Fairness of Scientific Collaboration Networks in Academic Conferences

Ritwik Dutta and Radu Marculescu

***Abstract*—This paper outlines our research in academic social networks. We analyze the complex collaboration structure among the participants of academic conferences, and outline a methodology to characterize a conference in terms of its important personality metrics. In particular, we focus on networks capturing the collaboration between the program committee members and the regular authors of a conference, including co-authorship and institutional linkages, and derive a quotient to quantify the *fairness* of scientific collaborations for the conference. Using a simple model, we show how the *stability* of a conference can be evaluated as an inverse measure of the fairness quotient. We have found that the fairness factor of a conference is a dynamic property that can vary from one year to another based upon the co-authorship network of the conference. Our findings offer fundamental insights into the evolution of special communities or interest groups that tend to form in an academic network with the sole purpose of promoting interests with regard to professional growth of individuals in these communities at the cost of denying privileges to individuals outside the communities. We have presented our results in the light of two popular conferences that we have analyzed over a period of five years, but our methodology is applicable to any conference and over any number of years, and can be extended to a study of other complex social networks as well.**

***Keywords*—ForceAtlas, Gephi, network communities, network graphs, OpenOrd, scientific collaboration, social networks.**

I. Introduction

ACADEMIC conferences are collaborative and multi-site initiatives that involve students, researchers, educators, and organizations. Such a conference is usually characterized by scientific collaborative networks of people, institutions, and organizations that are largely independent, geographically distributed, and diverse in terms of their operating environment, society, and culture. Given the importance of academic conferences in influencing professional careers in terms of hiring and promotions, it is imperative that fair and stringent policies should exist with respect to authorship credit, paper review, paper acceptance, plagiarism, and resubmission of previous work for publications in an academic conference.

For a conference, particularly a selective one, the review process needs to be unbiased, balanced, and free of conflicts of interest. In this regard, the Technical Program Committee (TPC) plays a crucial role in terms of its goals, composition, and responsibility as a steering committee to handle the TPC's influence on reviewers, authors and TPC-authored papers. A "common interest clique" formed between one or more TPC members and regular authors can jeopardize the ideals on which a conference is founded; moreover, the existence of such a special-interest community or group can unfairly promote professional growth of individuals in the community at the cost of denying the same privileges to those outside the community.

The work presented in this paper concerns coming up with a systematic analysis methodology to characterize academic conferences in terms of some important personality metrics. The metric that we have focused on is *fairness* and we have outlined a method to derive the *fairness quotient* of a conference. Using a simple model, we also show how the *fairness stability* of a conference can be evaluated as an inverse measure of the fairness quotient. We have found that the fairness factor of a conference is a dynamic property that can vary from one year to another based upon the co-authorship network of the conference. Our method is generic and can handle a complex set of variables, hence can be extended to look at other criteria as well as study other complex social networks.

We refer to the network of collaboration among the participants of a conference as the Conference Collaboration Network (CCN). A three-step approach is followed to identify the fairness factor of such a CCN. We first target a select set of conferences and obtain the relevant collaboration data. Next, we extract patterns from the collaboration data. Finally, we draw inferences about the nature of the conference from the extracted patterns. Whereas much of previous work has addressed academic conduct, integrity, and authorship order, our paper is one of the first to analyze the impact of the program committee on authorship. In light of the variability, ambiguity, and uncertainty that exist today, we believe our results can be used to guide and optimize policies regarding the rules of engagement in different academic events and networks.

The paper is organized as follows. Section II describes past work in the area of social group network analysis. In Section III we clarify certain terminologies and explain our research objectives. Section IV details on the effective representation of academic networks in a way that helps easy visualization of data. Section V presents our derivation of the fairness factor and stability of a conference. Section VI incorporates the results, and Section VII draws the conclusions.

Ritwik Dutta is with Archbishop Mitty High School, San Jose, CA 95129 USA (e-mail: ritzymail@gmail.com).

Radu Marculescu is with the Electrical & Computer Engineering Department, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: radum@cmu.edu).

II. PAST WORK

The CCNs that we have analyzed are interdependent networks. Past work has addressed interdependent networks as well. For example, Danziger's paper [1] on interdependent networks, both in terms of their static and dynamic properties, has shown that multiple networks depending on each other, through connections that defined each network as a "layer" with inter-layer connections providing the order and structure of this interdependence, could be used for interpretation of complex and interconnected data. In the context of social and scientific networks, evolution and collaboration within the network is important for our research. One such analysis [2] of social networks, that included the Cornell University Library archive and customer-to-customer phone calls for a 4-million-user mobile network, has been used to demonstrate that these networks are constantly evolving in terms of emergence and disappearance of communities as well as their dynamics. In the area of collaborative networks, a different analysis [3] of scientific collaborations for MEDLINE, the Los Alamos e-Print Archive, SPIRES, and NCSTRL has shown the differences in the average number of papers-per-author and authors-per-paper in four scientific areas – astrophysics, condensed matter, high-energy, and computer science; it was found that these collaboration networks were highly clustered and specific areas (such as high-energy-physics) had noticeably different amounts of authors-per-paper. Similarly, experimentation [4] with ResearchGate data, using network centrality and PageRank algorithm, used academic connections keyed by the "interest" field to help rank authors in a collaborative network in terms of *influence* on the network and *value*, while categorizing them by discipline. On similar lines, Barbasi *et al* [5] did a study of journal collaborations in mathematics and neuroscience, and showed that substantial differences exist in the rate of collaboration and the shift over time, of clustering and author separation as well as the increased fracturing in links from incoming authors, compared to those already in the network.

Our research differs from those mentioned in the previous paragraph primarily in the focus of our analysis. Previous research has focused mainly on analyzing the properties of the networks (*e.g.* centrality, separation, degree, community size, *etc.*), determining how they changed over time, and drawing results and conclusions from those properties. We have, however, used various network properties for the computation of the values of *fairness* and *stability* of an academic social network. We have also attempted to contextualize and interpret the fairness and stability values rather than just tabulating the results and showing their variance.

III. TERMINOLOGY AND OBJECTIVES

In our work, a CCN is represented as a regular network graph comprising *nodes* and *edges*. Each node represents a participant in the conference, either as an author, or a Technical Program Committee (TPC) member, or both. Each edge connects two nodes, either as co-authors or as participants from the same institution. The *degree* of a node is the number of connection edges it has. A *community* in the graph is defined as a *group of nodes* with dense connections between them. A CCN has an inherent, three-level *hierarchy*: authors, TPC members, and those who are both authors and TPC members.

When we started our analysis of conferences, we first investigated if any meaningful information about *co-authorship communities* could be obtained via network analysis of co-authorship edges in the graphical representation of the conference data. We felt that the results were inconclusive because while raw collaboration data did show the structure of the conference and its communities, it did not explain why the structure was the way it was. A similar problem surfaced in trying to measure the *success* of a participant in a conference using general and specific properties of the collaboration network and the network nodes. The main issue was that defining success solely in terms of collaborative values proved to be difficult and not entirely accurate. Defining success in terms of co-authorship is not a functional procedure; depending on the goals of the author and the personal ideals of success, the definition can vary. Since it was not be possible to determine the author's end goal from the data, we decided to not focus on determining the success of authors.

While concepts such as co-authorship communities and success are not possible to be determined based solely on collaboration data, *fairness* is something that is verifiable using only the network properties of the CCN. The method we use to determine fairness depends on analyzing the inter-level connections of multi-level collaboration. Initially, we resorted to manual analysis of both numerical and visual properties of the CCN in parallel. Ultimately, the visual properties proved to be more useful for the manual analysis, and based on the conclusions drawn, we set out to find various numerical properties of the CCN that could be combined to obtain a fairness quotient. Of course, the fairness quotient has to be size-agnostic, so that it is consistent across all conferences, large or small, and once such a quotient is obtained, conferences can be automatically rated for fairness without the need for manual analysis any more.

IV. NETWORK REPRESENTATION

A. Data Procurement

The obvious first step in procuring the CCN data is deciding upon a conference. Several different conferences (DAC, ICCAD, NOCS, CODES, EMSOFT) were selected for analysis and the relevant conference citation data - paper titles, authors, and affiliations - were gathered for a time period of five years (2010 to 2014). The original plan for collecting the data was to use the IEEE Xplore [6] API to automatically download all of the citations for each conference for the chosen time frame. However, due to API limitations, the data had to be gathered manually by using the command search feature that yielded data separately for each year of each conference; the data was exported using the "CSV export" feature. Next, the information regarding the TPC members was obtained from the website of each individual conference, and manually inserted into the CSV files. We decided to do a thorough analysis for the NOCS[1] and ICCAD[2] first before considering other conferences.

B. Data Processing

The conference data in CSV format was processed using a Python program we developed that generated the network graph. A node was created for each participant, and the node was tagged with the role (TPC or author or both), the institution of affiliation (university or company), and the name of the participant. The program also created edges between the nodes when they represented authors who were connected via co-authorship, institution, or even social groups (the last item representing the class where people from different institutions co-authored).

C. Data Visualization

During our analysis, different visualization methods were utilized to picture and understand the network properties better. The network graph itself was analyzed and visualized in Gephi [7], an open-source graph visualization platform. A script was written to automatically color the nodes and the edges in the network graph according to their tags, as shown in

Figure *1*. These colors helped understand the network more easily. However, simply coloring the network was not enough to clarify its meaning, as the issue of a good graphical layout still remained to be solved.

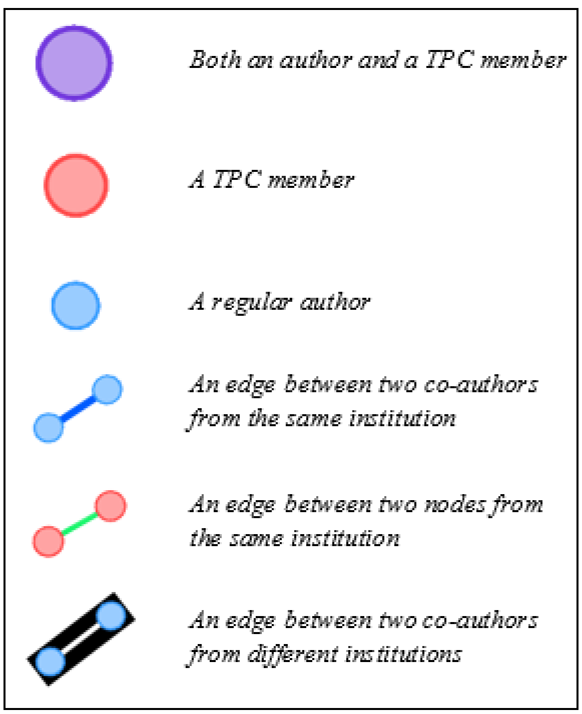

Figure 1. Color keys for graphs

D. Data Layout

At first, the default layout algorithms present in Gephi were used for force-directed layout of the network graph. However, this did not lead to good results because the layout did not show a clear visualization of the connections and the relationships in the graph. The nodes in the graph were far too compressed to allow any information to be gleaned from a visual inspection. The default algorithms used for layout in Gephi are the ForceAtlas [8] and the Yifan Hu [9] algorithms. Although neither satisfied our visualization requirements, the former turned out to be better than the latter. Figure 2 is a sample network layout using the ForceAtlas algorithm on the NOCS 2010 conference network.

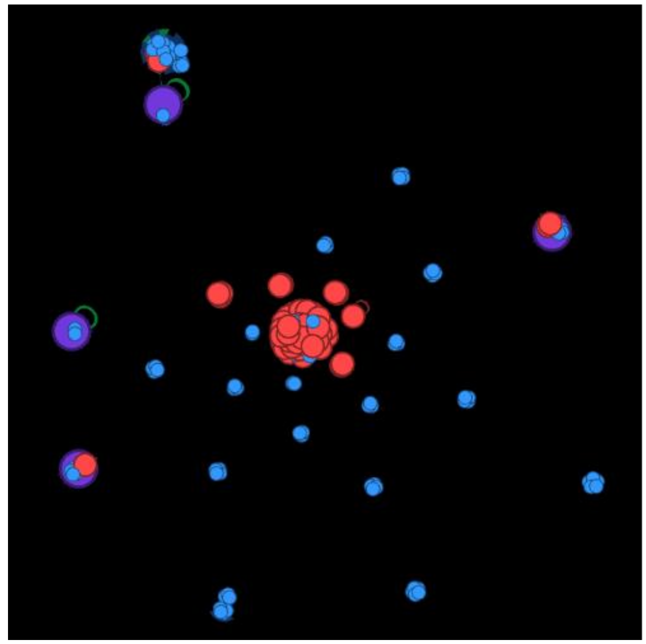

Figure 2. ForceAtlas algorithm driven layout

E. Layout Improvement Algorithms

Since the previously-mentioned algorithms were not very suitable, we searched for and found a potential alternative candidate on the internet – the OpenOrd [10] algorithm, available as a plug-in for Gephi. OpenOrd uses a five-stage simulated annealing process and an edge-cutting technique to generate layouts that scale well to large and complex networks. Figure 3 is a sample network layout using the OpenOrd algorithm on the NOCS 2010 conference network. In terms of the layout, the OpenOrd algorithm provided almost everything we wanted; it provided a good visualization of the connections in the network and also made it clear how "central" and "important" various nodes were, which was required to understand the hierarchical relationships between the nodes. However, OpenOrd was still not ideal for our purposes due to the visual clutter it created owing to the absence of any shape or form in the layout it portrayed.

Another layout algorithm that is available in and works with Gephi is Fruchterman-Reingold [11], one of the first algorithms for force directed drawing. Unfortunately, it did not satisfy our requirements when acting alone; the layouts created allowed visual understanding, but they missed the

[1] NOCS: International Symposium on Networks-on-Chip
[2] ICCAD: International Conference On Computer Aided Design

clear visual structure that lends itself easily to analysis. Fortunately, when Fruchterman-Reingold was applied to the graph in conjunction with the OpenOrd algorithm, it removed the previously-mentioned clutter and generated an easily understandable layout that still retained the important properties. Figure 4 is a sample network layout generated by applying the Fruchterman-Reingold algorithm to the NOCS 2010 OpenOrd network layout seen in Figure 3.

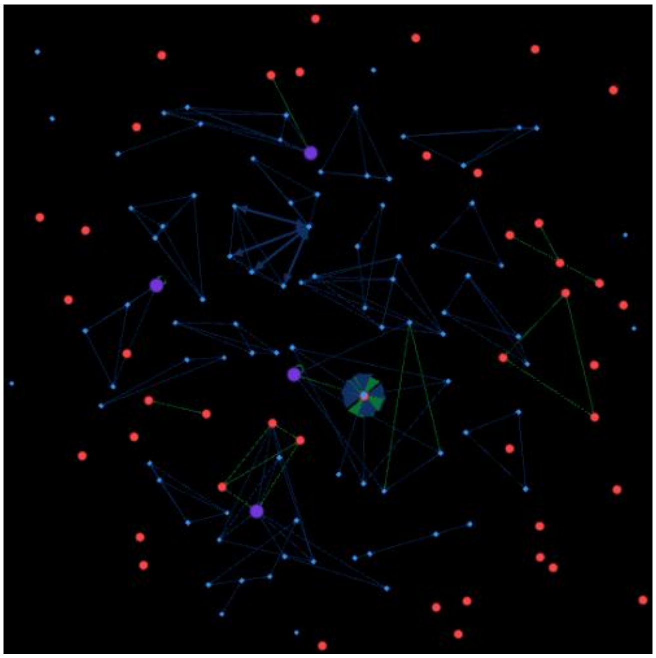

Figure 3. OpenOrd algorithm

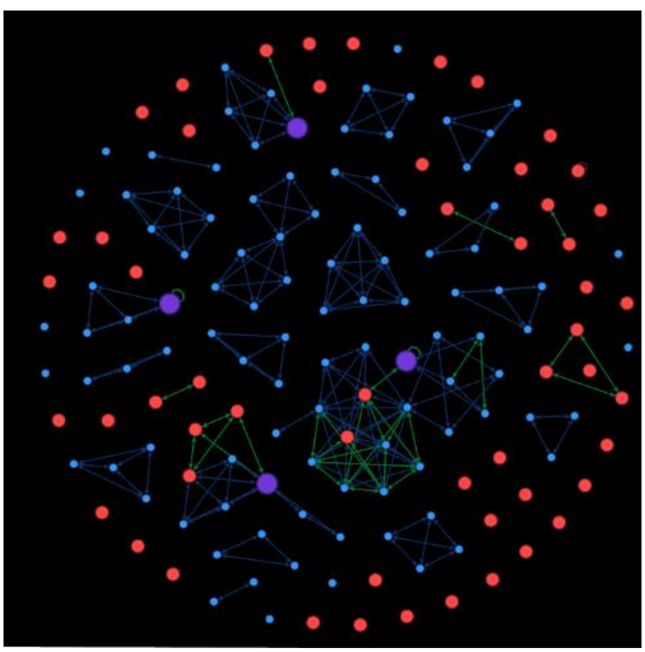

Figure 4. Fruchterman-Reingold applied after OpenOrd

As can be seen in Figure 4, the network layout obtained by using the two algorithms in sequence was almost perfect. But, while the graph shows connections between the nodes very clearly, it does miss the clarity of multi-level visualization. Visualizing connections between the different *levels* of nodes (*e.g.* authors, TPC members, TPC-authors[3]) is a key component because it makes the relationships and influences between nodes more visually apparent.

Getting a multi-level representation, however, was not difficult. We discovered a layout algorithm called Network Splitter 3D [12] on the Gephi plugin marketplace that allowed a network to be split into levels based on a *z-level* tag for each node, and then rotated about the *x*-axis to visualize the levels. When this splitter was applied to the network graph obtained after applying the Fruchterman-Reingold and OpenOrd algorithms, it produced an easy-to-understand, multilevel visualization that showcased all the relevant network properties and met our requirements. Figure 5 is a sample layout showing the result of applying the Network Splitter 3D algorithm on the NOCS 2010 Fruchterman-Reingold and OpenOrd network layout seen in Figure 4. With the CCN graphs represented in 3D, the next step was to make use of the visualization to perform different types of analysis on the networks.

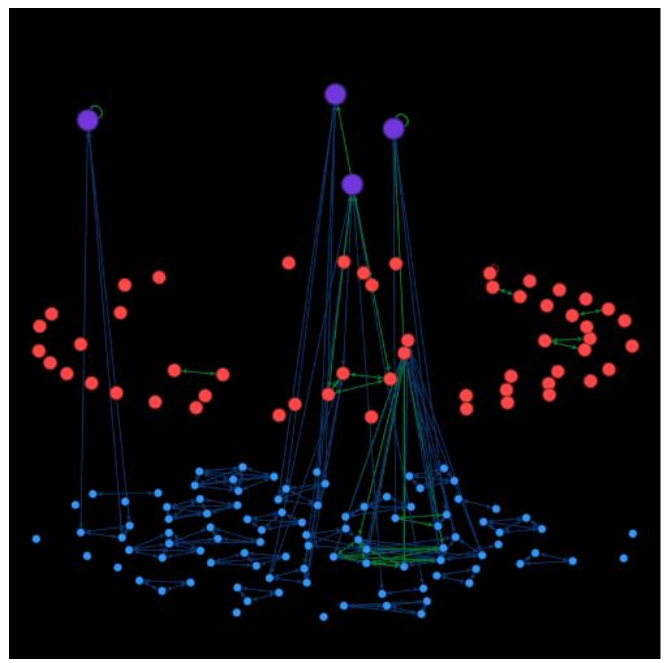

Figure 5. The Network Splitter 3D multilevel visualization applied to Fruchterman-Reingold and OpenOrd combination

V. NETWORK ANALYSIS

We wanted to understand how TPC members influenced authorship. In that regard, it was not only the regular authors but the TPC members as well who played a role in our concept of "fairness" in a network. In order to find a numerical explanation for this fairness metric in relation to the TPC-authors, all the relevant properties of the network and sub-networks were analyzed. These included the average degree of

[3] A TPC-author is a TPC member who also authored a conference paper.

authors, average degree of TPC members, number of communities, average path length, *etc*. These values alone did not function as indicators of fairness in themselves, but when applied to different segments of the network, they could be converted to a discrete fairness value.

After the network properties mentioned above were extracted for a conference, the focus shifted to trying to determine how they could be used to represent the fairness of the conference. To do so, a clear definition of fairness was needed. After consulting a variety of network topologies for different conferences, it was determined that an "unfair" conference would be one where relationships that were conflicts of interest either gave some authors advantages, or others disadvantages, in areas such as *paper acceptance likelihood* and *co-authorship desirability*.

There is an undeniable likelihood of papers authored by TPC members to be accepted due to the fact that by virtue of being a part of the steering committee, the TPC members are already regarded as more "elite" and influential than the other authors. Therefore, if a TPC member were to write a paper like any other author, that paper would have a higher probability of being accepted. This observation led us to the idea of comparing the authorship groups of TPC members to those of the other authors. If the groups were abnormally large, with TPC members having an average of ten co-authors versus an average of three for other papers for example, it would imply that they were most likely included due to the influence of their name and level, rather than actual contributions to the paper itself.

A. Fairness Quotient

To have a quantitative measure of fairness, we first extracted two community-related values from the conference collaboration network and combined them as follows:

- TPC-author community size mean (c_{tpc-a}) = the average size of the communities which TPC-authors were part of,
- Author community size mean (c_a) = the average size of the communities that the normal conference authors were part of, and
- TPC to author community ratio (r_c) = the ratio of the average size of TPC-author communities to the average size of normal author communities.

Next, we extracted from the network two different values related to the degree of the nodes, and combined them as follows:

- TPC-author degree mean (d_{tpc-a}) = the average degree of TPC-authors in the graph; this is different from community size because community sizes do not take into account connections outside of communities, which are frequently institutional connections that stretch across the network between communities,
- Author degree mean (d_a) = the average degree of all authors in the graph, and
- TPC to author degree ratio (r_d) = the ratio of the average degree of TPC-authors to the average degree of the general conference authors.

Finally, the total fairness value, that we call the *fairness quotient* (v_f) of a conference, was calculated from the TPC to author community ratio (r_c) and the TPC to author degree ratio (r_d) as follows:

- $\boldsymbol{r_c} = \varphi_1 \times \frac{c_{tpc-a}}{c_a}$
- $\boldsymbol{r_d} = \frac{d_{tpc-a}}{d_a}$
- $\boldsymbol{v_f} = \varphi_1 \times \frac{r_c + r_d}{2} = \varphi_1 \times \frac{\frac{\varphi_1 \times c_{tpc-a}}{c_a} + \frac{d_{tpc-a}}{d_a}}{2}$

Note: φ_1 is a re-ranging parameter for which we have used a value of 10.

VI. RESULTS

Table 1 shows the computed fairness quotient ($\boldsymbol{v_f}$) for each year of the two selected conferences (NOCS and ICCAD) over a period of five years. As can be seen, the "fairness" of conferences can vary over time.

Table 1. Fairness quotients from 2010 to 2014

Conference	Year	v_f
NOCS	2010	3.2
	2011	3.0
	2012	4.4
	2013	1.9
	2014	2.2
ICCAD	2010	4.1
	2011	5.7
	2012	6.1
	2013	2.3
	2014	6.5

In order to determine the correctness of our method of determining fairness of conferences, we compared our analytical results with the fairness determined from a manual visual analysis of the same conferences. For the NOCS conferences, we found $\boldsymbol{v_f}$ to be an excellent measure of the fairness for any particular year.

A. Average Fairness Interpretation

To understand the significance (and the problem) of averages, let us consider the following average values of the fairness quotient:

- $\overline{v}_{f-total}$ = average fairness quotient calculated for NOCS and ICCAD over five years,
- $\overline{v}_{f-nocs}$ = average fairness quotient calculated for NOCS over five years,
- $\overline{v}_{f-iccad}$ = average fairness quotient calculated for ICCAD over five years.

The values of the average fairness quotients are shown in Table *2*.

Table 2. Fairness averages for datasets

Average quotient	Value
$\overline{v}_{f-total}$	3.96
$\overline{v}_{f-nocs}$	2.94
$\overline{v}_{f-iccad}$	4.97

The table shows that the average fairness considering all of the data (for the two conferences) is 3.96. However, averaging a value across both conferences is not very meaningful; computing the quotient $\overline{v}_f$ for each conference and comparing different conferences based on the quotient make sense because that allows us to interpret the fairness differences of conferences in terms of the values on the same scale. Since the fairness quotient is computed using the ratios of metrics ($\boldsymbol{r_c}$ and $\boldsymbol{r_d}$) on a particular subset and not the entire network, the value of the quotient retains a consistent meaning independent of the size of the conference. Hence, looking at the differences in the average fairness values of NOCS and ICCAD is more meaningful than considering a total average. The average fairness of NOCS computed over the five years from 2010 to 2014 is 2.94. For ICCAD, the same computation returns a result of 4.97. ICCAD is thus about 69% more fair than NOCS.

B. Stability Interpretation

Let us represent the fairness change of a conference, between two consecutive years of interest, by Δ_f and the average fairness change over a period of time as $\overline{\Delta_f}$. Clearly, for the subsequent years y_1 and y_2, the fairness change is given by

$$\Delta_{f_{y_1 \to y_2}} = \left| v_{f_{y_1}} - v_{f_{y_2}} \right|$$

Therefore, for a range of years, y_1 to y_m, divided into n intervals, the *average fairness change* can be calculated as

$$\overline{\Delta_f} = \frac{\Delta_{f_{y_1 \to y_2}} + \Delta_{f_{y_2 \to y_3}} + \cdots + \Delta_{f_{y_{m-1} \to y_m}}}{n}$$

The average fairness value is next normalized using σ, the standard deviation of the $\Delta_{f_{y_j \to y_{j+1}}}$ values over the n year intervals. This yields a metric for the *standardized fairness change* and serves to standardize the meaning or interpretation of the average fairness change across all datasets. The standard deviation is calculated as

$$\sigma = \sqrt{\frac{\sum_{j=1}^{m-1} \left(\Delta_{f_{y_j \to y_{j+1}}} - \overline{\Delta_f} \right)^2}{n}}$$

After normalization, we obtain an expression for the standardized fairness change as

$$\overline{\Delta_{fs}} = \frac{\overline{\Delta_f}}{\sigma}$$

Stability being an inverse measure of variation or change, we define the *fairness stability* of a conference as

$$S_f \propto \frac{1}{\overline{\Delta_{fs}}},$$

which yields

$$S_f = \varphi_2 \times (\overline{\Delta_{fs}})^{-1} = {}^{\varphi_2}\!/\!_{\overline{\Delta_{fs}}}$$

Note that φ_2 is the constant of proportionality that we have also used as a re-ranging parameter with a value of 10. s_f is scaled by 10 in our examples to normalize its scale with respect to the other variables such as v_f.

C. Stability Analysis

Let us consider the stability of NOCS and ICCAD, as shown in

Table 3 and
Table *4* respectively, in terms of the following metrics:

- $\boldsymbol{s_{f-nocs}}$ = stability of NOCS calculated for the year intervals 2010-2011, 2011-2012, 2012-2013, and 2012-2013, and
- $\boldsymbol{s_{f-iccad}}$ = stability of ICCAD calculated for the year intervals 2010-2011, 2011-2012, 2012-2013, and 2012-2013.

From the NOCS fairness value analysis, we see that the fairness value changes between the years (relative to the annual fairness values) do not show large swings, but there are still noticeable fluctuations. These fluctuations are in the intervals of 2011 to 2012 and 2012 to 2013, and the data shows that a brief period of greater-than-average fairness emerged in the conference and subsequently vanished, leaving the conference more unfair than it was in 2010, before the temporary upturn had even started. When $\boldsymbol{s_{f-nocs}}$ is computed across the five consecutive single-year intervals from 2010 to 2014, we get a result of 8.54. Since we have already analyzed and characterized NOCS, a new conference with a stability value different from $\boldsymbol{s_{f-nocs}}$ will now allow us to characterize the new conference's stability relative to NOCS. When $\boldsymbol{s_{f-iccad}}$ is computed across the same five consecutive single-year intervals as NOCS, we get a stability value of 6.28. At a first glance, this would imply that ICCAD has more swings in fairness than NOCS making it less stable, but that is not the case; it is just that the swings that do happen, tend to be larger. For NOCS, the swings greater than the standard deviation are on an average 1.77 larger. For ICCAD, the swings larger than the standard deviation are larger by an average factor of 2.03. This also shows that ICCAD's significant fairness changes are about 14% larger than those of NOCS.

Table 3. NOCS stability computation

Parameter	Years	2010	2011	2012	2013	2014
v_f		3.2	3.0	4.4	1.9	2.2
$\Delta_{f_{y_j \to y_{j+1}}}$	2010-2011	0.2				
	2011-2012		1.4			
	2012-2013			2.5		
	2013-2014				0.3	
$\overline{\Delta_f}$	2010-2011, 2011-2012, 2012-2013, 2012-2013	$\overline{\Delta_f} = \frac{0.2 + 1.4 + 2.5 + 0.3}{4} = 1.1$				
σ	2010-2011, 2011-2012, 2012-2013, 2012-2013	$\sigma = \sqrt{\frac{(0.2-1.1)^2 + (1.4-1.1)^2 + (2.5-1.1)^2 + (0.3-1.1)^2}{4}} = 0.94$				
s_{f-nocs}	2010-2011, 2011-2012, 2012-2013, 2012-2013	$s_{f-nocs} = \varphi_2 / \overline{\Delta_{fs}} = \varphi_2 / \left(\frac{\overline{\Delta_f}}{\sigma}\right) = \frac{10}{\left(\frac{1.1}{0.94}\right)} = 8.54$				

Table 4. ICCAD stability computation

Parameter	Years	2010	2011	2012	2013	2014
v_f		4.1	5.7	6.1	2.3	6.5
$\Delta_{f_{y_j \to y_{j+1}}}$	2010-2011	1.6				
	2011-2012		0.4			
	2012-2013			3.8		
	2013-2014				4.2	
$\overline{\Delta_f}$	2010-2011, 2011-2012, 2012-2013, 2012-2013	$\overline{\Delta_f} = \frac{1.6 + 0.4 + 3.8 + 4.2}{4} = 2.5$				
σ	2010-2011, 2011-2012, 2012-2013, 2012-2013	$\sigma = \sqrt{\frac{(1.6-2.5)^2 + (0.4-2.5)^2 + (3.8-2.5)^2 + (4.2-2.5)^2}{4}} = 1.57$				
$s_{f-iccad}$	2010-11, 2011-12, 2012-13, 2012-13	$s_{f-iccad} = \varphi_2 / \overline{\Delta_{fs}} = \varphi_2 / \left(\frac{\overline{\Delta_f}}{\sigma}\right) = \frac{10}{\left(\frac{2.5}{1.57}\right)} = 6.28$				

The preceding discussion demonstrates the efficacy of using the standard deviation (σ) to normalize or standardize the average fairness change ($\overline{\Delta_f}$) before calculating the stability (S_f). Let us go through a simple example for illustrative purposes. For NOCS, an example change of 1.6 is significant compared to any of the calculated annual fairness values (*e.g.*, 2.2 in 2014). For ICCAD, where the values of fairness (*e.g.*, 6.2 in 2014) are larger, a change of 1.6 is not as significant. NOCS and ICCAD being similar conferences, let us now imagine a very different conference with an average fairness value ($\overline{v}_f$) of 51 - roughly an order of magnitude higher. For such a conference, a change of 7 (for example) in either direction is virtually insignificant compared to the fairness value itself. But, a change of 7 in either direction for NOCS and ICCAD would mean that something unexpected has happened. The imaginary conference with an average fairness of 51, however, will most likely have fluctuations that are far bigger than those of either NOCS or ICCAD, because fairness changes tend to scale relative to the fairness values themselves, thus giving this particular conference (with the average fairness of 51) a far higher standard deviation. Therefore, normalizing the data using the standard deviation gives the computed stability value ($s_{f-conference}$) the same meaning regardless of the dataset it was calculated with.

NOCS and ICCAD were initially chosen as the conferences for our analysis because they were popular technical conferences expected to be stable and fair. The overall stability values for these conferences, calculated over five years, show that NOCS and ICCAD, while not extremely so, are still unstable in the context of our definition of fairness; a visual analysis does not yield this level of insight that is provided by the computed stability and fairness numbers.

VII. Conclusion

We defined and analyzed *fairness* and *stability* of scientific collaborations in academic conferences by using co-authorship and institutional attributes of papers appearing in NOCS and ICCAD over a five year period (2010-2014). The conference data obtained from IEEE Xplore was converted into a network graph, from which we computed numerical values of fairness and stability for the two conferences, using a methodology that we proposed; the methodology makes use of the network properties of the collaboration network graph of the authors.

Our calculations led to interesting results that were not otherwise apparent from a simple visual analysis. When NOCS was initially selected as a conference for our analysis, we expected that the changes or fluctuations in its number of authors over the years would cause it to be unstable compared to ICCAD whose largely unchanging author-base would provide the conference great stability. However, our computed stability values showed that despite having similar number of authors every year, ICCAD was actually less stable than NOCS.

We analyzed only two conferences for five years in this paper, but we aim to investigate the properties of a larger number of conferences over a longer time period in order to study if shifts in the scientific community affect the fairness and the stability of a conference. It would be very interesting to include the information about external reviewers in our network graph and perform a more precise calculation of fairness, and a more in-depth analysis of stability. We believe that our techniques and results can be used to guide and optimize policies regarding the rules of engagement in different academic events and networks. Given the generality of our methodology, it would be interesting to extend our analysis to the examination of other social networks as well.

Acknowledgment

The authors would like to thank the Department of Electrical and Computer Engineering at Carnegie Mellon University for providing an opportunity to collaborate and pursue the investigation outlined in the paper.

References

[1] M. M. Danziger, A. Bashan, Y. Berezin, L. M. Shekhtman and S. Havlin, "An Introduction to Interdependent Networks," *Communications in Computer and Information Science,* vol. 438, pp. 189-202, 2014.

[2] G. Palla, A.-L. Barabasi and T. Vicsek, "Quantifying social group evolution," *Nature,* vol. 446, 2007.

[3] M. E. J. Newman, "The structure of scientific collaboration networks," *Proceedings of the National Academy of Sciences of the United States of America,* vol. 98, no. 2, pp. 404-409, 16 January 2001.

[4] A. Kadriu, "Discovering Value in Academic Social Networks: A Case Study in ResearchGate," in *Information Technology Interfaces (ITI),* 2013.

[5] L. A. Barabasi, H. Jeong, Z. Neda, E. Ravasz, A. Schubert and T. Vicsek, "Evolution of the social network of," *Physica A: Statistical Mechanics and its Applications,* vol. 311, pp. 590-614, 2002.

[6] "IEEE Xplore," [Online]. Available: http://ieeexplore.ieee.org/Xplore/home.jsp.

[7] "The Open Graph Viz Platform," [Online]. Available: http://gephi.github.io/ .

[8] M. Jacomy, T. Venturini, S. Heymann and M. Bastian, "ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software," *PLOS One,* 10 June 2014.

[9] Y. Hu, "Efficient, High-Quality Force-Directed Graph Drawing," *The Mathematica Journal,* vol. 10, no. 1, 2006.

[10] S. Martin, W. M. Brown, R. Klavans and K. W. Boyack, "OpenOrd: an open-source toolbox for large graph layout," in *SPIE*, 2011.

[11] T. M. J. Fruchterman and E. M. Reingold, "Graph Drawing by Force-directed Placement," *Software - Practice and Experience,* vol. 21, no. 11, pp. 1129-1164, November 1991.

[12] A. Barão, "The Gephi Network Splitter 3D Layout," 2014. [Online]. Available: http://www.relationalcapitalvalue.com/gephiplugins.html.