How Connected Are Our Conference Review Boards?

Sourav S Bhowmick
Nanyang Technological University, Singapore
assourav@ntu.edu.sg

ABSTRACT

Dense co-authorship network formed by the review board members of a conference may adversely impact the quality and integrity of the review process. In this report, we shed light on the topological characteristics of such networks for three major data management conference venues. Our results show all these venues give rise to dense networks with a large giant component. We advocate to rethink the traditional way review boards are formed to mitigate the emergence of dense networks.

1. INTRODUCTION

Conference review process enables us to rapidly vet our research results through peer review and then quickly share them. Our professional societies (*e.g.,* ACM) strive to ensure that our conference review boards are of high quality and continue to serve the review process effectively. Specifically, program committee (PC) chairs of our review boards have endeavoured to improve the quality by selecting experts to ensure high coverage of all topics of interest, diversifying members along various dimensions (*e.g.,* gender, location, experience), monitoring reviews, among others. Despite these efforts, there has been anecdotal evidence on the existence of collusion rings and violation of a venue's conflicts-of-interest (COI) policy [1,11,12] that undermine the quality, fairness, and integrity of the review process. Hence, it is paramount to look beyond these traditional strategies (topic coverage, diversity) to enhance the quality of review processes.

Co-authorship (*i.e.,* collaboration) relationship is one of the key pillars of COI policies for all major venues. Intuitively, given two sets of review board members with similar topic coverage and diversity, R_1 and R_2, it is superior to choose R_1 over R_2 if the co-authorship network formed by R_1 is significantly less dense than that of R_2. This is because a dense network may adversely impact the review process in at least three ways. First, it increases the likelihood of (even inadvertent) undeclared COI and COI violations. An author (can be a review board member[1]) may have higher chance of prior co-authorship with *some* review board members of a dense subnetwork that may potentially give rise to COIs that are either unreported by the author or existing duration-based COI policies fail to capture them. To elaborate on the latter case, consider an author a and two reviewers m_1 and m_2 with strong co-authorship tie (*i.e.,* there is an edge (m_1, m_2) in the network). All are located in the same region. Suppose a has strong co-authorship tie with m_1. The likelihood of m_2 to have ties with a (*i.e.,* a wedge co-authorship pattern) may be higher in this scenario compared to the case where m_1 and m_2 are isolated nodes in the network. Furthermore, in some cases they may be part of a collusion ring [11]. A recent anecdotal evidence of such possible collusion involving co-authors is mentioned in [1].

Second, a denser network increases the likelihood of a set of reviewer board members reviewing a submission to be connected. For instance, in one of the submission cycle of a major data management venue, around 40% of the submissions have at least one co-authorship edge between assigned reviewers and 36% of these submissions have strong ties (10 or more papers). For some cases, this may not have any adverse impact on the review process. But for other cases, a group of connected reviewers may either collude together to reach a favourable decision for a submission authored by close ties (*e.g.,* [1]) or a junior member may be unduly influenced by an influential reviewer who has been their co-author[2].

Third, a dense network makes it challenging for a PC chair to assign unbiased reviewers for submissions where authors either have close ties with members of a subnetwork or they themselves are also review board members having high degree cen-

[1]The number of review board members who are also authors of submissions is often significant. For example, at least 58% of them are found to be authors in a recent major venue.

[2]This behaviour can be explained by the *social impact theory* in social psychology [8].

trality or clustering coefficient in the network.

Despite the potential impact of co-authorship network topology on the review process, to date there has not been any systematic study that shed light on the characteristics of these networks in data management venues. In this report, we take a concrete step to this end. Our study revealed the existence of dense networks with small-world characteristics in all major data management venues. We report various topological features of these networks and their implications on the review process. We conclude by advocating the need to depart from the traditional approach of review board formation to a data-driven, *system-based* approach to mitigate the emergence of dense co-authorship networks.

2. DATASET

We consider the review boards of recent editions of three major data management venues, SIGMOD 2022, VLDB 2022, and ICDE 2022, for our study. The lists of members (meta-reviewers and reviewers) were received from the PC chairs of respective venues. Table 1 reports the statistics. In practice, a review board can be dynamic in nature with the addition and removal of board members during the review process for various reasons. Hence, we consider the aggregated review boards of VLDB and ICDE after the final submission cycles (*i.e.,* March 2022 and November 2021, respectively). For SIGMOD we use the review board of the first submission cycle (July 2021) since we did not receive any subsequent updates.

Table 1: Dataset.

Venue	# of meta-reviewers	# of reviewers
SIGMOD 2022	26	177
VLDB 2022	33	200
ICDE 2022	22	194

We *manually* retrieved the DBLP addresses of all review board members using Google search. Each review board member for a given venue is uniquely identified by their email address or *DBLP name* which is unique in DBLP[3].

3. NETWORK GENERATION

The *co-authorship network* $C = (V, E)$ of a review board is an undirected, labeled, weighted graph where V is a set of review board members and E is a set of co-authorship edges between them. Given a pair of review board members $u, v \in V$, $(u, v) \in E$ iff u and v have co-authored one or more articles. A node $u \in V$ is labeled with an unique identifier of the review board member and $(u, v) \in E$ is labeled with a *weight* w representing the number of co-authored articles by u and v. We automatically generate the co-authorship network of a review board from DBLP by leveraging the co-authorship network generation component of CLOSET [5], a state-of-the-art COI detection and management system. Specifically, for each review board member it retrieves the corresponding XML version of their DBLP page and extracts all co-authors who are members of the review board and computes the frequencies of co-authorship. Then the network is constructed from it. Each node is labeled with the corresponding DBLP name of the member.

Table 2: Network properties.

Prop.	Venue	C_M	C_R	C	ER
$\|V\|$	SIGMOD	26	177	203	177
	VLDB	**33**	**200**	**231**	200
	ICDE	22	194	216	194
$\|E\|$	SIGMOD	14	286	483	279.77
	VLDB	**41**	378	**702**	377.65
	ICDE	20	**411**	652	412.55
$\langle k \rangle$	SIGMOD	1.0769	3.2316	4.7586	3.1612
	VLDB	**2.4848**	3.78	**6.0779**	3.7765
	ICDE	1.8182	**4.2371**	6.0370	4.2531
ρ	SIGMOD	0.0431	0.0184	0.0236	0.018
	VLDB	0.0777	0.019	0.0264	0.019
	ICDE	**0.0866**	**0.022**	**0.0281**	0.022
$\langle c \rangle$	SIGMOD	0.0577	0.1407	0.2177	**0.0142**
	VLDB	**0.1162**	0.2142	**0.3216**	**0.0166**
	ICDE	0.1015	**0.2609**	0.2999	**0.0205**
M	SIGMOD	**13**	31	20	9.18
	VLDB	5	27	14	5.37
	ICDE	8	**35**	**27**	3.88
L	SIGMOD	11	142	180	168.1
	VLDB	**29**	**170**	**218**	195.44
	ICDE	15	156	186	191.06
Φ_ℓ	SIGMOD	0.058	0.164	0.264	0.0172
	VLDB	**0.131**	0.253	**0.402**	0.017
	ICDE	0.033	**0.323**	0.395	0.0222

For each review board, we generate three types of co-authorship networks, *meta-reviewer network*, *reviewer network*, and *review board network*. In a meta-review network $C_M = (V_M, E_M)$, V_M represents the set of meta-reviewers. On the other hand, in a reviewer network $C_R = (V_R, E_R)$, V_R represents the set of reviewers. The review board network $C = (V, E)$ is the aggregated network of meta-reviewers and reviewers, *i.e.,* $V = V_M \cup V_R$.

4. NETWORK PROPERTIES

In this section, we analyse various properties of the co-authorship networks of the three venues.

Global properties. We first report the global topological properties of the networks. Specifically, we compute the network size, average degree (denoted by $\langle k \rangle$), density (denoted by ρ), average clustering coefficient (denoted by $\langle c \rangle$), number of connected components (denoted by M), the size of the largest connected component (*i.e.,* giant component) (denoted by L), and the local network efficiency [10][4] (denoted by Φ_ℓ) of the three types of co-

[3] In DBLP, homonyms are distinguished from one another by a unique numerical suffix to their name.

[4] Efficiency of a network measures how efficiently it exchanges

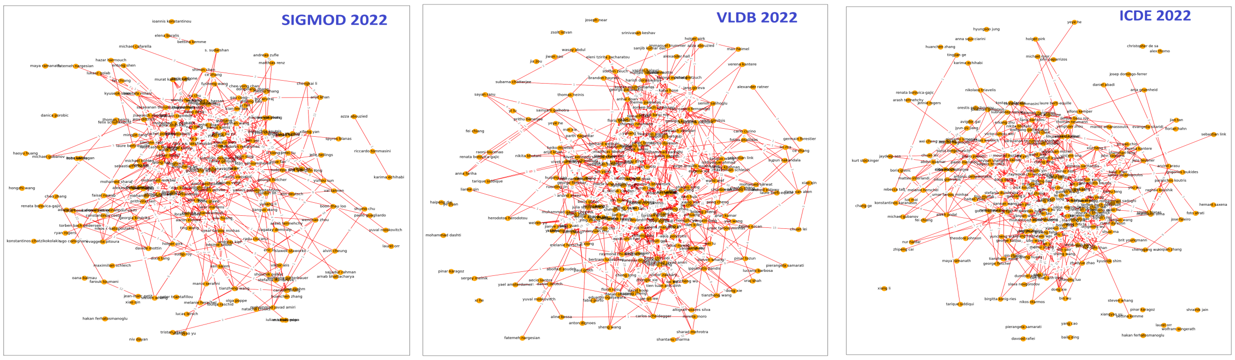

Figure 1: Review board networks.

authorship networks for each venue. Table 2 reports the results. We can make the following observations. First, although the size of the review boards of all venues is similar, the number of edges in the three types of co-authorship networks of VLDB and ICDE are significantly higher than that of SIGMOD. That is, the review boards of the former have more collaborative links compared to the latter. Second, the densities are high with ICDE networks being the highest among all. Third, the average clustering coefficients of all venues are significantly higher than the corresponding $\langle k \rangle/N$ values, demonstrating small-world properties of these networks [9]. For instance, $\langle k \rangle/N = 0.019$ for C_R of VLDB2022 which is significantly lower than $\langle c \rangle = 0.2142$. Specifically, the high values of $\langle c \rangle$ indicate well-connectedness of the neighborhood of a reviewer in these networks. Observe that ICDE and VLDB networks have higher $\langle c \rangle$ than that of SIGMOD. Fourth, the number of connected components in all networks is low hovering between 14-35 in C_R and C. Interestingly, the size of the giant components in C_R and C is large for all venues. For example, around 94% of the nodes in C of VLDB are part of the giant component! Lastly, C_R of ICDE2022 has the highest local efficiency. Figure 1 depicts the review board networks of these venues. Note that the goal here is not visual clarity as graphs with more than 100 nodes look like a hair-ball [13]. Instead, the intention here is to visually appreciate the denseness of these networks and the existence of a large giant component.

Comparison with random network. The above results demonstrate high average clustering coefficient (resp. local network efficiency) in C_R and C highlighting strong co-authorship ties and information exchange between the neighborhoods of review board members. *Can these properties emerge by chance?* We utilize the reviewer network C_R to answer this question. We construct a random network of similar size and average degree of C_R and compute these measures of the network. Specifically, we construct an Erdos-Renyi (ER) network using the $G(N, p)$ model [6]. We set $N = |V_R|$ and $p = \langle k_R \rangle / N - 1$ [6] where $\langle k_R \rangle$ is the average degree of C_R. This enables us to generate an ER network with size and density similar to C_R. We generate 100 instances of the ER network and compute the average values of the seven topological properties. The last column in Table 2 reports the results. Observe that the avg. clustering coefficients and local efficiency of the ER networks are around an order of magnitude smaller than the corresponding values in C_R. Hence, randomness cannot explain them and it may represent some "signature of order, requiring a deeper explanation" [4]. The ER network also demonstrates the co-existence of a giant component and isolates. The emergence of giant component is expected as $1 < \langle k \rangle < ln|V|$ in these networks [4,7]. However, its size is larger than the ones in C_R.

Degree distribution. Figure 2 depicts the degree histograms of reviewer networks (C_R) of the three venues. Observe that all venues have several high-degree reviewers (*i.e.,* degree more than 7).

PC overlap. Next, we look at the amount of overlap between the review boards. Since VLDB has the most onerous multi-cycle review process, we compute the overlap between the nodes in the *reviewer networks* of (VLDB 2022, SIGMOD 2022) and (VLDB 2022, ICDE 2022). We focus on the reviewer network since reviewing fatigue is encountered most by the reviewers. The number of common reviewers are 40 and 34, respectively. To get an understanding of the trend of the overlap, we also compute the overlap between the 2021 and 2023 editions of VLDB and SIGMOD. There are 38 and 60 common reviewers in 2021 and 2023, respectively, showing an upward trend in the size of common reviewers.

Weight thresholding. A co-authorship network may contain strong and weak ties where the

information [10]. Local efficiency is the average efficiency of the local subgraphs. Small-world networks have high global and local efficiency.

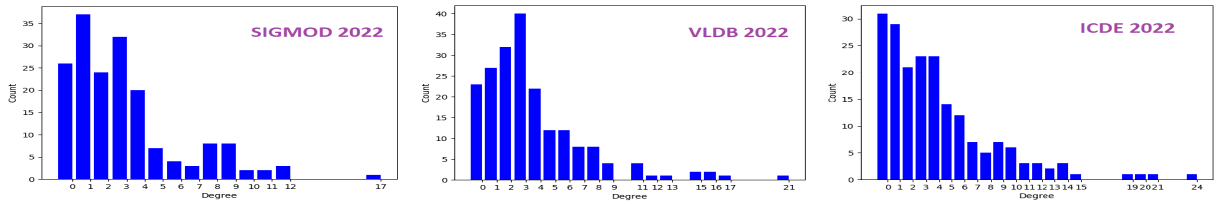

Figure 2: Degree histograms of reviewer networks.

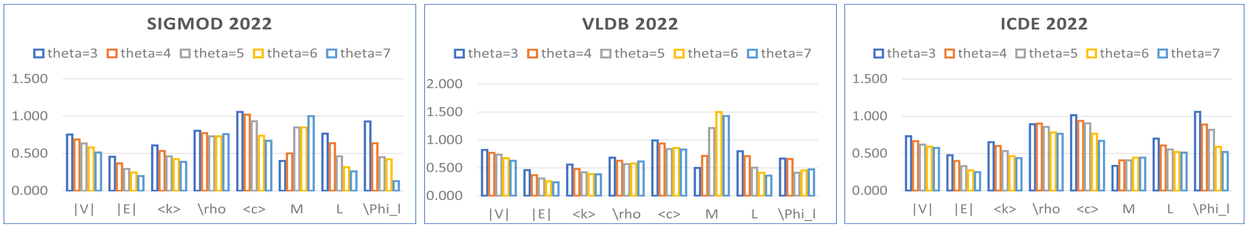

Figure 3: Network properties vs weight thresholding (Reviewer networks).

strength of a tie is measured using frequency of co-authorship (*i.e.,* edge weight). In recent times, the frequency of co-authorship has been incorporated in the COI policies of several venues (*e.g.,* ICDE 2020, ICDE 2021, VLDB 2023). Hence, in this set of experiments we study the topological properties of the subnetworks with strong ties in a co-authorship network using *weight thresholding*. That is, we sparsify the network by using an *edge weight threshold* θ and observe how the topological properties of subnetworks with strong ties evolve. Given θ and a co-authorship network C, we remove all edges with $w < \theta$ along with nodes that have no edge with $w \geq \theta$ resulting in a sparsified network $C_\theta = (V_\theta, E_\theta)$ where $V_\theta \subseteq V$ and $E_\theta \subseteq E$. Then, for each property P we compute P_θ/P where P_θ is the value of the property P in C_θ and P is the corresponding value in C. We use the reviewer network C_R of each venue as the original network and vary θ from 3 to 7.

Figure 3 plots the results. Observe that the densities of C_θ, remain relatively robust. Importantly, for VLDB 2022 and ICDE 2022 a significant part of the networks (around 60% of the nodes) maintain strong ties when $\theta = 7$. Furthermore, $\langle c \rangle$ and Φ_ℓ of VLDB 2022 remain very stable with increasing θ. Given that both these venues employ only duration-based COI (*i.e.,* 2 years) [2,3], they may not be able capture many of these strong ties in their COI policies. In summary, subnetworks with strong ties and stable topological features have significant presence in the co-authorship networks of some venues that may make their COI policies inadequate for the intended purpose.

5. DISCUSSIONS

We revealed that the review boards of 2022 edition of three major data management venues form dense co-authorship network that cannot be modeled by random networks. Note that the topological properties of these networks are qualitatively similar to several other editions as well. For instance, ρ, $\langle c \rangle$, and Φ_ℓ of C_R of 2021 edition of these venues vary between 0.02-0.033, 0.23-0.24, and 0.28-0.29, respectively. Similarly, for SIGMOD 2023 and VLDB 2023 these values vary between 0.021-0.023, 0.22-0.23, and 0.280-0.283, respectively. While such features are highly desirable for social networks, as remarked earlier, in a peer review setting they may potentially create challenges for PC chairs to assign unbiased reviewers to submissions.

We believe that the manifestation of dense co-authorship networks of our review boards is mainly due to the traditional "manual" approach of review board formation where a candidate reviewer is invited independent of other candidates primarily based on the recommendations from meta-reviewers and PC chairs. We advocate that it is important to advance towards a *system-based* approach where the goal is to select a set of review board members that is diverse, covers the topics of a venue, but forms a *sparse* co-authorship network with low average clustering coefficient and a smaller giant component and network efficiency. Naturally, this is infeasible to achieve purely manually (by PC chairs) or purely automatically in practice. Hence, it is paramount to build *data-driven*, *PC chair-in-the-loop* tools that can facilitate it. A possible approach can be to randomly choose reviewers from a large pool of candidates, while satisfying the topic constraints. Exploring this and related options are avenues of future work.

Acknowledgement: We thank Divesh Srivastava for his insightful comments on the article and the paradigm of data-driven PC formation.

6. REFERENCES

[1] https://twitter.com/benrozemberczki/status/1549677286642917377, Last accessed: 1st August, 2022.

[2] IEEE ICDE 2022 Call for Research Papers. https://icde2022.ieeecomputer.my/research-track/#conflict-of-interest, Last accessed: 1st August, 2022.

[3] PVLDB vol 15 - Submission guidelines. http://vldb.org/pvldb/vol15-submission/, Last accessed: 1st August, 2022.

[4] A. L. Barabasi. Network Science. *Cambridge University Press*, 2016.

[5] S. S. Bhowmick. CLOSET: Conflict of Interest Detection and Management System. https://personal.ntu.edu.sg/assourav/research/DARE/closet.html, 2019, Last accessed: 1st August, 2022.

[6] P. Erdos, A. Rényi. On random graphs. *I. Publicationes Mathematicae (Debrecen)*, 6:290-297, 1959.

[7] P. Erdos, A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung.* Acad. Sci., 5:17-61, 1960.

[8] B. Latané. The psychology of social impact. *American Psychologist*, 36 (4): 343–356, 1981.

[9] Y. Li, Y. Shang, Y. Yang. Clustering coefficients of large networks. *Inf. Sci.* 382-383: 350-358, 2017.

[10] V Latora, M Marchiori. Efficient behavior of small-world networks. *Physical review letters*, 2001.

[11] M. L. Littman. Collusion rings threaten the integrity of computer science research. *Commun. ACM*, 64(6):43-44, 2021.

[12] M. Winslett, R. Snodgrass. We need to automate the declaration of conflicts of interest.*Commun. ACM*, 63(10), 2020.

[13] V. Yoghourdjian, D. Archambault, S. Diehl, T. Dwyer, K. Klein, H. C. Purchase and H.-Y. Wu. 2018. Exploring the limits of complexity: a survey of empirical studies on graph visualization. *Visual Informatics* 2, 4 (2018), 264-282.