Co-authorship and community structure in the DAFx conference proceedings : 1998–2016

Wilson, AD

CO-AUTHORSHIP AND COMMUNITY STRUCTURE IN THE DAFX CONFERENCE PROCEEDINGS: 1998–2016

Alex Wilson

Acoustics Research Centre,
University of Salford
Greater Manchester, UK
a.wilson1@edu.salford.ac.uk

ABSTRACT

This paper presents the co-authorship network of the DAFx conference series, from its inception in 1998 to the present, along with subsequent analysis. In total 1,281 unique authors have contributed 1,175 unique submissions to this conference series. The co-authorship network is revealed to contain a large weakly connected component containing 667 authors (≈52% of the total network). The size of this component compares well to previous studies of other conference series of similar age and scope. Within this connected component, 24 communities were detected using the Louvain method. While some communities have formed based on geographic proximity, links between communities are observed. This shows a high level of collaboration in the network, possibly due to the speciality of the conference and the movement of academics throughout Europe.

1. INTRODUCTION

The DAFx conference series, which began in 1998 as a workshop series, evolved from the Action-G6 of the European COST programme, named "digital audio effects". Its objectives were originally described as follows [1].

1. To compare the different methods of the European teams in terms of algorithms, implementations and musical use.
2. To bring together the knowledge of different European teams in the domain of digital sound processing, here designated as "digital audio effects", in a form which can be made available inside and outside of the teams themselves.

This second objective refers to the development of research networks, within Europe and also beyond. This paper aims to examine the breath and structure of the resultant network after a period of two decades. To this end a co-authorship network was created, indicating which authors of DAFx submissions have authored papers together. From this, the nature of collaboration within the DAFx community was examined. More clearly, the goals of this study were the following:

1. Collect bibliographic data from the entire DAFx conference series, 1998–2016
2. Create co-authorship network from this data
3. Identify connected components
4. Identify communities within the largest connected component

[1] http://www.consilium.europa.eu/uedocs/cms_data/docs/dynadoc/out/cost/en/cost_at_g6.pdf

2. BACKGROUND AND LITERATURE REVIEW

A graph is a structure which describes the relationships between a set of nodes. The links between nodes are called edges. Graphs are often categorised as being undirected or directed: a graph can be called undirected when the edges between pairs of nodes have no directional information, or directed when the edge describes a one-way connection between nodes. This paper will consider co-authorship graphs exclusively.

2.1. Co-authorship networks

Co-authorship is one of the more frequently-investigated forms of scientific collaboration. It has been noted that the proportion of single-authored papers, across various scientific disciplines, has decreased, from 25% in 1980 to 11% in 2000 [1]. Multi-author papers also gain more citations, including self-citations [2].

Scientific collaboration networks are typically represented as *undirected, unweighted* graphs, i.e. what is represented is a simple binary classification of whether or not two individuals have collaborated on a paper. This means that a lot of important information is usually absent, such as the number of co-authored works by the pair, and the assumption is made that the partnership is completely equal. In real collaboration networks one can often see that a pair will co-author numerous works or that the share of work is not equal. When the number of co-authored works is included it is typically represented by the weight of an edge between two nodes. It has been shown that *weighted, undirected* co-authorship networks have a high correlation with social networks, themselves influenced by geographic proximity [3, 4].

A number of strategies exist for creating *directed* co-authorship networks, which encode information about the partnership between the authors — which author should be given priority. These include "first author takes all" or "last author takes all" strategies but these both assume no relationships between intermediate authors in the list of authors. More nuanced approaches have been developed, which include interactions between intermediate authors in the author list [5].

The connected components of a graph G are the set of largest subgraphs of G that are each connected. A co-author network may consist of many individual connected components. A *bridge* is a node whose removal would cause the number of components to increase. In a co-authorship network, this is an author who has co-authored works with member of otherwise disparate and unconnected communities. These bridges can be formed as a result of the movement of researchers from one research group to another.

These concepts are often best understood in the context of ones own discipline [6]. Hence, this paper is concerned with the total co-authorship network of the DAFx proceedings and pays particular attention to the largest connected component.

2.2. Bibliometrics of DAFx proceedings 1998–2009

A previous submission to the DAFx conference series examined the bibliometrics of the series, for the first twelve years [7]. The following is a brief summary of that work.

- Background to the DAFx conference series — list of locations and organisers
- There had been 722 submissions by 767 unique authors.
- Number of submissions per year, individual and cumulative
- Authorship distribution — number of authors per paper (the modal value was 2)
- A list of the most cited papers
- A list of the 20 most frequent authors
- Confirmation that the conference submissions followed Lotka's law, shown by a log-log plot of publications against number of authors.

The current submission attempts not to repeat any of these contributions, save for necessary updates to data after an additional seven years of submissions.

2.3. Bibliometrics & scientometrics of other conference series

Another paper by the author of the initial DAFx study focussed on the IEEE Transactions on Software Engineering [8]. This included further types of analysis not reported in the study of DAFx, such as collaboration between countries. Of course patterns of international collaboration can change with time. When considering the period of 1990-2000, the growth of the European Union and its funding for science was credited as a significant change in the scientific environment [9]. It was in this climate that the DAFx series began.

After its first nine years, a bibliometric analysis of the ISMIR (International Society for Music Information Retrieval) conference series was published [10]. This included a limited co-author network, in which only 22 authors were labelled. This paper suggested that the European research labs were tightly interconnected. The ISMIR community does contain some overlap with DAFx, in terms of authors. A recent follow-up examined the role of female authors in the ISMIR conference proceedings [11]. At this time, after 16 years, the total number of unique authors was 1,910. This paper included a brief co-author analysis. Nine clusters of authors are shown — clusters containing female authors with at least five co-authors. This showed that the two most prolific female authors were members of the same cluster. However, it is not clear whether these clusters can be connected to one another, i.e. whether they were drawn from a large connected component or numerous connected components. It can be inferred from [10] that some connections can be made (as a component cannot become smaller over time), yet insufficient data is presented.

It has been reported that the main component in the co-authorship network of the PACIS (Pacific Asia Conference on Information Systems) reached a size of 663 authors after a period of 15 years, which accounted for 33% of the total network [12]. This compared well to a value of 29% in a study of the ECIS (European Conference on Information Systems) after a period of twelve years [13]. Within ECIS, the second largest component contained only 37 authors, indicating how the main component grows by merging with smaller components, through the process of collaboration.

Herein, comparative data is presented for the DAFx conference series, after 19 years. The number of authors in the largest connected component was of particular interest, as was the detection of communities within that component.

3. DATA COLLECTION AND PROCESSING

Bibtex files from each of the first 12 conferences were already available [7]. Unfortunately, the file for DAFx-98 only listed one author per paper — this was manually corrected. JabRef[2] was used for Bibtex editing. Bibtex files for the seven subsequent conferences were created manually for this paper, from the conference proceedings listed on each conference website. Several editions of the DAFx conference proceedings (from 2005 to 2012) are indexed on Scopus and DAFx-03 is indexed on Web Of Science. These entries were checked against the manually created entries.

From the combined proceedings it was possible to construct a co-authorship network, showing which authors had directly collaborated in the writing of a paper. The following is a description of the process by which the network was created from the Bibtex data.

- All individual Bibtex files were merged and this file was imported into Network Workbench (NWB), a tool for network analysis and scientometrics [14]. The co-author network was extracted from the Bibtex file using the supplied routine.
- The names of authors will not always be consistent throughout all of their publications. This can be due to use of initials in place of full names, deliberate changes in name, reversal of first-name/family-name conventions, or simply human error in transcription. The merging of duplicate nodes is crucial for the accurate determination of graph metrics. Possible duplicates were highlighted by applying the Jaro distance metric [15, 16]. Through trial and error, authors were merged at 0.92 similarity providing the first two letters were in common, and it was noted when two entries were above 0.85 similarity. The data was then manually inspected. Where node labels only contained a first initial and not the full name, Google Scholar was used to identify whether this node matched any others. If 'both' authors had similar co-authors and wrote about similar topics, then the likelihood of them being duplicates was considered great enough to make a correction.
- The network was updated by merging nodes that were flagged as duplicates.
- The graph was split into separate, comma-separated files: one for nodes and one for edges.
- These node and edge tables were imported into Gephi [17]. Visual inspection of the graph was able to identify further duplicate nodes. An iterative approach was taken to the correction of duplicate nodes, by repeating these steps until all had been accounted for. The need for an iterative approach to data cleaning has also been described previously [12]. Of course, these methods were not able to detect any deliberate changes in name, such as by marriage.

[2] `http://www.jabref.org/`

Table 1: *Summary statistics of entire co-author network*

Number of nodes	1281
Number of edges	2058
Average degree	3.213
Average weighted degree	3.911
Connected components	269
...which are isolates	132
...which are dyads	58
...which are triads	37
The largest connected component consists of	667 nodes.

With the final network represented as a list of nodes and a list of the edges between them, this data was loaded in Gephi, for further processing to determine the connected components and perform community detection. Connected components were found using a depth-first search method [18]. Community detection was performed using the Louvain method [19]. With this additional data added as node attributes, the full set of nodes and edges was exported to .CSV files. Further processing and visualsation was performed using Matlab R2016a. In all calculations, the network was assumed to be undirected, i.e. the order of co-authors in a paper was not considered at this time.

4. RESULTS

After 19 years of conference proceedings, the number of unique authors has reached 1,281, from 1,175 submissions authored. Table 1 displays a summary of the total network. For all graph plots (Figures 1 and 2), node positioning was achieved using a force-directed layout [20]. As shown in Table 1, there are 269 connected components but this number includes 132 isolates (authors with degree = 0), 58 dyads (a pair of authors, each with degree = 1), 37 triads (three authors all connected to one another, with degree = 2) and other such small, highly-connected groups. Such a connected component can represent an individual paper, such as one paper submitted in 2012 which had 8 authors. Where a component represents a single paper, the component will be a complete graph, as each node connects to each other. The largest connected component (shown in Figure 1) contains 667 nodes, over half of the total nodes, making it roughly 44 times larger than the next largest connected component. This suggests that...

a) new contributors to the conference proceedings are authors who are known to other authors in the network, such as their new students or colleagues.

b) when smaller components begin to form it is not long before they merge with the main component. This forming of new collaborative bonds would be a natural consequence of authors meeting at conferences. The next largest connected components contain less than 15 nodes — there may be a critical mass a component reaches before it joins with another.

As shown in Table 1, there are 132 isolates in the network. An isolate is a node with a degree of 0, i.e. an author with no co-authors. In this conference, isolates make up roughly 10% of all authors. This number includes many of the keynote and tutorial submissions, which are usually credited to one author, frequently a local author from a related field not directly involved in DAFx. With an mean value of 62 papers per year, three keynote speakers and three tutorial presenters would ensure a figure of 10% isolates, were these speakers to be unique each year. Of course, the creation of research networks takes time, and so authors whose first DAFx submissions were single-authored and took place recently may remain isolates for a number of years. Additionally, some authors prefer to work without co-authors. Overall, the importance of these contributions should not be discounted. Only seven authors have made more than one contribution without having had a co-author. These are listed in Table 2.

Table 2: *Isolates (authors without co-authors) who have made more than one submission.*

Name	N_{works}
Sinan Bökesoy	3
Niels Bogaards	2
Christian Müller-Tomfelde	2
David Kim-Boyle	2
Richard Hoffmann-Burchardi	2
Angelo Farina	2
Tor Halmrast	2

4.1. Community detection

Within the main component, 24 communities were uncovered using the Louvain method [19], having between 4 and 88 members. Qualitative analysis of these communities reveals a clear geographic influence on collaborative patterns.

- Figure 2a displays the largest community, formed by the collaborations between some of the most frequently-contributing authors. Many USA-based authors are members of this community.
- Figure 2b appears to show pre-predominately researchers based in France and Canada. As shown in Fig. 1, this community could be broken down into smaller sub-communities (or 'cliques') in each continent..
- Figure 2c contains many individuals who were affiliated with Queen Mary University of London at the time of submission.
- Figure 2d describes a community of predominately French researchers and individuals with whom they have collaborated while based in France. In contrast to Fig. 2b, this community is focussed on IRCAM in Paris.

Each of these communities has hosted a DAFx conference, to which their large number of nodes can be at least partly-attributed (or *vice-versa*). Naturally, while the centre of each community may show a strong geographic influence, less-frequent collaborators in other regions are located further from the centre. Geographic proximity does facilitate academic collaboration (as described in section 2.1) but it is one of a number of factors.

Concerning the origins of the conference series as a means of disseminating knowledge within Europe, it can be seen in Figure 1 that a number of the 24 communities detected in the main component are of researchers beyond the continent. As mentioned above, the larest communities contain many North American-based authors. Additionally, the 17th largest community comprises of authors based in Taiwan, and is connected to the rest of the main

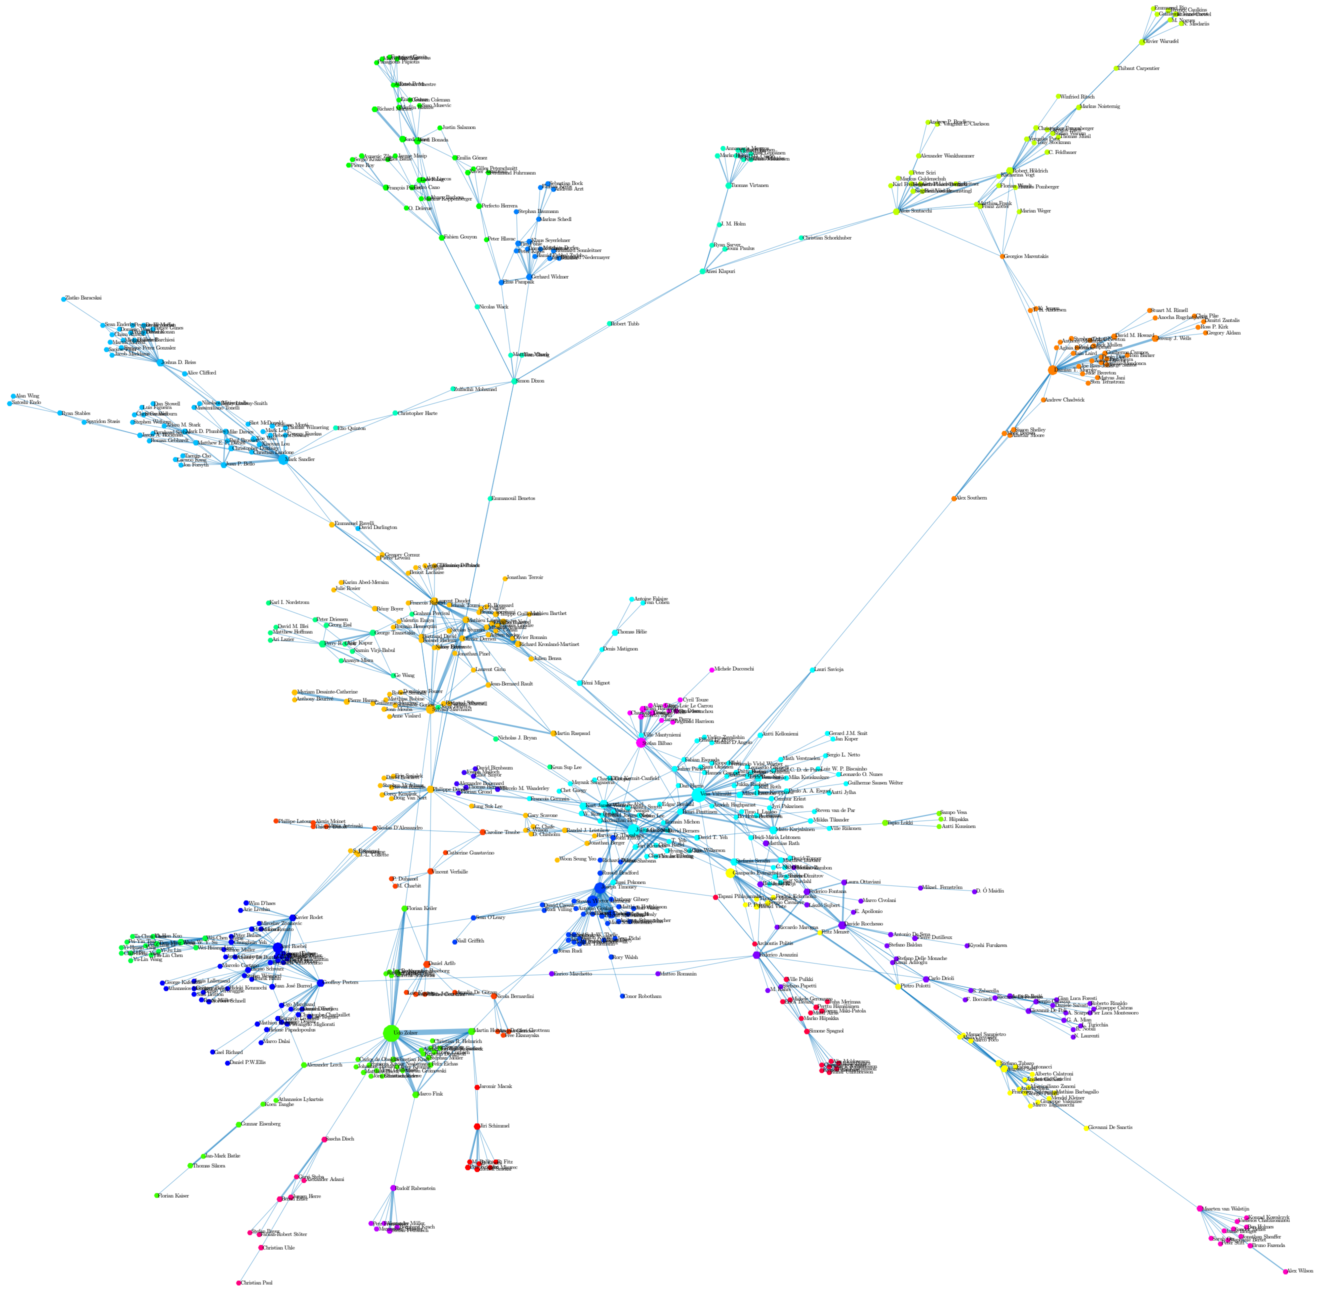

Figure 1: *Largest connected component, consisting of 667 nodes. Node size is proportional to number of works created by that author. Edge thickness is proportional to the number of co-authored works between nodes. Colour represents the communities detected using the Louvain algorithm [19]. Node positioning was achieved using a force-directed layout [20].*

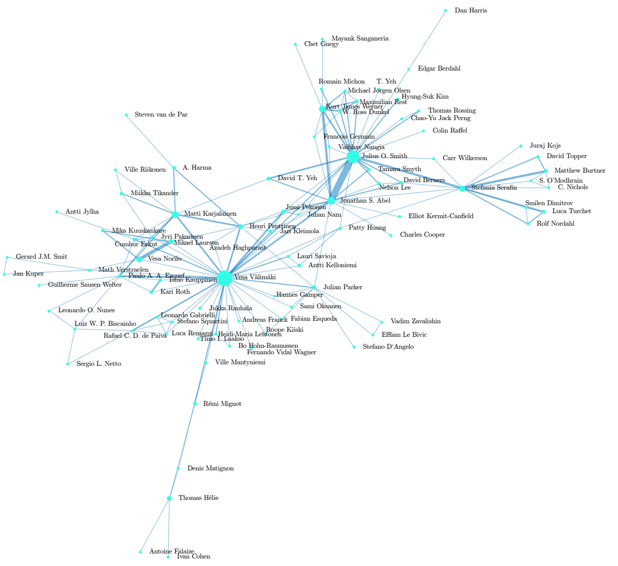

(a) $N_{\text{nodes}} = 88$

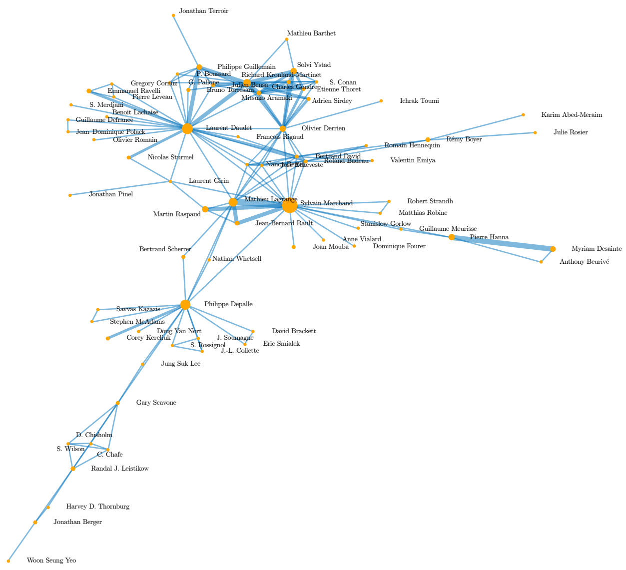

(b) $N_{\text{nodes}} = 73$

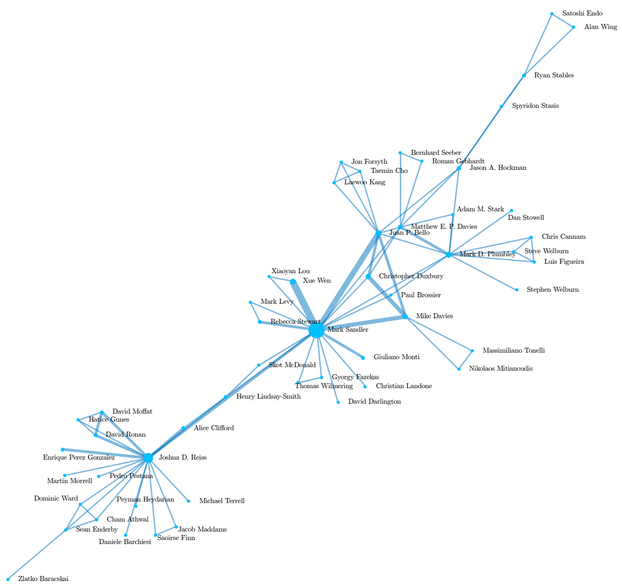

(c) $N_{\text{nodes}} = 53$

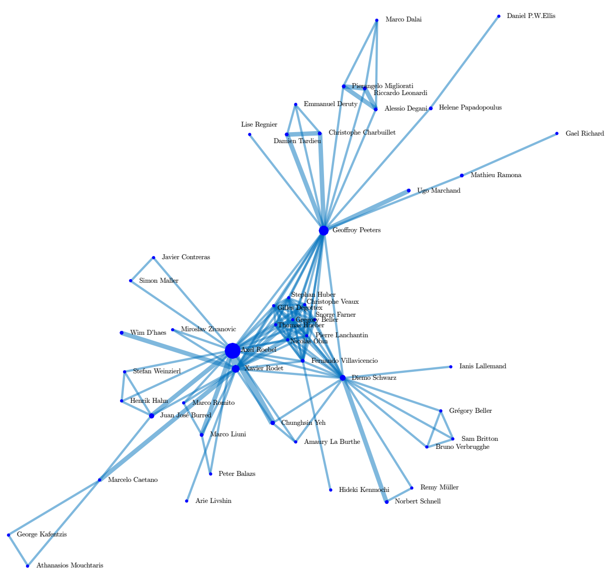

(d) $N_{\text{nodes}} = 49$

Figure 2: *The four largest communities detected within the main component. Note that most nodes shown here have many more edges beyond these communities. Colours used are as in Figure 1. Node positioning was achieved using a force-directed layout [20], run separately for each community.*

Table 3: *Top 10 authors (within the main component) by weighted degree*

Name	N_{works}	Degree	Weighted degree
Vesa Välimäki	36	40	62
Julius O. Smith	28	29	58
Joseph Timoney	24	22	56
Udo Zölzer	44	29	55
Victor Lazzarrini	24	28	54
Jonathan S. Abel	15	22	42
Damian T. Murphy	18	27	36
Laurent Daudet	11	28	35
Mark Sandler	21	22	34
Axel Röbel	20	27	34

component via connections to the community in Fig. 2d. Of course, the main component shown in Figure 1 contains only 52% of the total number of authors. While the other connected components are relatively small, each containing less than 15 nodes, these include active research communities in Japan and China.

4.2. Node degree

In an undirected graph, the *degree* of a node refers to the number of other nodes to which it is connected. In a co-authorship network this is simply the number of co-authors. When ranking co-authors by degree this gives (perhaps) undue preference to authors who have managed to author many works but perhaps contributing little to each. If the edge weights are considered, then the *weighted degree* takes into account the number of times a pair of co-authors have worked together. Table 3 shows the ten authors with the greatest weighted degree.

4.3. Centrality

In attempting to measure the influence of nodes within a network a number of centrality measures have been developed [21]. This section will report on three of these: *closeness*, *betweenness* and *eigenvector centrality*. These three scores were calculated using Matlab R2016a.

4.3.1. Closeness

Closeness centrality uses the inverse sum of the distance from a node to all other nodes in the graph. Assuming that not all nodes can be reached (which is true for this network), the centrality of node i is:

$$\text{closeness}(i) = \left(\frac{A_i}{N-1}\right)^2 \frac{1}{D_i} \tag{1}$$

Here, A_i is the number of nodes that can be reached from node i (not counting i itself), N is the number of nodes in the graph G, and D_i is the sum of distances from node i to all reachable nodes. If no nodes are reachable from node i, then closeness(i) is zero. This expression assumes that all edge weights are equal to 1. The reciprocal of the actual edge weights (the number of co-authored works) were introduced as the ‘cost’ used in the centrality calculations. This is suitable because one can deduce that co-authors with many co-authored works exchange information more readily

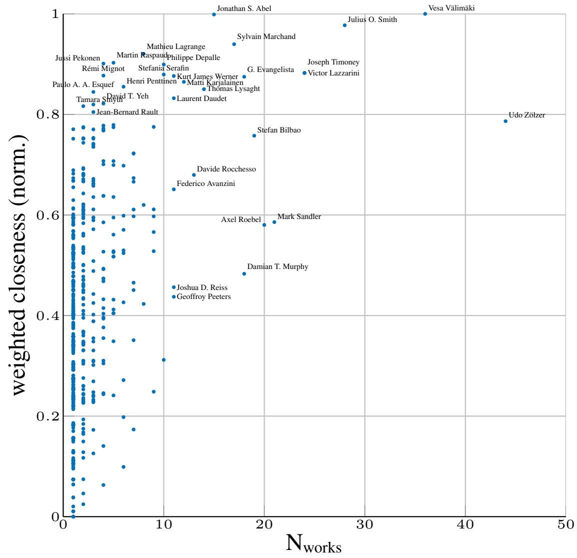

Figure 3: *Scatterplot showing number of submissions vs. weighted closeness, which is normalised to the range [0 1]. The greater the value of closeness, the fewer steps (on average) are required to reach another author in the network.*

than those with fewer co-authored works, and so the ‘cost’ associated with traversing this edge is lower. The authors with the highest weighted closeness centrality values are displayed in Fig. 3. Closeness generally increases with the number of submissions but there are notable exceptions, as highlighted.

4.3.2. Betweenness

The betweenness measure illustrates the importance of an author by means of assessing the flow of “traffic” that passes through that node. This is achieved by measuring the number of shortest paths from all nodes to all other nodes which involve passing through the node in question.

$$\text{betweenness}(i) = \sum_{s,t \neq i} \frac{n_{st}(i)}{N_{st}} \tag{2}$$

$n_{st}(u)$ is the number of shortest paths from source s to target t that pass through node i, and N_{st} is the total number of shortest paths from s to t. As with closeness, the reciprocal of the edge weights was used as a cost in calculating weighted betweenness. The values of weighted betweenness centrality are displayed in Fig. 4.

4.3.3. Eigenvector centrality

The eigenvector centrality measure assumes that when a node is connected to other high-scoring nodes that this counts for more than a connection to a lower-scoring node.

When eigenvector centrality was computed without weighting, the top 10 authors were not just all members of the same community (shown in Fig. 2d) but all co-authors of the same publication. For many of these authors, this has been their sole contribution to DAFx. When eigenvector centrality is calculated with edge weights taken into account, the results are plotted in Fig. 5.

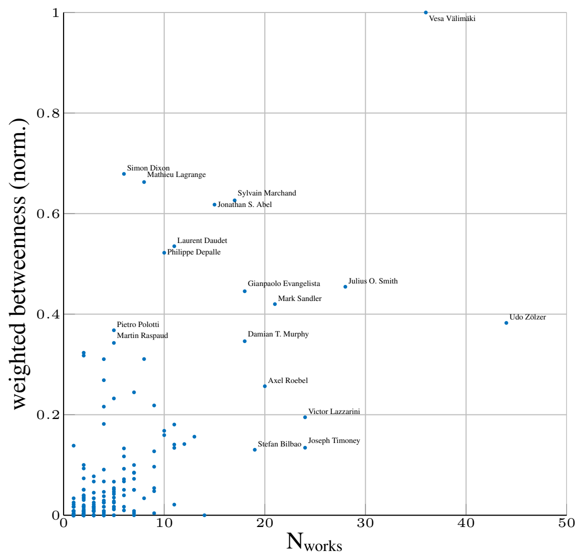

Figure 4: *Scatterplot showing number of submissions vs. weighted betweenness, which is normalised to the range [0 1]. High betweenness scores indicate an author who co-authors submissions with a large number of communities.*

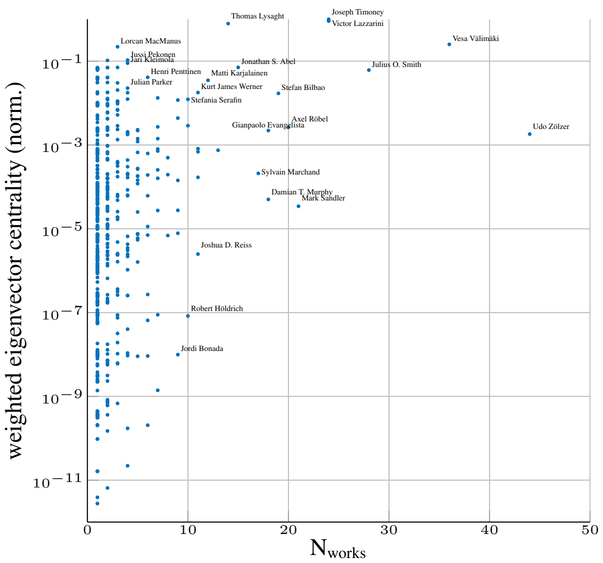

Figure 5: *Scatterplot showing number of submissions vs. weighted eigenvector centrality, normalised to the scale [0 1]. Note the use of logarithmic scaling.*

The three authors with the highest score according to this metric are all from Maynooth University, and have frequently collaborated. There are many infrequent collaborators with relatively high scores, while some frequently contributing authors in more isolated communities (according to this metric) are also highlighted.

5. DISCUSSION

The presentation in this paper of the DAFx co-authorship network allows for a number of possibilities: an author may use the network to assess their own contributions to the conference series and the contributions of their research group(s). In doing so, new collaborative opportunities can be located.

There is still an open question as to whether the density of co-authorship networks is increased by the formation of long-term collaborative bonds or if it is occasional and short-term collaborative efforts that causes the network to expand. At this stage, a series of questions is posed to the community, for discussion at conferences and beyond.

1. How could the network be used to identify potential collaborators?
2. Could more be done to integrate isolated communities?
3. There are very few frequently-contributing authors who lack co-authors. Could more be done to support authors who do not collaborate?
4. Would the size of the network be increased by the use of a double-blind review process, in which submissions are evaluated without knowing the names of the authors?

As shown by Table 3 and Figs. 3, 4 and 5, the relationship between these measures of node importance and more straightforward measures, such as the number of works, is not clear. It is possible for an author with few works to be considered highly important to the network as a whole. Each centrality measure has strengths and weaknesses and provides different insights into the topology of the network. Closeness centrality rewards authors who are prominent in communities which are themselves well connected to others, often achieved by high degree scores — the top three authors are members of the community in Fig. 2a. Betweenness centrality highlighted the efforts of a number of authors who have worked with a number of communities and authors who act as bridges — six communities are represented in the top 10 authors. In contrast, eigenvector centrality rewards groups of authors who frequently collaborate with one another.

Figures 1 and 2d show that large communities are not always so well-connected to the rest of the network. Does this indicate that these communities are large enough to be self-sustaining and in less need of outward collaboration?

One limitation in this study is that author order was not taken into account. Doing so would result in a directed network, allowing for a more sophisticated analysis particularly regarding centrality measures. However, establishing the relative contributions of each author in a paper is not a trivial task. As described in section 2.1, there are a variety of strategies which can be employed, but this task remains a focus of on-going study.

6. CONCLUSIONS

The aim of the study was to examine the nature of community and collaboration within the DAFx conference proceedings, after two decades. By collecting Bibtex archives for each conference a co-authorship network was created. This revealed a large connected component — 52% of all authors are connected to one another by a number of intermediate co-authors. Twenty-four communities were detected within this component, heavily influenced by geographical proximity. Communities are connected by the movement of researchers between research groups and by the interactions and discussions at conferences. This network could be displayed online, allowing authors interactive access to the data. This would also facilitate the updating of the data with each new set of conference proceedings.

There have now been two papers explicitly describing the DAFx conference proceedings: one on basic bibliometrics and one describing the details of the network. There is scope for further work, in the context of DAFx and also more generally. The Bibtex entries for the DAFx conference could contain information on submission type: keynote, tutorial, oral presentation or poster presentation. It would be interesting to examine whether the choice of poster or oral presentation has an impact on the formation of collaborative links. Of course, DAFx does not exist in isolation, and its contributors also make submissions to other conference series. Larger co-authorship networks can be constructed by the merger of related conference proceedings. This could reveal the extent of the overlap and the interdisciplinary nature of the research groups involved.

7. REFERENCES

[1] Wolfgang Glänzel and András Schubert, “Analysing scientific networks through co-authorship,” in *Handbook of quantitative science and technology research*, pp. 257–276. Springer, 2004.

[2] Wolfgang Glänzel and Bart Thijs, “Does co-authorship inflate the share of self-citations?,” *Scientometrics*, vol. 61, no. 3, pp. 395–404, 2004.

[3] Katy Börner, Shashikant Penumarthy, Mark Meiss, and Weimao Ke, “Mapping the diffusion of scholarly knowledge among major US research institutions,” *Scientometrics*, vol. 68, no. 3, pp. 415–426, 2006.

[4] Howard D White, Barry Wellman, and Nancy Nazer, “Does citation reflect social structure?: Longitudinal evidence from the globenet interdisciplinary research group,” *Journal of the American Society for information Science and Technology*, vol. 55, no. 2, pp. 111–126, 2004.

[5] Jinseok Kim and Jana Diesner, “A network-based approach to coauthorship credit allocation,” *Scientometrics*, vol. 101, no. 1, pp. 587–602, 2014.

[6] Howard D White and Katherine W McCain, “Visualizing a discipline: An author co-citation analysis of information science, 1972-1995,” *Journal of the American society for information science*, vol. 49, no. 4, pp. 327–355, 1998.

[7] Brahim Hamadicharef, “Bibliometric study of the DAFx proceedings 1998-2009,” in *Proceedings of International Conference on Digital Audio Effects*, 2010, pp. 427–430.

[8] Brahim Hamadicharef, “Scientometric study of the IEEE transactions on software engineering 1980-2010,” in *Proceedings of the 2nd International Congress on Computer Applications and Computational Science*. Springer, 2012, pp. 101–106.

[9] Caroline S Wagner and Loet Leydesdorff, “Mapping the network of global science: comparing international coauthorships from 1990 to 2000,” *International journal of Technology and Globalisation*, vol. 1, no. 2, pp. 185–208, 2005.

[10] Jin Ha Lee, M Cameron Jones, and J Stephen Downie, “An analysis of ISMIR proceedings: Patterns of authorship, topic, and citation.,” in *Proceedings of the 10th International Conference on Music Information Retrieval*, 2009, pp. 57–62.

[11] Xiao Hu, Kahyun Choi, Jin Ha Lee, Audrey Laplante, Yun Hao, Sally Jo Cunningham, and J Stephen Downie, “WiMIR: An informetric study on women authors in ISMIR,” in *Proceedings of the 17th International Conference on Music Information Retrieval (ISMIR)*, New York, USA, 2016.

[12] France Cheong and Brian J Corbitt, “A social network analysis of the co-authorship network of the pacific asia conference on information systems from 1993 to 2008,” *PACIS Proceedings*, 2009.

[13] Richard Vidgen, Stephan Henneberg, and Peter Naudé, “What sort of community is the european conference on information systems? a social network analysis 1993–2005,” *European Journal of Information Systems*, vol. 16, no. 1, pp. 5–19, 2007.

[14] NWB Team et al., “Network workbench tool. indiana university, northeastern university, and university of michigan,” 2006.

[15] Matthew A Jaro, “Advances in record-linkage methodology as applied to matching the 1985 census of Tampa, Florida,” *Journal of the American Statistical Association*, vol. 84, no. 406, pp. 414–420, 1989.

[16] Matthew A Jaro, “Probabilistic linkage of large public health data files,” *Statistics in medicine*, vol. 14, no. 5-7, pp. 491–498, 1995.

[17] Mathieu Bastian, Sebastien Heymann, Mathieu Jacomy, et al., “Gephi: an open source software for exploring and manipulating networks.,” *ICWSM*, vol. 8, pp. 361–362, 2009.

[18] Robert Tarjan, “Depth-first search and linear graph algorithms,” *SIAM journal on computing*, vol. 1, no. 2, pp. 146–160, 1972.

[19] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre, “Fast unfolding of communities in large networks,” *Journal of statistical mechanics: theory and experiment*, , no. 10, 2008.

[20] Thomas MJ Fruchterman and Edward M Reingold, “Graph drawing by force-directed placement,” *Software: Practice and experience*, vol. 21, no. 11, pp. 1129–1164, 1991.

[21] Tore Opsahl, Filip Agneessens, and John Skvoretz, “Node centrality in weighted networks: Generalizing degree and shortest paths,” *Social networks*, vol. 32, no. 3, pp. 245–251, 2010.