

# A Decade of Database Research Publications

Sherif Sakr<sup>1</sup>    Mohammad Alomari<sup>2</sup>

<sup>1</sup> NICTA and University of New South Wales, Australia  
[ssakr@cse.unsw.edu.au](mailto:ssakr@cse.unsw.edu.au)

<sup>2</sup> University of Sydney, Australia  
[miomari@it.usyd.edu.au](mailto:miomari@it.usyd.edu.au)

**Technical Report**  
**UNSW-CSE-TR-1022**  
**November 2010**

THE UNIVERSITY OF  
NEW SOUTH WALES



School of Computer Science and Engineering  
The University of New South Wales  
Sydney 2052, Australia

## **Abstract**

We analyze the database research publications of four major core database technology conferences (SIGMOD, VLDB, ICDE, EDBT), two main theoretical database conferences (PODS, ICDT) and three database journals (TODS, VLDB Journal, TKDE) over a period of 10 years (2001 - 2010). Our analysis considers only regular papers as we do not include short papers, demo papers, posters, tutorials or panels into our statistics. We rank the research scholars according to their number of publication in each conference/journal separately and in combined. We also report about the growth in the number of research publications and the size of the research community in the last decade.

# 1 Introduction

The database management technology has played a vital role in the advancements of the information technology field. Database researchers are one of the key players and main sources to the growth of the database systems. They are playing a foundational role in creating the technological infrastructure from which database advancements evolve. The impact of research scholars in the community is often measured by their number of publications in top-tier research venues and the number of citations they receive, i.e. how frequently their publications are referenced by other publications (e.g. H-index [11], g-index [6]). In principle, there is a direct relationship between the tier rank of a research venue and its number of citations which is commonly determined as the *impact factor* [3]. The success of a research scholar in publishing his research results in a top-tier venue increases his chances of having his work being widely received by his peers in the community and consequently to be more frequently cited by them.

In general, achieving an accurate, fair and insightful citation-based analysis is a very challenging task due to the difficulty of parsing and extracting the citation meta data from the research articles. Recently, some online services have been introduced to capture the citation information of research publications (e.g. MS Libra<sup>1</sup>, Google Scholar<sup>2</sup>). However, the information provided by these services suffer from some anomalies such as: incompleteness and duplication. Therefore, preparing a high quality citation information for a pool of research publications requires an extensive amount of manual labor work. Moreover, citation-based analysis methods tend to consider only the explicit citation relationships as indicated in the reference parts of the articles. In practice, it is impossible for authors of any article (including this one) to cite all the related publications of their work but they are normally only able to cite only a fraction of them. Therefore, the final decision of selecting the set of papers to be referenced usually depends on many scientific and non-scientific factors. For example, it has been shown that citations tend to have problems like biased-citation, self-citation, or positive vs. negative citation [13, 19]. One common situation is that article introductions are usually citing related survey papers. Therefore, survey papers usually have citation counts that are many times more than any original work in its corresponding topic (e.g. according to Google Scholar, at the time of writing this paper, the two surveys: [9] has 883 citations and [15] has 2169 citations). Some studies have also shown that different citation choices correspond to different citation impact [17].

Complementary to a previous work which mainly considered ranking the research scholars based on their citation counts [16], in this paper, we focus on ranking the research scholars by the count of their research publications in top-tier venues. We selected a set of top-tier database research venues which are generally considered as the most representative, influential and prestigious in the database community. In particular, we analyzed the database research publications of four major core database technology conferences (SIGMOD, VLDB, ICDE, EDBT), two main theoretical database conferences (PODS, ICDT) and three database journals (TODS, VLDB Journal, IEEE TKDE) over a 10 years period (2001 - 2010). In general, we believe that research fields are better pre-

---

<sup>1</sup><http://academic.research.microsoft.com/>

<sup>2</sup><http://scholar.google.com>.

sented by their own venues rather than by multi-disciplinary venues. Therefore, we did not include some important conferences (e.g. CIKM, WWW) and journals (e.g. Information Systems) in the scope of this study.

In principle, some could argue that the number of publications may have become a less insightful or less significant metric due to the explosion of the number of conferences and journals in recent years [12]. Therefore, to remedy this argument, we considered only top-tier venues which are well-known with their very low *acceptance rates*. These *prestigious* venues are conducting highly selective review processes that mainly aims of ensuring that they are turning out high quality papers. Hence, these papers are usually expected to attract considerable attention (and citations) from other researchers in the community [5]. In fact, the distribution of our selected venues (6 conferences and 3 journals) is compatible with the fact that database researchers - and computer scientists in general - are considering prestigious conferences as favorite tools for presenting original research work in contrast to the general case of many other scientific disciplines where journal papers are routinely considered to be superior than conference papers [2, 8]. For example, it has been shown that the two top database conferences (SIGMOD and VLDB) receive many more citations per paper than the two top database journals (TODS and VLDB J.) [16]. In practice, the general culture in the computer science community is that journal papers are used to present deeper versions of papers that already have been presented at conferences. One of the main reasons behind this is that the review process of journal papers are usually very long. The *turnaround time* (the interval between the submission date of a manuscript and the date of having the editorial decision) for conferences is often less than a third of that of journals [18]. Since the field of computer science research tends to be fast paced, conferences provide a great chance for *timestamping* the latest research findings earlier which allows the knowledge to be publicly shared more rapidly.

In general, we are witnessing a continuous growth in the database field. That is mainly due to the continuous introduction of new application domains (e.g. web applications, mobile applications, cloud computing, sensor networks) with varying features and requirements on their data management aspects. In practice, data has become mobile, flexible, mirrored in a variety of logical and physical forms, evolving, being concurrently modified and replicated, dynamically generated and later reintegrated in very large repositories for further analysis and processing [7]. Therefore, there are many more researchers are entering the field to tackle these challenges and hence more research papers are being published. In this paper, we also study the growth rate on the size of contributing research community and the number of research publications in the last decade.

The input data of this study has been extracted from the XML records of the famous DBLP computer science bibliography<sup>3</sup>. Our analysis considers only regular papers as we do not include short papers, demo papers, posters, tutorials or panels into our statistics. We made the detailed results of our study accessible on the web<sup>4</sup>

---

<sup>3</sup><http://dblp.uni-trier.de/xml/>

<sup>4</sup><http://www.cse.unsw.edu.au/~ssakr/DBStatistics/index.html>

## 2 Study Results

### 2.1 Top Publishers of Database Research Venues


As we previously stated, in this study, we focus on measuring the number of publications in top-tier publication venues as one of the main indicators to evaluate the impact of a research scholar in the community and the quality of his research production. In this paper, we present the most important results of our study. For full detailed results, we refer the reader to the web page of this study.

Figures 2.1, 2.2, 2.3 illustrate the top publishers of the database research venues during the period between 2001 and 2010. Figure 2.1 represents the top publishers of the core database technology conferences: VLDB (Figure 2.1(a)), SIGMOD (Figure 2.1(b)), ICDE (Figure 2.1(c)) and EDBT (Figure 2.1(d)). Figure 2.2 represents the top publishers of the theoretical database conferences: PODS (Figure 2.2(a)) and ICDT (Figure 2.2(b)). Figure 2.3 represents the top publishers of the main database journals: VLDB journal (Figure 2.3(a)), TODS journal (Figure 2.3(b)) and TKDE (Figure 2.3(c)). The research scholars in these figures can be indicated with one of the following two symbols:


- The (+) symbol indicates that the research scholar appears on the correspondingly top publishers list of the same research venue for the former decade (1991 - 2000).
- The (\*) symbol indicates that the research scholar appears on the *ultimate* top publishers list of the same research venue in all of its editions since its origin.

For example, in Figure 2.1(a), *Divesh Srivastava* and *H. V. Jagadish* are indicated that they appear in the top publishers of the VLDB conference since its origin (1975 - 2010). However, only *H. V. Jagadish* is indicated that he appears on top publishers list of the VLDB conference on the former decade. Figure 2.4 illustrates aggregate lists of the top publishers for database research venues according to their focus: core database technology conference (Figure 2.4(a)), theoretical database conferences (Figure 2.4(b)) and database journals (Figure 2.4(c)). Several remarks can be observed from the reported results for these database research venues. Some key remarks are given as follows:


- There are distinctly 42 (non-distinctly 72) research scholars in the top publishers lists of the four core database technology conferences. There are distinctly 34 (non-distinctly 41) research scholars in the top publishers lists of the three main database journals. In combination, there are 63 distinct research scholars on the seven venues. These results show a clear overlap between the list of these top database research venues.
- Three research scholars appear on the top publishers list of *all* core database technology conferences. Namely, *Philip S. Yu*, *Nick Koudas* and *Yufei Tao*. In addition, *Philip S. Yu* appears on the top publishers lists of the VLDB journal and TKDE. *Yufei Tao* appears on the lists of the TODS and TKDE while *Nick Koudas* appears only on the list of TODS.
- Six research scholars appear on the top publishers list of *three* (out of four) core database technology conferences. Namely, *Divesh Srivastava*,




(a) VLDB



(b) SIGMOD




(c) ICDE




(d) EDBT

Figure 2.1: Top Publishers in Major Core Database Technology Conferences




(a) PODS




(b) ICDT


Figure 2.2: Top Publishers in Major Theoretical Database Conferences



(a) VLDB J.



(b) TODS



(c) TKDE

Figure 2.3: Top Publishers in Major Database Technology Journals

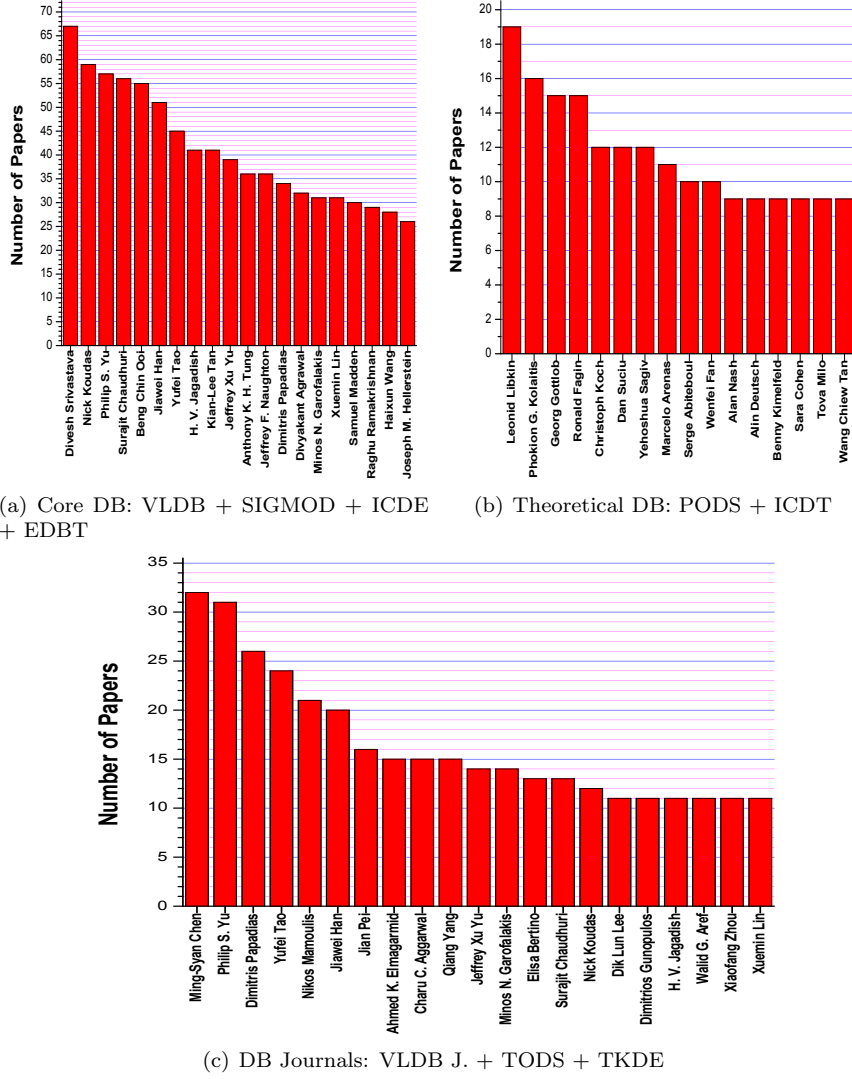


Figure 2.4: Aggregate Lists of Top Publishers for Database Research Venues



*Beng Chin Ooi, Surajit Chaudhuri, Jiawei Han, Jeffrey Xu Yu and H. V. Jagadish.* In addition, *Beng Chin Ooi* appears on the lists of the VLDB Journal and TKDE. *Jiawei Han* appears on the top list of TKDE. *Jeffrey Xu Yu* appears on the top list of the VLDB Journal. *Surajit Chaudhuri* and *H. V. Jagadish* appears on the top list of TODS.

- Eight research scholars appear on the top publishers list of *two* core database technology conferences. Namely, *Kian-Lee Tan, Anthony K. H. Tung, Haixun Wang, Dimitris Papadias, Jeffrey F. Naughton, Raghu Ramakrishnan, Divyakant Agrawal* and *Samuel Madden*. In addition, *Dimitris Papadias* appears on the top publishers lists of TODS and TKDE.
- There are 32 distinct research scholars in the top publishers list of the two theoretical database conferences (PODS and ICDT). Seven research scholars appear on the lists of both conferences. Namely, *Leonid Libkin, Marcelo Arenas, Phokion G. Kolaitis, Yehoshua Sagiv, Benny Kimelfeld, Christoph Koch* and *Ronald Fagin*.
- Seven research scholars have joint appearance on the top publishers list of at least one of the theoretical database conferences in addition to another appearance in at least one the top publishers list of a core database technology conference or a main database journal. Namely, *Victor Vianu* (ICDT, TODS), *Phokion G. Kolaitis* (PODS / ICDT, TODS), *Ronald Fagin* (PODS / ICDT, TODS), *Johannes Gehrke* (PODS, SIGMOD), *Wang Chiew Tan* (PODS, TODS), *Dan Suciu* (PODS, VLDB) and *Wenfei Fan* (PODS, VLDB).
- *Ming-Syan Chen* has the highest total number of publications in the major database journals in one year. In 2008, he has published 9 papers (5 papers in TKDE and 4 papers in VLDB Journal).
- *Philip S. Yu* has the highest total number of publications in the major database conferences in one year. In 2009, he has published 13 papers (6 papers in VLDB, 5 papers in ICDE and 2 papers in SIGMOD).
- *Divesh Srivastava* is the top publisher in the aggregate list of all core database technology conferences (Figure 2.4(a)). He published 67 papers in total with an average of about 7 papers per year. On the other side, he published only 5 papers in the main database journals. Therefore, he does not appear in the aggregate list of the main database journals (Figure 2.4(c)). Ten research scholars appear in both of the aggregate lists for top publishers on core database technology conferences and database journals. Namely, *Philip S. Yu* (with total of 88 papers), *Nick Koudas* (71 papers), *Jiawei Han* (71 papers), *Surajit Chaudhuri* (69 papers), *Yufei Tao* (69 papers), *H. V. Jagadish* (62 papers), *Dimitris Papadias* (60 papers), *Jeffrey Xu Yu* (53 papers), *Minos N. Garofalakis* (45 papers) and *Xuemin Lin* (42 papers).
- *Yannis Papakonstantinou* and *Dan Suciu* had at least one paper in each of the studied nine major database venues in the last decade.
- Table 1 shows the most important co-authorship relations between research scholars in the top lists of the database research venues. For

Deg.	Authors	# Pub.
2	Yufei Tao and Dimitris Papadias	34
2	Divesh Srivastava and Nick Koudas	33
2	Divyakant Agrawal and Amr El Abbadi	30
2	Vivek R. Narasayya and Surajit Chaudhuri	22
2	Beng Chin Ooi and Anthony K. H. Tung	16
2	Haixun Wang and Philip S. Yu	16
2	Xuemin Lin and Wei Wang	16
2	Xuemin Lin and Jeffrey Xu Yu	14
3	B. Gedik, P. S. Yu and K. Wu	9
3	D. Agrawal, A. El Abbadi and A. Metwally	7

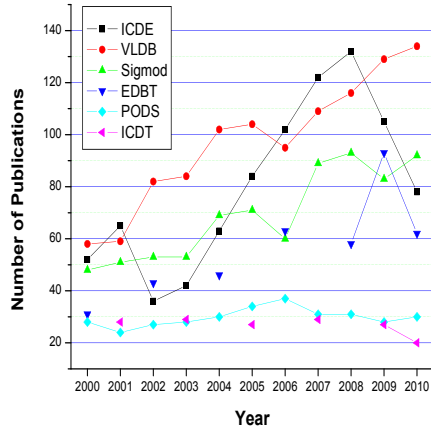
Table 2.1: Top Co-authorship Relationships

example, *Yufei Tao* and *Dimitris Papadias* have participated in the co-authorship of 34 regular paper in the different database research venues. The degree column (Deg.) indicates the number of the research scholars participating in the relationship.

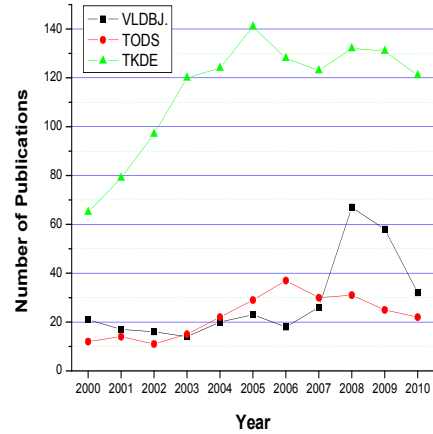
## 2.2 The Growth in number of Publications and Database Community Size

The topics of the database field is continuously growing. Therefore, there are more researchers who are entering the research community and more research papers are being published [4]. In our study, we determined the number of regular publications for all of our considered publication venues for the ten years period of 2001 - 2010. Moreover, we determined the number of unique authors for the publications of each venue as a measure of its contributing community size. Figure 2.5 presents an overview of the growth in the number of publications in the database research venues while Figure 2.6 presents an overview of the growth in the number of unique authors (participating community size). Combining the results of both figures show that the number of research publications and unique authors in core database technology conferences and database journals has on average nearly doubled in number. On the contrary for the theoretical database conference (PODS and ICDT), there was no clear increase either on the number of publications nor on the number of authors. They kept having an average of around 30 papers and 75 authors per conference over the whole decade.

In principle, the number of regular research publications for core database technology conferences cannot continue growing in proportion to the size of the community. Therefore, most of the conference have introduced other forms of publications such as: posters, short papers and demo papers in order to provide a chance for a wider part of the community to present their work and to continue attracting and focusing the researchers to participate in a small set of top conferences as there are always limits on the number of conferences that researchers can attend. For example, the 2002 edition of the ICDE conference first introduced the acceptance of demo papers, the 2003 edition introduced the acceptance of poster papers and the 2009 edition introduced the acceptance of 4

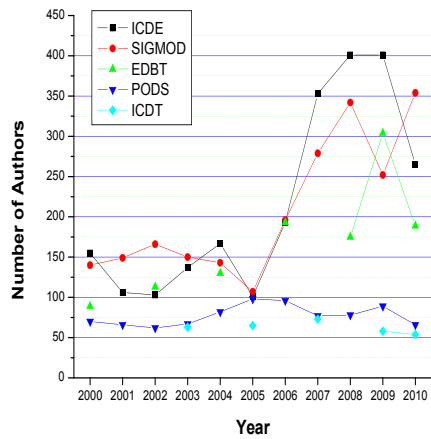


(a) Conferences

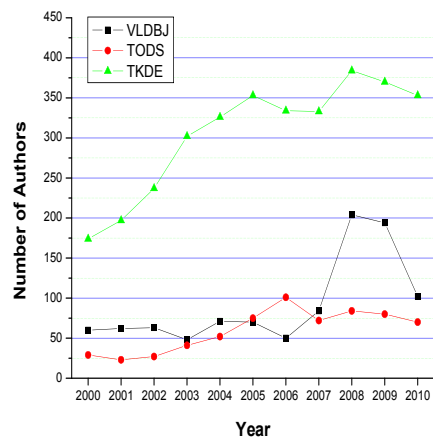


(b) Journals

Figure 2.5: Growth in Number of Publications



(a) Conferences



(b) Journals

Figure 2.6: Growth in Number of Authors

pages short papers. We believe that having more journal papers could be a good solution to absorb this continuous increase of research publications without the need to increase the number of conferences or to increase the number of accepted papers in the current conferences.

One of the main reasons behind the increase in the number of publications in the database community is the continuous introduction of new research challenges which is relevant to the scope of the community. For example, XML has started to be introduced as a hot research topic for the database research community in the early of the last decade. Moro et al [14] referenced a list of more than 100 publications in a survey paper that provides an overview of *some* of the work that have been done in different aspects for XML data management. Recently, the topic of large scale data management on cloud computing and parallel data processing (e.g. MapReduce) have been introduced and they attract a lot of interest from the database research community [1]. As a consequence, a new series of research conferences, the ACM Symposium on Cloud Computing, has been started in 2010 [10]. This series is co-sponsored by the ACM Special Interest Groups on Management of Data (ACM SIGMOD) and on Operating Systems (ACM SIGOPS). The conference will be held in conjunction with ACM SIGMOD and ACM SOSP Conferences in alternate years.

### 3 Conclusions

Research is a competitive endeavor. Research scholars usually have multiple goals to achieve and it is therefore reasonable that their impact must be judged by multiple criteria. We believe that ranking of research scholars based on the count of their publications in top-tier research venues can be an insightful indicator in a comprehensive assessment process. Other important factors such as: invitations to program committees of prestigious conferences, membership on editorial boards of high quality journals, grant funding and awards can be also good indicators for evaluating the impact of research scholars.

In this paper, we presented a detailed study for the publications of 6 major database conferences and 3 major database journals in the period between 2001 and 2010. The results of our study reveals the fact that the number of research publications per year and the community size has nearly doubled through the last decade. The results also show a considerable overlap between the top publishers lists of the core database technology conferences and the database journals. The results are also compatible with the fact that the researchers in the database community tend to prefer publishing their work in prestigious conferences rather than in major database journals. The average publication rate for top publishers in conference venues highly exceed their average publication rate in the major database journals. In principle, we believe that conference publications will remain as an attractive way to gain a quick publicity for new research findings. However, the number of conferences or the number of accepted publications per conference can not continue increasing as this will limit the value of these venues gradually. Therefore, we believe that journal papers will remain as the best way to document and archive significant pieces of research which can not fit within the 12-page limit of conferences. The community should continue pushing towards achieving the switch to the culture of highly evaluating the journal papers over the conference papers [2]. One of the valuable trials in this

direction is the introduction of the The Proceedings of the VLDB Endowment (PVLDB)<sup>1</sup> which aims of providing *journal-like* experience to authors of the VLDB submissions.

## Bibliography

- [1] Divyakant Agrawal, Sudipto Das, and Amr El Abbadi. Big Data and Cloud Computing: New Wine or just New Bottles? *PVLDB*, 3(2):1647–1648, 2010.
- [2] Anastassia Ailamaki, Laura M. Haas, H. V. Jagadish, David Maier, M. Tamer Özsu, and Marianne Winslett. Time for Our Field to Grow Up. *PVLDB*, 3(2):1658, 2010.
- [3] Mayur Amin and Michael Mabe. Impact Factors: Use and Abuse. *Perspectives in Publishing*, pages 1–6, 2000.
- [4] Philip A. Bernstein, David J. DeWitt, Andreas Heuer, Zachary G. Ives, Christian S. Jensen, Holger Meyer, M. Tamer Özsu, Richard T. Snodgrass, Kyu-Young Whang, and Jennifer Widom. Database Publication Practices. In *VLDB*, pages 1241–1246, 2005.
- [5] Jilin Chen and Joseph A. Konstan. Conference paper selectivity and impact. *Commun. ACM*, 53(6):79–83, 2010.
- [6] Leo Egghe. Theory and practise of the g-index. *Scientometrics*, 69(1):131–152, 2006.
- [7] Daniela Florescu. Database Research for the Current Millennium. In *ICDE*, page 866, 2004.
- [8] Lance Fortnow. Viewpoint - Time for computer science to grow up. *Commun. ACM*, 52(8):33–35, 2009.
- [9] Alon Y. Halevy. Answering queries using views: A survey. *VLDB J.*, 10(4):270–294, 2001.
- [10] Joseph M. Hellerstein, Surajit Chaudhuri, and Mendel Rosenblum, editors. *Proceedings of the 1st ACM Symposium on Cloud Computing, SoCC 2010, Indianapolis, Indiana, USA, June 10-11, 2010*. ACM, 2010.
- [11] Jorge Hirsch. An index to quantify an individual’s scientific research output. *Proceedings of the National Academy of Sciences*, 102(46):16569–16572, 2005.
- [12] Mikalai Krapivin, Maurizio Marchese, and Fabio Casati. Exploring and Understanding citation-Based Scientific Metrics. *Advances in Complex Systems*, 13(1):59–81, 2010.
- [13] M. H. MacRoberts and Barbara R. MacRoberts. Problems of citation analysis. *SCIENTOMETRICS*, 36(3):435–444, 1996.

---

<sup>1</sup><http://www.vldb.org/pvldb/>

- [14] Mirella M. Moro, Vanessa P. Braganholo, Carina F. Dorneles, Denio Duarte, Renata de Matos Galante, and Ronaldo dos Santos Mello. XML: some papers in a haystack. *SIGMOD Record*, 38(2):29–34, 2009.
- [15] Erhard Rahm and Philip A. Bernstein. A survey of approaches to automatic schema matching. *VLDB J.*, 10(4):334–350, 2001.
- [16] Erhard Rahm and Andreas Thor. Citation analysis of database publications. *SIGMOD Record*, 34(4):48–53, 2005.
- [17] Xiaolin Shi, Belle L. Tseng, and Lada A. Adamic. Information Diffusion in Computer Science Citation Networks. In *ICWSM*, 2009.
- [18] Richard T. Snodgrass. Journal Relevance. *SIGMOD Record*, 32(3):11–15, 2003.
- [19] Su Yan and Dongwon Lee. Toward alternative measures for ranking venues: a case of database research community. In *JCDL*, pages 235–244, 2007.