** The pair $(\mathcal{G}, f)$ is said to be group unanimous (GU) if the following condition holds for every possible profile $\boldsymbol{\pi}$: If there is a non-empty set of papers $\mathcal{P}' \subset \mathcal{P}$ such that for every triplet of reviewer and papers (r_i, p_x, p_y) such that $p_x, p_y \in \mathcal{P}_i, p_x \in \mathcal{P}', p_y \in \mathcal{P}\backslash\mathcal{P}'$ we have $s_{ix} > s_{iy}$, then $f(\boldsymbol{\pi})$ must have $p_x \succ p_y$ for every pair of papers $p_x \in \mathcal{P}'$ and $p_y \in \mathcal{P}\backslash\mathcal{P}'$ such that at least one reviewer has reviewed both p_x and p_y.

Theorem 1 still holds with exactly the same proof (cf. Appendex B.2).

B Proofs

In the appendix, we provide the missing proofs of all the results from the main paper.

B.1 Proof of Proposition 1

Proof. Suppose $(\mathcal{G}, f)$ is PU, and $\mathcal{P}' \subset \mathcal{P}$ satisfies that every reviewer ranks the papers she reviewed from $\mathcal{P}'$ higher than those she reviewed from $\mathcal{P}\backslash\mathcal{P}'$. Now for every $p_x \in \mathcal{P}'$ and $p_y \in \mathcal{P}\backslash\mathcal{P}'$ and reviewer r_i such that r_i reviews both p_x and p_y, r_i must rank $p_x \succ p_y$ since otherwise the assumption of $\mathcal{P}'$ is violated. Since f is PU, we know that $f(\boldsymbol{\pi})$ must respect $p_x \succ p_y$ as well. This argument holds for every $p_x \in \mathcal{P}'$ and $p_y \in \mathcal{P}\backslash\mathcal{P}'$ that have been reviewed by at least one reviewer, and hence $(\mathcal{G}, f)$ is also GU. ∎

B.2 Proof of Theorem 1

We assume that the condition on the partitioning of the conflict graph, as stated in the statement of this theorem, is met. We begin with a lemma which shows that for any aggregation algorithm $\mathfrak{B}$, Contract-and-Sort is group unanimous.

Lemma 1. For any assignment and aggregation algorithms $\mathfrak{A}$ and $\mathfrak{B}$, the aggregation procedure Contract-and-Sort is group unanimous.

Proof of Lemma 1. Let $f(\widetilde{\boldsymbol{\pi}}) :=$ Contract-and-Sort$(\widetilde{\boldsymbol{\pi}}, \mathfrak{B})$, where $\widetilde{\boldsymbol{\pi}}$ is a preference profile. Define $\pi = f(\widetilde{\boldsymbol{\pi}})$. Let k denote the number of SCCs in $G_{\widetilde{\boldsymbol{\pi}}}$. Construct a directed graph $\widetilde{G}_{\widetilde{\boldsymbol{\pi}}}$ such that each of its vertices represents a SCC in $G_{\widetilde{\boldsymbol{\pi}}}$, and there is an edge from one vertex to another in $\widetilde{G}_{\widetilde{\boldsymbol{\pi}}}$ iff there exists an edge going from one SCC to the other in the original graph $G_{\widetilde{\boldsymbol{\pi}}}$. Let $\tilde{v}_1, \ldots, \tilde{v}_k$ be a topological ordering of the vertices in $\widetilde{G}_{\widetilde{\boldsymbol{\pi}}}$. Since $\tilde{v}_1, \ldots, \tilde{v}_k$ is a topological ordering, then edges can only go from $\tilde{v}_{j_1}$ to $\tilde{v}_{j_2}$ where $j_1 < j_2$. Now consider any cut $(\mathcal{P}_X, \mathcal{P}_Y)$ in $G_{\widetilde{\boldsymbol{\pi}}}$ that satisfies the requirement of group unanimity, i.e., all edges in the cut direct from $\mathcal{P}_X$ to $\mathcal{P}_Y$. Then there is no pair of papers $p_x \in \mathcal{P}_X$ and $p_y \in \mathcal{P}_Y$ such that p_x and p_y are in the same connected component, otherwise there will be both paths from p_x to p_y and p_y to p_x, contradicting that $(\mathcal{P}_X, \mathcal{P}_Y)$ forms a cut where all the edges go in one direction. This shows that $\mathcal{P}_X$ and $\mathcal{P}_Y$ also form a partition of all the vertices $\tilde{v}_1, \ldots, \tilde{v}_k$. Now consider any edge (p_x, p_y) from $\mathcal{P}_X$ to $\mathcal{P}_Y$. Suppose p_x is in component $\tilde{v}_{j_x}$ and p_y in component $\tilde{v}_{j_y}$. We have $j_x \neq j_y$, since $\mathcal{P}_X$ and $\mathcal{P}_Y$ forms a partition of all SCCs; also it cannot happen that $j_x > j_y$, otherwise $\tilde{v}_1, \ldots, \tilde{v}_k$ is not a topological ordering returned by f. So it must be $j_x < j_y$, and the edge (p_x, p_y) is respected in the final ordering. ∎

Proof of Theorem 1. Under the assumptions on μ, λ and sizes of $\mathcal{R}_C, \mathcal{R}_{\bar{C}}, \mathcal{P}_C, \mathcal{P}_{\bar{C}}$, it is easy to verify that there is a paper allocation satisfies $|\mathcal{P}_i| \leqslant \mu, \forall\, i \in [m]$ and each paper gets at least λ reviews. The strategyproofness of Divide-and-Rank follows from the standard ideas in the past literature on partitioning-based methods [Alon *et al.*, 2011]: Algorithm 1 guarantees that reviewers in $\mathcal{R}_C$ do not review papers in $\mathcal{P}_{\bar{C}}$, and reviewers in $\mathcal{R}_{\bar{C}}$ do not review papers in $\mathcal{P}_C$. Hence the fact that Divide-and-Rank is strategyproof trivially follows from the assignment procedure where each reviewer does not review the papers that are in conflict with her, as specified by the conflict graph $\mathcal{C}$. Given that all the other reviews are fixed, the ranking of the papers in conflict with her will only be determined by the other group of reviewers and so fixed no matter how she changes her own ranking. On the other hand, from Lemma 1, since Contract-and-Sort is group unanimous, we know that π_C and $\pi_{\bar{C}}$ respect group unanimity w.r.t. $\boldsymbol{\pi}_C$ and $\boldsymbol{\pi}_{\bar{C}}$, respectively. Since $\boldsymbol{\pi} = (\boldsymbol{\pi}_C, \boldsymbol{\pi}_{\bar{C}})$, it follows that π_C and $\pi_{\bar{C}}$ also respect group unanimity w.r.t. $\boldsymbol{\pi}$. Finally, note that there is no reviewer who has reviewed both papers from $\mathcal{P}_C$ and $\mathcal{P}_{\bar{C}}$, the interlacing step preserves the group unanimity, which completes our proof. ∎

B.3 Proof of Theorem 2

Proof. The proof of Theorem 2 is a direct formulation of our intuition in Section Impossibility of Pairwise Unanimity. Without loss of generality let $(p_1, \ldots, p_l)$ be the cycle not reviewed by a single reviewer, for $l \geqslant 3$. Hence there exists a partial profile $\boldsymbol{\pi}$ such that for all the reviewers who have reviewed both p_j and p_{j+1}, $p_j \succ p_{j+1}, \forall j \in [l]$ (define $p_{l+1} = p_1$). On the other hand, since for each reviewer, at least one pair (p_j, p_{j+1}) is not reviewed by her, the constructed partial profile is valid. Now assume f is PU, then we must have $p_1 \succ \cdots \succ p_l$ and $p_l \succ p_1$, which contradicts the transitivity of the ranking. ∎

B.4 Proof of Corollary 1

Proof. We prove each of the conditions in order.

Proof of part (i): If there is a cycle of size $\mu + 1$, then no reviewer can review all the papers in it since it exceeds the size of review sets. So there is no such cycle.

Proof of part (ii): The statement trivially holds for $\mu = 2$. For $\mu \geqslant 3$, Suppose there are two reviewers r_{i_1} and r_{i_2} such that $2 \leqslant |\mathcal{P}_{i_1} \cap \mathcal{P}_{i_2}| \leqslant \mu - 1$. Since $\mathcal{P}_{i_1} \neq \mathcal{P}_{i_2}$, there exist papers p_{j_1} and p_{j_2} such that $p_{j_1} \in \mathcal{P}_{i_1} \backslash \mathcal{P}_{i_2}$ and $p_{j_2} \in \mathcal{P}_{i_2} \backslash \mathcal{P}_{i_1}$. Also $|\mathcal{P}_{i_1} \cap \mathcal{P}_{i_2}| \geqslant 2$, and let $p_{j_3}, p_{j_4} \in \mathcal{P}_{i_1} \cap \mathcal{P}_{i_2}$. By definition it is easy to verify that $(p_{j_1}, p_{j_3}, p_{j_2}, p_{j_4})$ forms a cycle that satisfies the condition in Theorem 2, and hence $(\mathcal{G}, f)$ is not pairwise unanimous.

Proof of part (iii): Define a "paper-relation graph" $\mathcal{G}_p$ as follows: Given a paper-review assignment $\{\mathcal{P}_i\}_{i=1}^m$, the paper-relation graph $\mathcal{G}_p$ is an undirected graph, whose nodes are the *distinct* sets in $\{\mathcal{P}_i\}_{i=1}^m$; we connect two review sets iff they have one paper in common. Note that by (ii), each pair of distinct sets has at most one paper in common.

We first show that $(\mathcal{G}, f)$ is pairwise unanimous, then $\mathcal{G}_p$ must necessarily be a forest of cliques; in other words, all cycles in $\mathcal{G}_p$ are essentially cliques. To see this, not losing generality consider a cycle in $\mathcal{G}_p$ as $(\mathcal{P}_1, ..., \mathcal{P}_l)$. Also, suppose $\mathcal{P}_1 \cap \mathcal{P}_2 = \{p_1\}, \mathcal{P}_2 \cap \mathcal{P}_3 = \{p_2\}, ..., \mathcal{P}_l \cap \mathcal{P}_1 = \{p_l\}$. We consider two cases:

- $p_1 = p_2 = \cdots = p_l$. In this case $(\mathcal{P}_1, ..., \mathcal{P}_l)$ is a clique.
- At least two papers are different; suppose $p_1 \neq p_2$. We have $p_2 \notin \mathcal{P}_1$, since otherwise $p_1, p_2 \in \mathcal{P}_1 \cap \mathcal{P}_2$, which violates our assumption that $|\mathcal{P}_i \cap \mathcal{P}_j| \leqslant 1$ for every $i \neq j$. So there is a cycle that contains p_1, p_2, and reviewers of $\mathcal{P}_1$ have not read p_2. This violates the condition in Theorem 2. So this cannot happen, and all cycles are cliques.

We now use this result to complete our proof. Consider the union of all sets of papers that form vertices of $\mathcal{G}_p$. We know that this union contains exactly n papers since each paper is reviewed at least once. Now let k_p denote the number of distinct review sets (that is, number of vertices of $\mathcal{G}_p$), and let $\mathcal{P}_{i_i}, ..., \mathcal{P}_{i_{k_p}}$ denote the vertices of $\mathcal{G}_p$. Now from our previous analysis we know that for a clique C in $\mathcal{G}_p$, every review set has one common paper, and all other papers in the review set are different from other sets. So each clique contains $n_C(\mu - 1) + 1$ papers, where n_C is the size of the clique. Now let $m^{(C)}$ be the number of edges between cliques, and $n^{(C)}$ be the number of cliques. Notice that $\mathcal{G}_p$ is a forest of cliques, we have

$$\begin{aligned} n &= \sum_{C \in \mathcal{G}_p} |C| - \sum_{C, C' \in \mathcal{G}_p} |C \cap C'| \\ &= \sum_{C \in \mathcal{G}_p} (n_C(\mu - 1) + 1) + n^{(C)} - m^{(C)} \\ &= k_p \mu - \sum_{C \in \mathcal{G}_p} n_C + n^{(C)} - m^{(C)} \geqslant k_p(\mu - 1) + 1. \end{aligned}$$

The first equality is from the inclusion-exclusion principle, and that the union of three or more cliques is empty, since otherwise $\mathcal{G}_p$ cannot be a forest of cliques. Putting k_p on the left side gives the result $k_p \leqslant \frac{n-1}{\mu-1}$. ■

B.5 Proof of Proposition 2

Proof. Fix some ranking of papers within each individual set P_1, P_2, P_3 and P_4 (e.g., according to the natural order of their indices). In the remainder of the proof, any ranking of all papers always considers these fixed rankings within these individual sets. With this in place, in what follows, we refer to any ranking in terms of the rankings of the four sets of papers.

Suppose there is one such f that satisfies group unanimity and weak strategyproofness for $\mathcal{G}$, and consider the following 4 profiles:

(1) $r_1 : P_1 \succ P_2, r_2 : P_2 \succ P_3, r_3 : P_3 \succ P_4$
(2) $r_1 : P_2 \succ P_1, r_2 : P_3 \succ P_2, r_3 : P_4 \succ P_3$
(3) $r_1 : P_2 \succ P_1, r_2 : P_2 \succ P_3, r_3 : P_3 \succ P_4$
(4) $r_1 : P_2 \succ P_1, r_2 : P_3 \succ P_2, r_3 : P_3 \succ P_4$

By the property of GU, profile (1) leads to output $P_1 \succ P_2 \succ P_3 \succ P_4$, whereas (2) leads to output $P_4 \succ P_3 \succ P_2 \succ P_1$. Now compare (1) and (3): The output of (3) must have P_2 at the top and satisfy $P_3 \succ P_4$, by the property of GU. So the output of profile (3) must be one of i) $P_2 \succ P_1 \succ P_3 \succ P_4$, ii) $P_2 \succ P_3 \succ P_1 \succ P_4$, or iii) $P_2 \succ P_3 \succ P_4 \succ P_1$. Now note that only reviewer r_1 changes ranking across profiles (1) and (3), and hence by WSP the position of at least one paper in the output of profile (3) must be the same as in that of profile (1). This makes iii) infeasible, so the output of (3) must be either i) or ii). Similarly, the output of (4) is either $P_3 \succ P_4 \succ P_2 \succ P_1$ or $P_3 \succ P_2 \succ P_4 \succ P_1$. Now comparing (3) and (4): only r_2 changes ranking, but none of the four papers can be at the same position no matter how we choose the outputs of (3) and (4). This yields a contradiction. ■