Technical University of Crete
Electronic and Computer Engineering Department

How good is your work?
A New Conference Classification Approach for
Judging the Quality of a Publication

by

Orestis - Stavros Loizides

Diploma Thesis

Advisor: Assistant Professor Polychronis Koutsakis

December 2012

Acknowledgements

Working on this thesis was a pleasant but arduous experience which taught me values and supplied me with qualifications for my career and my personal life as well. The combined difficulty of the significant number of papers on the subject that I had to read, the amount of research work I had to conduct and finally having to write down the results of my thesis led me to work often at hours where NBA games are played on the other side of the Atlantic. However, in this route I was not alone. There were people who stood by me, either making a path for me or escorting me there.

First, I would like to thank from the bottom of my heart my advisor, Professor Polychronis Koutsakis. With your continuous and effective guidance you helped me a lot to achieve my goals. In the years of our acquaintance and especially over the last year of working together, your support, insistence and patience were an inspiration to me beyond any personal expectations. The way you work and treat people around you is a model for me and makes me feel that I acquired, apart from a fruitful cooperation, a friend as well.

I would also like to thank all my friends in Greece and America who stood by me and helped me with every question, understanding and support I asked for. More personally, I would like to thank Constantinos Loukakes, Manolis Yiannopoulos,Thanasis Malisovas and Katerina Tzedakes who experienced my stress and anxiety every day over the last year. You were the source of strength and self- confidence that I needed when I felt "lost".

Finally, I thank my parents and my siblings for their unconditional love and support which they give me every day. Even when you ran out of patience with me, you kept believing in me and rewarding my efforts.

Abstract

The lack of a clear method to judge a researcher's contribution has recently led Nelakuditi et al. to propose a new metric, called Peer Reputation (PR) metric. PR ties the selectivity of a publication venue with the reputation of the first author's institution. The authors who proposed PR compute it for a number of networking research publication venues and argue that PR is a better indicator of selectivity than a venue's Acceptance Ratio (AR). We agree that PR is an idea towards the right direction and that it offers substantial information that is missing from AR. Still, we argue in this thesis that PR is not adequate by itself in giving a solid evaluation of a researcher's contribution. In our study, we discuss and evaluate quantitatively the points on which PR does not sufficiently serve its purpose and then we propose a new, finer-grain conference classification approach which partially incorporates PR, but aims to alleviate its shortcomings. To achieve this, we have gathered data from close to 3000 papers from 12 top-tier conferences (54 venues in total) belonging to different research fields (networking, informatics, electronics), between 2006-2011. In our work we use three different rankings of doctoral programs in USA and three world university rankings, to study how the different rankings influence our results.

Table of Contents

Introduction

The lack of a clear method to judge a researcher's contribution has recently [1] led to the proposal of a new metric, called Peer Reputation (PR) metric. PR ties the selectivity of a publication venue with the reputation of the first author's institution. In [1], the authors compute PR for a number of networking research publication venues and argue that PR is a better indicator of selectivity than a venue's Acceptance Ratio (AR). We agree that PR is an idea towards the right direction and that it offers substantial information that is missing from AR. Still, we argue in this thesis that PR is not adequate by itself in giving a solid evaluation of a researcher's contribution. In our study, we discuss and evaluate quantitatively the points on which PR does not sufficiently serve its purpose and then we propose a new, finer-grain conference classification approach which partially incorporates PR, but aims to alleviate its shortcomings. To achieve this, we have gathered data from close to 3000 papers from 12 top-tier conferences (54 venues in total) belonging to different research fields (networking, informatics, electronics), between 2006-2011. In our work we use three different rankings of doctoral programs in USA and three world university rankings, to study how the different rankings influence our results.

We propose the use of 4 metrics for quantifying the level of a conference. Two of these metrics are associated with PR, one is associated with the quality of the paper as denoted by the paper's citations during its lifetime since publication and the last metric is associated with the recognition of the authors of the papers presented in each conference. We use these 4 metrics to produce a vector associated with each conference, and then, via the k-means clustering algorithm, we classify the conferences into clusters of excellence.

We believe that our approach can bring a fresh look into the current conference evaluation system.

1. Evaluation of PR

1.1 AR versus PR

The authors in [1] argue that a new metric for the qualitative assessment of a researcher’s contribution in a specific discipline is needed. To make their point, they refer to the shortcomings of the use of conference AR (authors self-filtering can lead to highly selective conferences having similar ARs with less selective ones) and of the use of the number of citations that a publication receives (citations of conferences are not well-tracked, new publications in conferences and journals take time to attract citations). The authors in [1] argue that the selectivity of a publication venue is a function of the reputations of the authors' affiliating institutions, and for simplicity they represent each paper by the affiliation of its first author, which is mapped to a rank. Then, the PR of the conference is a function of this rank set, and defined as

$\mathrm{PR}(P)=\mathrm{gist}(\{\mathrm{rank}(\mathrm{inst}(\mathrm{first}(p))): \forall p \in P\})$,

where P is a publication venue, p is paper, first() gives the first author, inst() maps an author to the affiliation, and rank() is the given function for mapping an institute to a rank. Their PR metric, which is proposed for making “snap judgments” on a publication quality, conveys the selectivity of a conference with a tuple, say <1/3, 20>, indicating that 1/3 of the papers at that conference are from the top 20 universities. PR is evaluated over 18 venues, in [1], 16 of which are conferences and 2 are journals. Our work in this thesis focuses only on conferences.

It is also stated in [1] that although PR is not a perfect metric to assess the quality of a publication, it provides a coarse-grain measure of the selectivity of a conference or a journal, and can potentially be more helpful than AR. We agree with the latter comment, not only because of the aforementioned shortcomings associated with AR, but also because AR does not give the extent of competition among peers for a conference publication, since there is no indication of the quality of rejected or accepted papers (this point is not made in [1]). Hence, PR can indeed be very helpful in highlighting the work of researchers from lower ranked universities who are publishing in top venues. However, we believe that, despite the improvement it offers over AR, PR also has some important shortcomings as a metric, and we conducted an extensive study to ascertain whether our points are valid.

1.2 Weaknesses of PR

As mentioned above, we intuitively found PR and the results in [1] to have some practical shortcomings. There is also a "philosophical" issue regarding the use of PR that we wish to briefly mention before going into detail about the practical issues:

PR adopts the currently prevalent bias in academia, that an institution "makes" a researcher, i.e., a researcher working in a top-ranked university is expected to be better than a colleague affiliated with a lower-ranked university. Of course, many top-class researchers work in lower ranked universities. The number of IEEE/ACM Fellows in universities ranked below the top-20 is enough to question the "aloofness" of this bias. However, we will not delve further into this point in this thesis, as it is not a problem of PR but of the prevalent conception, as already mentioned.

The practical problems that we found concerning PR are the following:

1. The evaluation of PR is based, in [1], only on networking conferences, and is made only over two years (2008, 2009). This is not only limiting in itself, but also because the

universities' rankings used in [1] are based on the evaluation of *Computer Science* graduate schools, hence the rankings reflect the quality of a graduate program on a much larger set of disciplines than the single discipline (networking) that [1] focuses on. Hence, these rankings may not even be fully representative in terms of the networking field (e.g., excellence in another field may give one university the edge in Computer Science rankings over other universities which might have stronger networking graduate programs).

2. PR disregards the importance of the location of a conference. Depending on the continent where a conference takes place, submissions from far-away countries might be discouraged as travel could be too time- and energy-consuming, or too expensive.

3. PR is based only on the ranking of the university that the first author is affiliated with. Given the fact that many papers stem from the collaboration of authors from different universities, or from different departments of the same university, which are differently ranked in their respective fields, the choice of using just the first author seems to be an oversimplification.

4. A limitation that the authors themselves mention, but do not consider as important, is that they base PR solely on the ranking of US universities. The authors explain, in [1], that in their view this is not a serious limitation for popular networking conferences as these venues receive a high fraction of papers from USA universities. Our study, which includes rankings of world universities, shows that PR results can significantly change when universities outside the USA are taken into account.

All of the above limitations are discussed in Section 1.3.

1.3 An Extensive Evaluation of PR

In order to alleviate the first weakness of PR and get a more representative set of results from more fields of Computer Science, we studied 12 top-tied conferences between 2006 and 2011. In total, we studied 54 venues, 9 of which focus on Informatics (I), 23 on Electronics (E) and 22 on Networking (N). Four of the networking venues are the networking symposiums of flagship IEEE conferences (ICC, WCNC)[1] which however have a much higher AR than the rest of the conferences under study (note that the ARs for ICC and WCNC refer to the whole conference, not just the networking symposiums). A total of 2783 papers were used in our study.

In this first part of our thesis we present, in Table 1, the data of 30 conferences of Informatics, Electronics and Networking between 2008 and 2011. The AR values for the conferences used in our study have been taken from [6-10] and from numerous other sources on the web, in order to double-check the AR values reported. In the very few cases that our results differ from those in [1], this is caused by the fact that we round up to the higher integer while the authors in [1] round up to the lower (e.g., for the 31 accepted papers in Mobihoc 2009, our 1/3 PR results use the 11th highest ranked authors' affiliation, while [1] used the 10th).

#	Conference	#papers	AR	PR: 1/4	PR: 1/3	Field
1	MOBIHOC 2008, HONG KONG	44	14.7	28	44	N
2	MOBIHOC 2009, NEW ORLEANS	31	17.7	44	44	N
3	MOBIHOC 2010, CHICAGO	26	16.6	28	35	N
4	MOBIHOC 2011, PARIS	25	19.7	13	20	N

[1] For ICC, we used the Wireless Networking Symposium and the Next Generation Networking and Internet Symposium, for both 2010 and 2011. For WCNC, we used the Networks Track for both years.

5	IMC 2008, ATHENS	31	17.3	10	11	N
6	IMC 2009, CHICAGO	41	22.4	13	20	N
7	IMC 2010, MELBOURNE	47	22.3	14	20	N
8	IMC 2011, BERLIN	42	19.1	14	35	N
9	SENSYS 2008, RALEIGH	25	16.3	5	5	N
10	SENSYS 2009, BERKELEY	21	17.6	1	17	N
11	SENSYS 2010, ZURICH	25	17.6	14	20	N
12	SENSYS 2011, SEATTLE	24	19.5	20	28	N
13	ICC 2010, CAPE TOWN	138	39.5	122	122	N
14	ICC 2011, KYOTO	165	38.5	122	122	N
15	WCNC 2010, SYDNEY	180	37.2	122	122	N
16	WCNC 2011, CANCUN	142	48.1	122	122	N
17	ISCA 2010, SAINT-MALO, FRANCE	44	18	10	11	E
18	ISCA 2011, SAN JOSE	40	19	8	11	E
19	MICRO 2010, ATLANTA	45	18.1	8	13	E
20	MICRO 2011, PORTO ALEGRE, BRAZIL	44	21	11	13	E
21	ASPLOS 2010, PITTSBURGH	32	17.7	7	8	E
22	ASPLOS 2011, NEWPORT BEACH, USA	32	21.1	8	11	E
23	FCCM 2010, CHARLOTTE	24	18.2	39	122	E
24	FCCM 2011, SALT LAKE CITY	21	17.7	53	99	E
25	ICDE 2010, LONG BEACH, USA	69	12.5	20	39	I
26	ICDE 2011, HANNOVER	98	19.8	122	122	I
27	VLDB 2008, AUCKLAND, NEW ZEALAND	101	17.1	14	39	I
28	VLDB 2009, LYON, FRANCE	108	17	28	122	I
29	VLDB 2010, SINGAPORE	120	20.1	20	58	I
30	VLDB 2011, SEATTLE	100	18.1	20	53	I

Table 1. Conference Data

The second limitation of PR (not taking into account the conference location) is evident by the results for the SenSys conference, between the years 2008 and 2010. The results reveal that when the conference took place in the USA, in 2008 and 2009, the 1/4 and 1/3 PR results showed a significantly larger selectivity in comparison to the results for 2010, when the conference took place in Europe. It should also be emphasized that the AR fluctuated very slightly for all three of these venues of SenSys; this agrees with the results in [1] and the rest of our results and indicates that there is no correlation between PR and AR, other than the obvious one that a much higher AR usually leads to lower selectivity. The results for SenSys 2011, which again took place in the USA, are slightly worse, but the AR was quite higher, therefore a direct comparison cannot be made. Another indication of the influence that the location of a venue can have on PR results is given by comparing the results for the four VLDB venues, which took place in four different continents. The VLDB 2009 conference, which took place in France and had the lowest AR among the four venues (therefore based on AR it would seem to be the *most* selective), appears to have been the *least* selective one based on the PR metric. A possible explanation for this result is the attractiveness of France as a location. This may have led to the increased number of submissions (highest number among the four VLDB editions studied), thus leading to the lowest AR but also to the acceptance of many papers written by non-US-based authors, hence the worse PR values (this result is therefore also tied with the fourth weakness of PR, which we discuss below).

Still, location is not always a critical factor. Counter-examples include the IMC and ISCA conference venues, where the change in location led to rather negligible changes in the PR results; the slightly better results for IMC 2008, in comparison to IMC 2009-2011, are connected to the much lower AR and number of accepted papers for that venue.

The above results seem to indicate that location is a factor, but its influence can be mitigated by other facts (e.g., a paper may be co-authored by researchers working in different locations, hence it is easier for one of them to present it). The efficient incorporation of location into a metric is a quite complex task which we intend to tackle in future work.

The co-authorship of papers leads us to the third weakness of PR (taking into account the affiliation of just the first author). To give an example of how this weakness, combined with the inherent "aloofness" of the metric, can lead to incorrect judgments regarding the quality of a venue, consider the following case: let's assume that a conference contains only papers sole-authored by IEEE/ACM fellows from universities that are ranked between 50-70. Suppose also that another conference contains papers where the first authors are post-docs from top-20 ranked universities. This second conference will be ranked by PR as a top-class venue, whereas the first conference will be considered a rather mediocre venue. This is obviously a wrong conclusion, but this is what someone using PR will conclude by making a "snap judgement". The above example takes the specific shortcoming of PR to the extreme, of course, but the shortcoming is clear. Therefore, a finer grain evaluation than that which is offered by PR is needed, in our view.

To alleviate this problem, we used the mean ranking of the universities that all authors of each paper are affiliated with. Figures 1 and 2 show, once again, that the results vary depending on the conference. It should be noted that the conference numbers on the xaxis follows that of Table 1, is completely ad-hoc and plays no role in the results other than allowing a visual comparison between the ranking of the first author and the mean ranking for the same venue.

As shown in the Figures, in quite a few venues the 1/4 and 1/3 PR first author and average results are almost identical. This is explained by the fact that a large number of papers were written from researchers who had the same affiliation. However, the inclusion of all authors can also lead to significant changes in the PR values and hence to very different "snap judgments", as shown in the Figures for some of the other conferences. This was true for 23% of the venues in terms of the 1/4 PR results, and for 40% of the venues in terms of the 1/3 PR results. In all but one of the venues these changes were related to a *decrease* in the 1/4 and 1/3 PR values (i.e., an increase in selectivity). This indicates that at least for these conferences, the first authors were on average affiliated with lower-ranked universities than their co-authors.

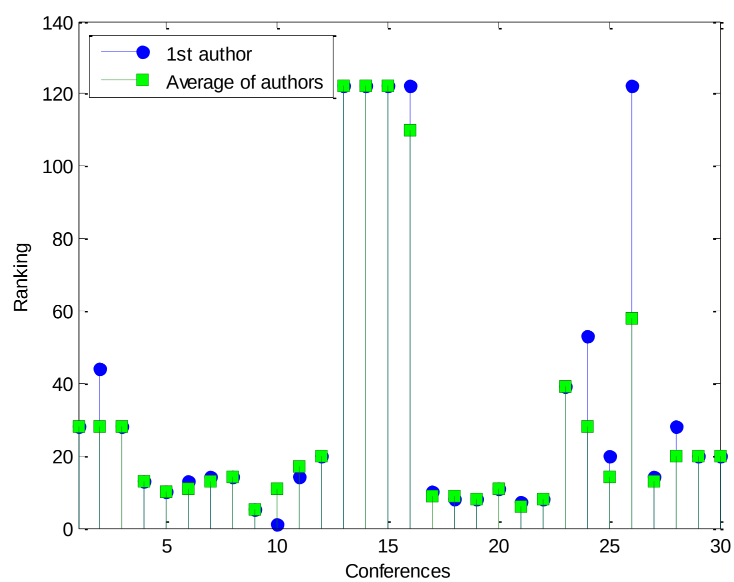

Figure 1. 1/4 PR results for the first authors and for the average over all authors

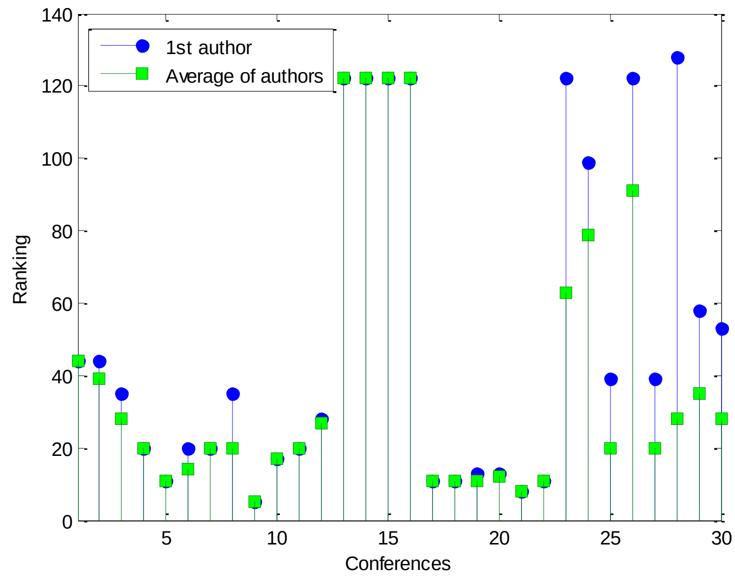

Figure 2. 1/3 PR results for the first authors and for the average over all authors

All of the above results have been produced, similarly to [1], with the US News and World Report rankings of graduate programs in Computer Science [2]. As it is well-known, university rankings are the subject of much controversy, and the rankings methods used are often questioned by universities and individual scholars. In order to alleviate the fourth weakness of PR and get a more clear view of how different rankings influence our results, we have used four additional rankings:

a) the latest (2010) National Research Council (NRC) rankings for the top 127 Computer Sciences graduate programs in the USA [3]

b) the NRC rankings for the top 136 Electrical and Computer Engineering graduate programs in the USA [3].

c) the specialized Academic Ranking of World Universities (ARWU) in Engineering/Technology and Computer Sciences - 2010 [4].

d) the specialized latest ARWU ranking in Computer Science for 2012 [5].

Over all the venues we studied, we found that 55% of the papers were written by authors who did not have an affiliation with a university in the USA (therefore, they are completely disregarded by PR, if they are first authors; otherwise they would be disregarded by PR anyway). This indicates the importance of taking world university rankings into account to ensure the accuracy and fairness of the results, and is further supported by the results presented in Figures 3-6, on the PR metric. These reveal that, despite the fact that the results show more or less similar *behavior* for all rankings, the PR values differ significantly depending on the ranking, hence very different "snap judgments" will be made by using PR when including non-US universities that in the case of ignoring them. The results also reveal, similarly to Figures 1-2, that for over 30% of the results, for each ranking, the use of the average PR leads to significantly different results than those when only the first author's affiliation is used.

We need to clarify that, similarly to [1], we treat all research institutions that are not included in a specific ranking as ranked lower than the lowest ranking (e.g, for the NRC Computer Science rankings, they get a rank equal to 128). This is the reason that the ICC and WCNC conferences (conferences 13-16) have an almost constant PR value in all the figures (i.e., the 1/4 and 1/3 PR values correspond to research institutions that are not included in the rankings).

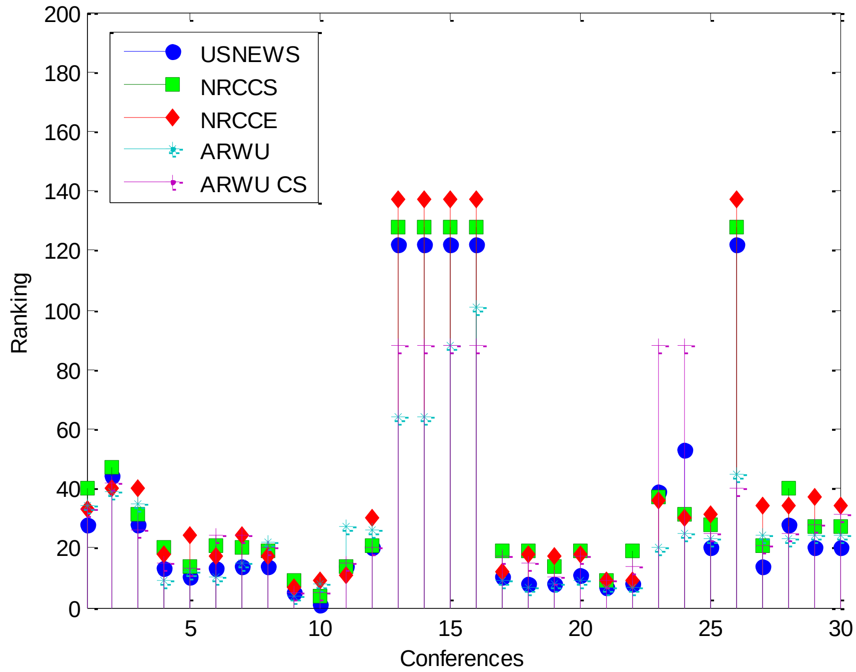

Figure 3. 1/4 PR results for the first author, with the use of 5 different rankings

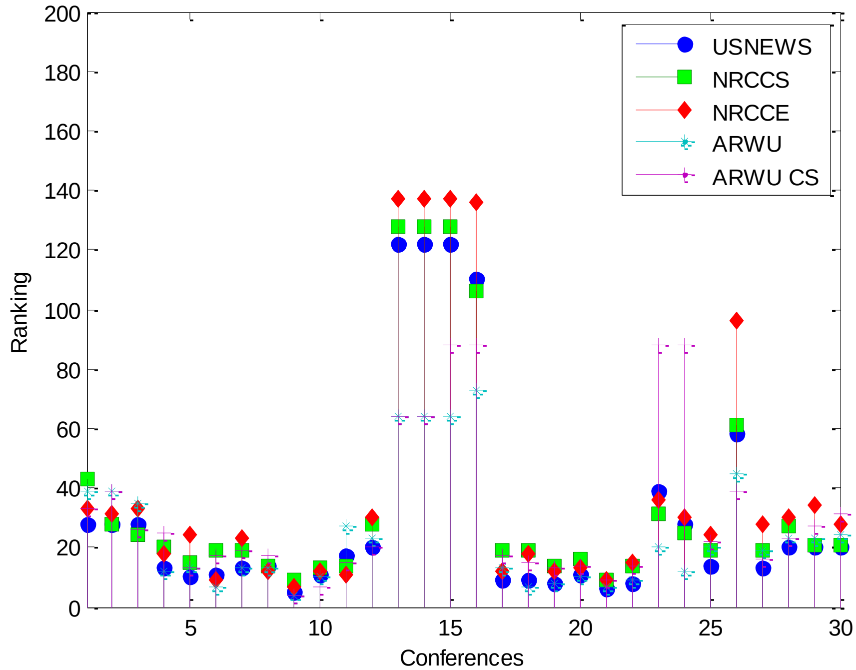

Figure 4. Average 1/4 PR results, with the use of 5 different rankings

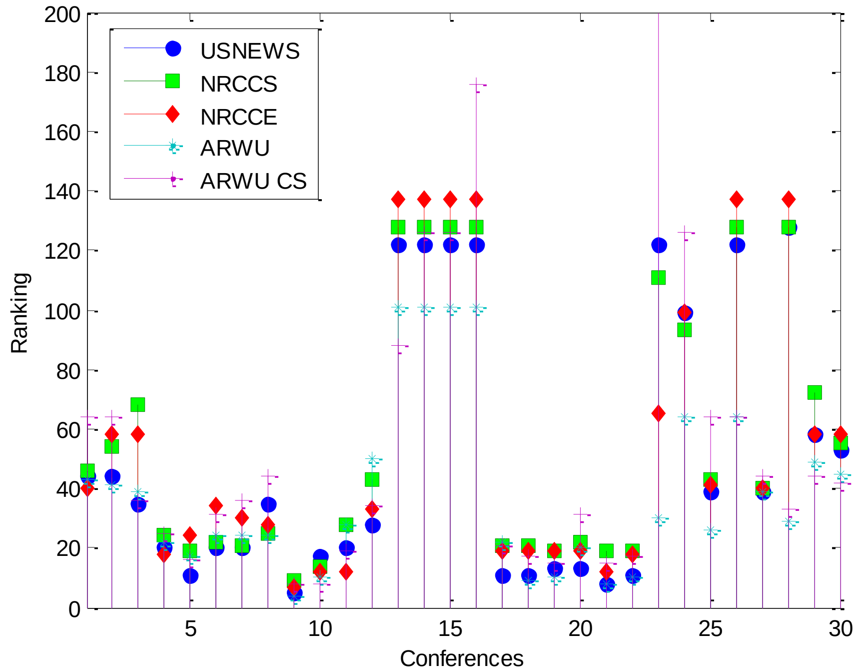

Figure 5. 1/3 PR results for the first author, with the use of 5 different rankings

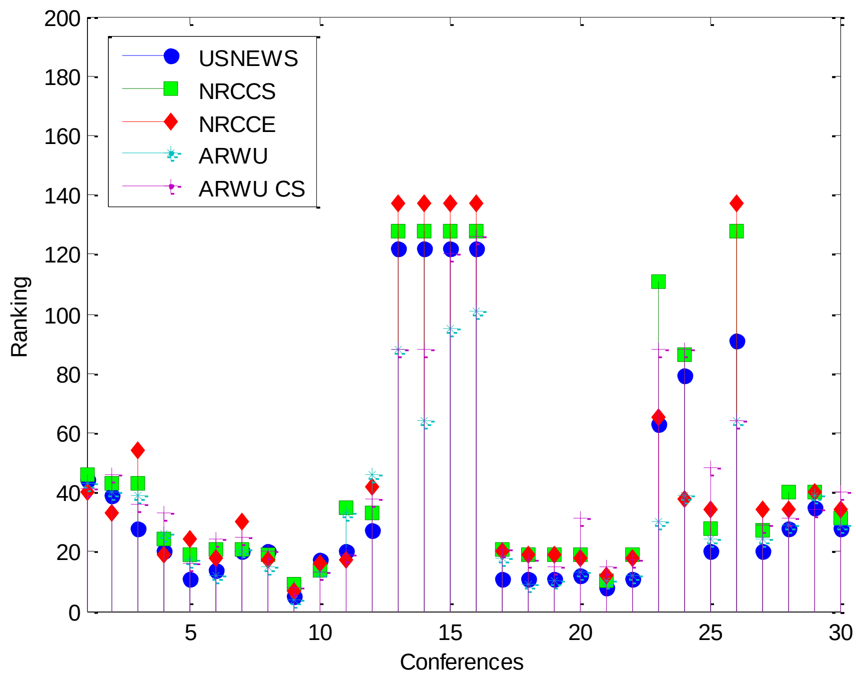

Figure 6. Average 1/3 PR results, with the use of 5 different rankings

Tables 2 and 3 show the average 1/4 and 1/3 PR results derived from each one of Figures 3-6, for all 30 venues. Table 2 includes the average results for just the first author, while Table 3 includes the average results for the mean PR for all authors. We present the results for all venues, and for the N, E, and I venues separately. The results presented in the Tables show once again the differences when using different rankings and that, on average, first authors are affiliated with lower ranked institutions than the rest of the authors in the accepted papers. These results also reveal that, when using any of the specialized rankings (i.e., with the exception of the more general ARWU one), the inclusion of all authors has minor effects on the results for networking conferences, larger effects on the results for electronics conferences and major effects on the results for informatics conferences.

Finally, another way of judging the influence of including all authors in the PR results is the relative ordering of conferences based on their PR values for the first author and for all authors, respectively. Our results once again confirm the significant effect of taking all authors into consideration, for some of the conferences: the SENSYS2009 conference, for example, which is the best among all 30 conferences in terms of its 1/4 PR value for the first author (US News ranking), falls to the 8th place when all authors are considered; another example is the VLDB09 conference, which is tied with 5 other conferences in the last place in terms of its 1/3 PR value (equal to 122) for the first author, but climbs to the 18th place (with a strong 1/3 PR value equal to 28) when all authors are considered.

	USNEWS		NRCCS		NRCCE		ARWU		ARWU_CS	
	1/4PR	**1/3PR**	**1/4PR**	**1/3PR**	**1/4PR**	**1/3PR**	**1/4PR**	**1/3PR**	**1/4PR**	**1/3PR**
ALL VENUES	22.1	38.9	26.8	44.4	27.7	41.5	18.9	28.2	25.6	43.1
N	17	24.9	21.7	31.1	22.5	29.5	20.1	27.2	20.2	32.1
E	18	36	20.9	40.6	18.6	33.8	11.5	21.6	32.3	55.4
I	40.8	77.2	48.8	82.2	54.6	82.6	27.8	41.4	27.8	49.8

TABLE 2. Averages for the 1/4 and 1/3 PR results for the first author, over all the venues.

	USNEWS		NRCCS		NRCCE		ARWU		ARWU_CS	
	1/4PR	**1/3PR**	**1/4PR**	**1/3PR**	**1/4PR**	**1/3PR**	**1/4PR**	**1/3PR**	**1/4PR**	**1/3PR**
ALL VENUES	17.6	26.6	21.6	35.8	24	32.2	18	25.2	24.8	32.7
N	16.5	22.1	20.5	27.3	20.3	26.4	19.6	25.8	19.6	26.6
E	14.8	25.8	18.4	38	18.1	26.1	10.8	17.6	32.3	36.5
I	25	38.8	29.4	52.6	42.4	52.8	26	36	25.4	41.2

TABLE 3. Averages for the mean 1/4 and 1/3 PR results for all authors, over all the venues.

2. A New Conference Classification Approach for Judging the Quality of a Publication

The results derived and presented in Section 1 were very useful not only in substantiating the shortcomings of PR, but also in helping us to identify the individual elements that need to be included in the evaluation of the quality of a conference publication. These three elements, constituting the three "dimensions" of our approach, are the Papers' Impact, the Authors' Recognition and the Conference Data (Figure 7).

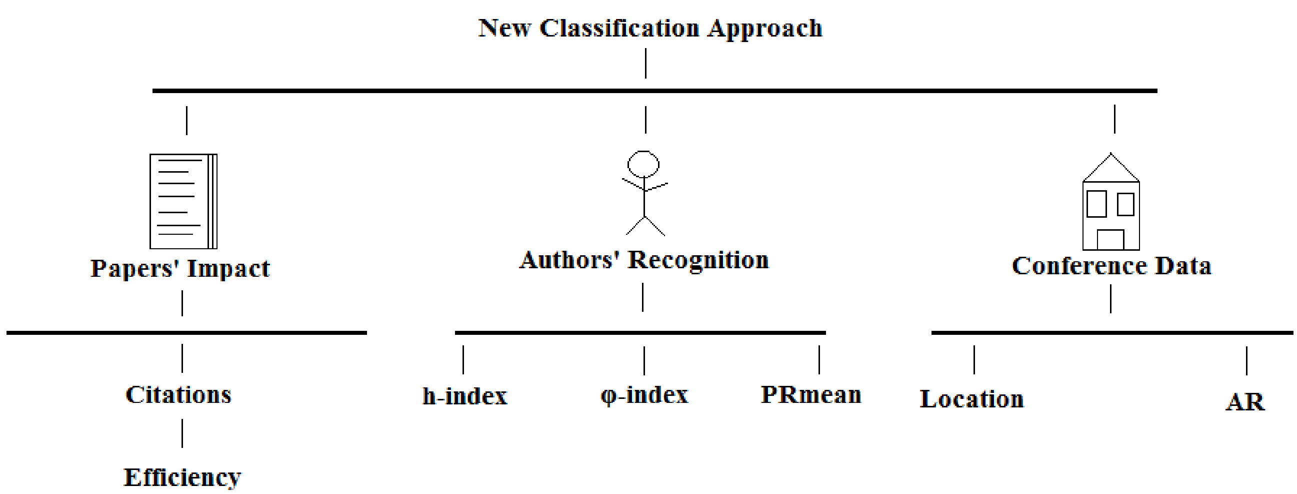

Figure 7. The three "dimensions" of our classification approach

2.1 Papers' Impact

Instead of focusing just on the affiliation of the authors of papers presented in top-tier conferences, as the PR metric does, we had the idea of studying the impact of accepted papers on the scientific community. This impact can be at least partially measured by the citations received by the papers presented in the conference and published in the conference proceedings.

2.1.1 Citations

We notice, from the data we gathered and present in Table 4, that in almost every conference for which we computed the mean number of citations of its papers, the number of citations increases as years go by. It should be noted that we are taking into account conferences from 2006 to 2009 as the number of citations that conferences get in the first couple of years after their publication is usually poor. We present results for a significant subset of the conferences studied in our work; this subset contains Networking and Electronics conferences.

Year	**MobiHoc**	**IMC**	**SenSys**	**ISCA**	**MICRO**	**ASPLOS**	**FCCM**
2006	94	85	154	87	69	84	28
2007	68	80	70	79	52	n/a	17
2008	46	34	82	60	25	50	12
2009	34	30	52	46	26	52	8

TABLE 4. Mean number of citations for Networking and Electronics conferences.

The aforementioned results are intuitively expected, as the longer a paper is published the more citations it is probably to get, until it reaches a point, usually after a number of years, when its contribution no longer coincides with the state-of-the-art and the number of citations will start decreasing. For the purposes of our work, we need a value which will be associated with the number of citations and which can be used in studying the quality of a publication venue. The *Efficiency* of a paper is defined in [12] as

Efficiency(p) = #citation / (2011 – publication_year)

where p is a paper, #citation is the total amount of citations a paper has received and publication_year is the current year we study.

We use 2011 as the default year in the denominator, in order to allow for at least a 2-year interval during which the papers can start accumulating citations. In this way, we have a means of quantifying knowledge that has been created, shared and distributed among authors. In Table 5 we present the mean Efficiency values of the papers presented in each conference.

Year	**MobiHoc**	**IMC**	**SenSys**	**ISCA**	**MICRO**	**ASPLOS**	**FCCM**
2006	19	17	31	17	14	17	6
2007	17	20	18	20	13	n/a	4
2008	15	11	27	20	8	17	4
2009	17	15	26	23	13	26	4

TABLE 5. Mean Efficiency for Networking and Electronics conferences.

By studying the results of Table 5 we can draw useful conclusions about which year of a conference had the greatest resonance with the scientific community. Furthermore, if we compare the Efficiency results with those of the average PR values, we can see that they constitute two entirely different elements. For instance, SenSys 2006 shows the highest mean Efficiency results when compared with the subsequent venues of the same conference (and when compared with all other conferences); however, this venue had the worst average 1/4 and 1/3 PR values (20, 41 respectively) of all other SenSys venues (2007: (8, 17), 2008: (8, 8), 2009: (9, 20)). Equally interesting is the fact that in 2006

SenSys had the lowest selectivity (19.7% AR) compared to the next three years (2007: 16.8% AR, 2008: 16.3% AR, 2009: 17.6% AR). Given that the 1/3 PR value of 41 for SenSys 2006 would indicate, according to [1], that this was a venue of just good quality, the above results show that the impact of a paper is not necessarily relevant to an author's affiliation, nor to the Acceptance Ratio of a conference.

2.2 Authors' Recognition

We can derive significant conclusions about an author's contribution to his field (and therefore the potential future impact of a paper he presents for a conference) by studying his/her h-index [15] and φ-index [11]. The first gives a strong indication of the impact of the author's past work, while the second indicates the author's partnership ability, which is usually accompanied with higher citation visibility. More specifically, the partnership ability φ-index takes into account the number of co-authors and the times each of them acted as co-author. An author is said to have a partnership ability φ, if with φ of his/her n co-authors he/she had at least φ joint papers, and with the rest (n-φ) co-authors he/she had no more than φ joint papers. Still, gathering all the required information for computing the φ-index of every author of the 2783 papers studied was out of the scope of this work. Therefore, the inclusion of the φ-index in our classification approach will be part of our future work.

We will discuss in Section 2.2.1 the importance of the h-index for our conference classification, but before doing so we present in Table 6 the **average** 1/4 PR and 1/3 PR values, based on the LEIDEN ranking, for all 54 venues used in our study. Table 6 also contains the AR values for all venues, with 3 exceptions for which we were unable to find the AR of the conference.

#	Conference	#papers	AR	PRmean: 1/4	PRmean: 1/3	Field
1	MOBIHOC 2006, ITALY	31	9.8	21	40	N
2	MOBIHOC 2007, CANADA	27	18.5	20	34	N
3	MOBIHOC 2008, HONG KONG	44	14.7	47	58	N
4	MOBIHOC 2009, NEW ORLEANS	31	17.7	62	91	N
5	MOBIHOC 2010, CHICAGO	26	16.6	50	58	N
6	MOBIHOC 2011, PARIS	25	19.7	30	53	N
7	IMC 2006, BRAZIL	34	21.9	27	34	N
8	IMC 2007, SAN DIEGO	39	24.4	26	35	N
9	IMC 2008, ATHENS	31	17.3	40	42	N
10	IMC 2009, CHICAGO	41	22.4	32	51	N
11	IMC 2010, MELBOURNE	47	22.3	20	30	N
12	IMC 2011, BERLIN	42	19.1	15	18	N
13	SENSYS 2006, COLORADO	24	19.7	20	41	N
14	SENSYS 2007, SYDNEY	25	16.8	8	17	N
15	SENSYS 2008, RALEIGH	25	16.3	8	8	N
16	SENSYS 2009, BERKELEY	21	17.6	9	20	N
17	SENSYS 2010, ZURICH	25	17.6	22	30	N
18	SENSYS 2011, SEATTLE	24	19.5	57	57	N
19	ICC 2010, CAPE TOWN	138	39.5	124	182	N
20	ICC 2011, KYOTO	165	38.5	151	192	N
21	WCNC 2010, SYDNEY	180	37.2	182	213	N
22	WCNC 2011, CANCUN	142	48.1	166	199	N
23	ISCA 2006, BOSTON	31	13	25	32	E
24	ISCA 2007, SAN DIEGO	46	23	22	32	E
25	ISCA 2008, CHINA	37	14	16	30	E

26	ISCA 2009, TEXAS	43	20	15	22	E
27	ISCA 2010, SAINT-MALO, FRANCE	44	18	15	30	E
28	ISCA 2011, SAN JOSE	40	19	24	32	E
29	MICRO 2010, FLORIDA	42	24	15	15	E
30	MICRO 2010, ILLINOIS	35	21.1	26	30	E
31	MICRO 2010, ITALY	40	19	22	30	E
32	MICRO 2010, NEW YORK	52	24.9	26	30	E
33	MICRO 2010, ATLANTA	45	18.1	19	32	E
34	MICRO 2011, PORTO ALEGRE, BRAZIL	44	21	23	26	E
35	ASPLOS 2006, SAN JOSE	38	24.1	28	30	E
	BI-ANNUAL					E
36	ASPLOS 2008, SEATTLE	31	24.4	7	16	E
37	ASPLOS 2009, WASHINGTON	29	18	19	28	E
38	ASPLOS 2010, PITTSBURGH	32	17.7	15	22	E
39	ASPLOS 2011, NEWPORT BEACH, USA	32	21.1	18	25	E
40	FCCM 2006, CALIFORNIA	25	*	34	52	E
41	FCCM 2007, CALIFORNIA	25	*	49	54	E
42	FCCM 2008, CALIFORNIA	24	28	47	54	E
43	FCCM 2009, CALIFORNIA	25	*	50	54	E
44	FCCM 2010, CHARLOTTE	24	18.2	57	79	E
45	FCCM 2011, SALT LAKE CITY	21	17.7	57	171	E
46	ICDE 2010, LONG BEACH, USA	69	12.5	49	62	I
47	ICDE 2011, HANNOVER	98	19.8	83	110	I
48	VLDB 2008, AUCKLAND, NEW ZEALAND	101	17.1	40	49	I
49	VLDB 2009, LYON, FRANCE	108	17	40	50	I
50	VLDB 2010, SINGAPORE	120	20.1	40	50	I
51	VLDB 2011, SEATTLE	100	18.1	44	56	I
52	CVPR 2009, MIAMI	61	4.2	19	32	CV
53	CVPR 2010, SAN FRANSISCO	78	4.5	27	50	CV

54	CVPR 2011, COLORADO	59	3.5	36	47	CV

TABLE 6. PRmean 1/4 and 1/3 results for all venues. *Acceptance Ratio not available.

2.2.1 H-index

An author has an index h if h of his/her N papers have at least h citations each, and the rest (N-h) papers have no more than h citations each [15]. For instance, if an author has four papers with 2,4,6 and 9 citations each then his/her h-index is equal to 3 as he/she has at least 3 papers with 3 citations each, while he/she has (4-3) = 1 paper with less than 3 citations. The h-index does not fully represent the dynamic of an author's work but can provide a strong indication of his/her recognition from his peers. The advantage of the h-index is that it combines an evaluation of quantity (number of papers) and quality (citations on each paper). In Table 7 we present the mean h-index of all authors for each conference of the subset of conferences used in Section 2.1.

Year	**MobiHoc**	**IMC**	**SenSys**	**ISCA**	**MICRO**	**ASPLOS**	**FCCM**
2006	17	20	22	16	15	15	10
2007	20	19	18	15	14	n/a	11
2008	15	15	19	16	15	16	8
2009	15	16	16	14	13	12	9

TABLE 7. Mean h-index for Networking and Electronics conferences' authors.

Table 8 presents the results for all four elements – PRmean(1/4), PRmean(1/3), Efficiency and h-index – for the SenSys conference between 2006 and 2009. Similarly to our comments in Section 2.1.1 on the Efficiency results, we observe that the 2006 version of the conference, which has the worst PRmean, included papers written by authors with the highest average h-index. This is an even more impressive result than that of Section 2.1.1. The reason is that, regarding the Efficiency metric, it could be argued that the fact that the SenSys 2006 papers received more citations was an outlier; what can not be disputed, however, is that by including papers written by authors with the highest average h-index and the worst PRmean, the SenSys 2006 proves that the PRmean value is insufficient, on its own, to reveal the quality of a venue.

Table 9 shows another very interesting result: the ISCA 2009 conference has the best PRmean values, the highest mean Efficiency and the lowest mean h-index of authors among all ISCA conferences held between 2006 and 2009. Again, this result demonstrates that no metric is able to characterize the quality of a conference on its own.

Conferences	**PRmean(1/4)**	**PRmean(1/3)**	**Efficiency**	**h-index**
SenSys06	**20**	**41**	31	**22**
SenSys07	8	17	18	18
SenSys08	8	8	27	19
SenSys09	9	20	26	16

TABLE 8. Results for SenSys symposium.

Conferences	**PRmean(1/4)**	**PRmean(1/3)**	**Efficiency**	**h-index**
ISCA06	25	32	17	16
ISCA07	22	32	20	15
ISCA08	16	30	20	16
ISCA09	15	22	**23**	**14**

TABLE 9. Results for SenSys symposium.

2.3 Conference Data

As discussed in Section 1, the Acceptance Ratio of a conference does not offer a solid indication on the quality of the conference. We have also shown in Section 1 that the conference location appears to be a factor influencing the PR results, at least in some of the venues. We discuss further the conference location factor in the following Section.

2.3.1 Location

In Section 1 we have shown the problems associated with choosing a US News Ranking for implementing PR. The restriction of ranking only US universities was only partially overcome by utilizing the ARWU rankings, because ARWU only ranks the top 100 universities worldwide, hence we had to rank equally every university outside the top 100 (with rank equal to 101). For this reason, we repeated the procedure of calculating PR values, this time with the LEIDEN ranking [14], which ranks the top 250 universities. As shown in Figures 8, 9 the LEIDEN ranking exhibits similar behavior (i.e., trend, but not actual PR values) with ARWU, with the significant difference that it leads to a larger deviation in values, as it offers the advantage of ranking a much larger pool of universities.

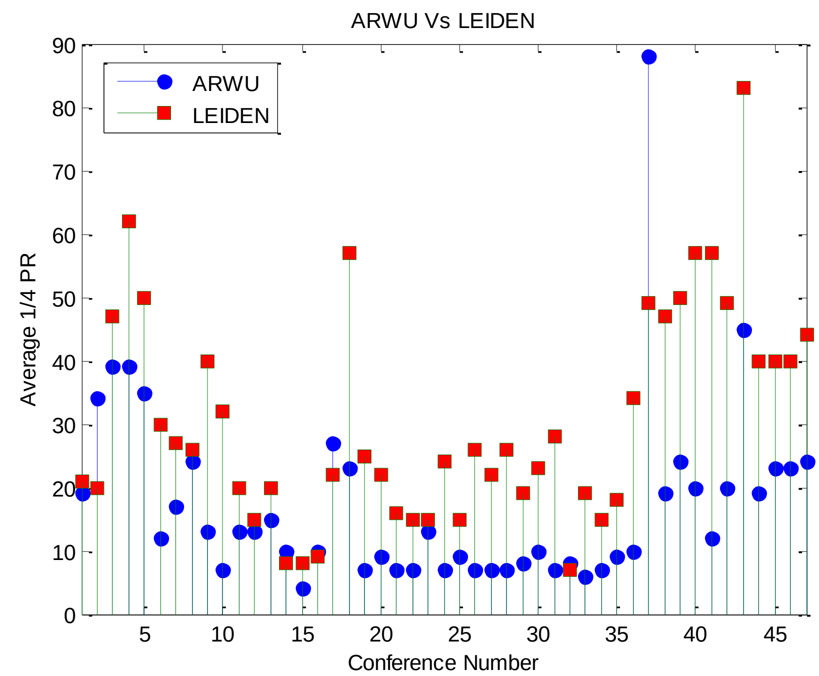

Figure 8. Average 1/4 PR results, with the use of ARWU and LEIDEN rankings

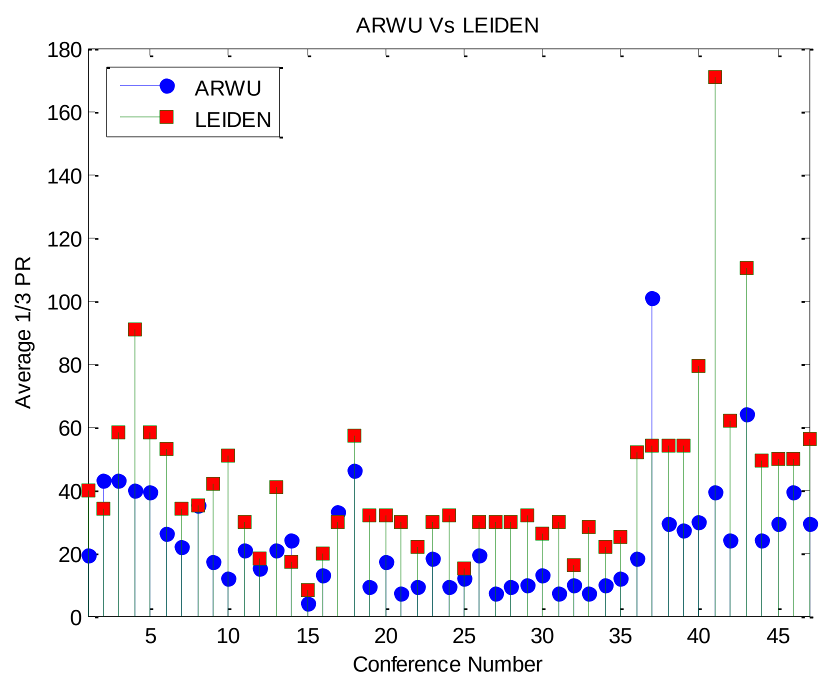

Figure 9. Average 1/3 PR results, with the use of ARWU and LEIDEN rankings

The most important result produced with the use of the LEIDEN ranking is shown in Tables 10-11: the average 1/4 and 1/3 PR results for conferences based in and outside America are identical. However, when adding the ICC and WCNC conferences in our calculations (Table 11) the results between conferences located in and outside America differ significantly. These results show that:

a) although location plays a significant role for individual venues, the use of the LEIDEN rankings for highly selective conferences leads *on average* to a minimization of the role of location in the PR results.

b) this, however, is not the case when adding to the pool of results large conferences, such as ICC and WCNC, which receive thousands of submissions from all over the world. The addition of these conferences leads to significantly better PR results for the conferences that take place within America. This result can be attributed to the fact that although these large conferences are well-known and respected in their relevant academic communities, a publication in these venues is not considered as "necessary" for prospective authors as in more selective conferences. Therefore, authors from top-ranked US universities may be unwilling to submit to these conferences if they take place, e.g., in Africa (see Table 11) or in Australia (see the difference in the respective PR values between Tables 10 and 11). On the contrary, authors from regions closer to these locations, affiliated with universities generally lower ranked than the top US ones, are likely to submit larger numbers of papers to these conferences.

The same qualitative result (i.e., conferences taking place in America having better PR results on average than those taking place outside America) is derived if we use the ARWU ranking instead of the LEIDEN one. When using ARWU, the PR results have the additional characteristic that many conferences have equal 1/3 PR results because ARWU ranks only the top 100 universities.

In Section 2.4.2 we will discuss the possible use of location for judging the quality of a conference venue.

Location Stats	**ALL**	**NUMBER OF CONFERENCES PER FIELD**			**AVERAGE PR RESULTS (LEIDEN)**	
		N	**E**	**I**	**1/4**	**1/3**
ALL	**47 (100%)**	**18**	**23**	**6**	**31**	**43**
NORTH AMERICA	**30 (63.83%)**	9	19	2	31	44
SOUTH AMERICA	**2 (4.26%)**	1	1		25	30
EUROPE	**9 (19.15%)**	5	2	2	32	45
ASIA	**3 (6.38%)**	1	1	1	34	46
AUSTRALIA	**3 (6.38%)**	2		1	23	32
AFRICA	**0 (0%)**	0			0	0
AMERICA	**32 (68.09%)**	10	20	2	31	43
WORLD (OUTSIDE AMERICA)	**15 (31.91%)**	8	3	4	31	42

TABLE 10. PR Results (LEIDEN ranking) according to location (without ICC and WCNC).

Location Stats	ALL	NUMBER OF CONFERENCES PER FIELD			AVERAGE PR RESULTS (LEIDEN)	
		N	E	I	1/4	1/3
ALL	**51 (100%)**	**22**	**23**	**6**	**40**	**55**
NORTH AMERICA	**30 (58.82%)**	9	19	2	31	44
SOUTH AMERICA	**3 (5.88%)**	2	1		72	86
EUROPE	**9 (17.66%)**	5	2	2	32	45
ASIA	**4 (7.84%)**	2	1	1	64	83
AUSTRALIA	**4 (7.84%)**	3		1	63	77
AFRICA	**1 (1.96%)**	1			124	182
AMERICA	**33 (64.71%)**	11	20	2	35	48
WORLD (OUTSIDE AMERICA)	**18 (35.29%)**	11	3	4	52	69

TABLE 11. PR Results (LEIDEN ranking) according to location (with ICC and WCNC).

2.4 Our Proposed Conference Classification Method

The results and the discussion in Sections 2.1-2.3 has shown that the (PRmean(1/4), PRmean(1/3), Efficiency, h-index) metrics can give us important information about a conference and its dynamics. A fifth important parameter is that of the conference location, on which we will comment further in Section 2.4.2.

Despite their importance, one common characteristic of all four metrics is their individual insufficiency. University rankings are controversial, the h-index can not reveal the impact of individual papers of an author (an author with an h-index of 20 can have a highest cited paper with 21 or 2000 citations, and this is not reflected in the index), and a high average Efficiency of a conference's papers may be heavily influenced by a few papers, whereas the majority may have low Efficiency. Hence, in our view it makes much more sense to try to *combine* the information offered from the four metrics than to try to base any classification on an isolated metric. Similarly, it makes much more sense to try to classify conferences into levels (Excellent, Good and Medium in our work) than to try to produce an "absolute" ranking of individual conferences. Actually, our approach offers the latter opportunity (of an "absolute" ranking) as well, and the relevant results will be presented, but we firmly believe that the classification into levels is a much fairer procedure both for the venues and the authors.

For this classification we use the k-means clustering algorithm [16].

2.4.1 K-means clustering

Clustering through the k-means algorithm aims to assign n patterns into k clusters based on the similarity between the patterns and the cluster centers; after a random initial partition, patterns are reassigned to the updated clusters (the centre of the patterns assigned to each cluster is used as a new cluster centre, also known as centroid). The procedure continues until there is no reassignment of any pattern from one cluster to another.

The k-means algorithm is popular because it is easy to implement, and its time complexity is O(n), where n is the number of patterns. A major problem of this algorithm is that it is sensitive to the selection of the initial partition and may converge to a local minimum if the initial partition is not properly chosen.

The clustering process begins when we define the number k of clusters in which we want to separate our sample – in our case k equals to 3 as we are going to separate our sample to Excellent, Good and Medium conferences. Then, we choose k centroids for the initial partition and compare the vector of values (explained below) associated with its conference with our centroids, via the Euclidian distance. We then follow the aforementioned steps of the algorithm, until no reassignment of conferences to clusters is possible.

2.4.2 *Classification Results*

We assign to every conference a vector of four values, each corresponding to one of the four metrics: [PRmean(1/4), PRmean(1/3), Efficiency, h-index]. Then, we use the k-means algorithm to classify every conference into the appropriate cluster (Excellent, Good, Medium) *separately for each metric of the vector* (we will discuss how the separate classifications can lead to a unified one).

We use two different methods to define our initial centroids, for the k-means algorithm.

In the first method, our initial centroids are derived from the function:

random(double(mean(Element_X))),

where Element_X is one of our four metrics' values, mean() gets the mean value for each Element, double() multiplies by two and random (Y) gives a random number from 1 to Y. We use the double() function because we have observed that the vast majority of values in each metric do not exceed the double of the mean.

In the second method we use as initial centroids the vector [μ-σ, μ, μ+σ] (where μ is the mean and σ is the standard deviation) for each one of the four elements.

Therefore, for every run, the first method chooses different initial centroids, whereas the second method always chooses the same initial centroids. Given the importance of the choice of the initial centroids, which we explained in Section 2.4.1, it is of significant interest to study whether the difference in the two methods of choosing initial centroids influences the classification results. Our results for the first method have been derived for 1000 different initial centroid choices.

After converging, the k-means algorithm has finished classifying the conferences into the three categories we have defined (Excellent, Good and Medium). Then, we introduce a point system where every conference takes points for every category to which the value of a metric of that conference belongs. Every category has its own weight (equal to 3 for Excellent, 2 for Good and 1 for Medium conferences). We multiply the weights by the points for each category for every metric and then we divide the result by the total amount of metrics (four) to determine the final classification of the conference. For example, ASPLOS 2009 has one Element in the Excellent category, two Elements in the Good category and one element in the Medium category. Therefore, it has a total Evaluation equal to ((1*w1) + (2*w2) + (1*w3)) / 4 = ((1*3) + (2*2) + (1*1)) / 4 = 2.

In Tables 12 and 13 we present the Networking and Electronics conferences for each centroids definition and we show how conferences are finally separated into these three categories.

Conferences	*PRmean (1/4)*	*PRmean (1/3)*	*Efficiency*	*h-index*	*Excellent*	*Good*	*Medium*	*Evaluation*
SENSYS08-RALEIGH	8	8	27	19	4	0	0	3.00
SENSYS07-SYDNEY	8	17	18	18	3	1	0	2.75
SENSYS09-BERKELEY	9	20	26	16	3	1	0	2.75
ISCA09-TEXAS	15	22	23	14	2	2	0	2.50
ASPLOS08-SEATTLE	7	16	17	16	2	2	0	2.50
MICRO06-FLORIDA	15	15	14	15	2	2	0	2.50
SENSYS06-COLORADO	20	41	31	22	2	2	0	2.50
ISCA08-CHINA	16	30	20	16	1.5	2.5	0	2.38
IMC07-SANDIEGO	26	35	20	19	1.5	2.5	0	2.38
MOBIHOC07-CANADA	20	34	17	20	1	3	0	2.25
IMC06-BRAZIL	27	34	17	20	1	3	0	2.25
ISCA07-SANDIEGO	22	32	20	15	0.5	3.5	0	2.13
ASPLOS09-WASHINGTON	19	28	26	12	1	2	1	2.00
MOBIHOC06-ITALY	21	40	19	17	0	4	0	2.00
ISCA06-BOSTON	25	32	17	16	0	4	0	2.00
MICRO07-	26	30	13	14	0	4	0	2.00

ILLINOIS								
ASPLOS06- SAN JOSE	28	30	17	15	0	4	0	2.00
MICRO08- ITALY	22	30	8	15	0	3	1	1.75
MICRO09- NEWYORK	26	30	13	13	0	3	1	1.75
IMC09- CHICAGO	32	51	15	16	0	3	1	1.75
IMC08- GREECE	40	42	11	15	0	3	1	1.75
MOBIHOC08- HONGKONG	47	58	15	15	0	2	2	1.50
MOBIHOC09- NEWORLEANS	62	91	17	15	0	2	2	1.50
FCCM06- CALIFORNIA	34	52	6	10	0	1	3	1.25
FCCM08- CALIFORNIA	47	54	4	8	0	0	4	1.00
FCCM07- CALIFORNIA	49	54	4	11	0	0	4	1.00
FCCM09- CALIFORNIA	50	54	4	9	0	0	4	1.00

TABLE 12. Clustering for Networking and Electronics conferences – First Centroid Initialization Method.

Conferences	*PRmean (1/4)*	*PRmean (1/3)*	*Efficiency*	*h-index*	*Excellent*	*Good*	*Medium*	*Evaluation*
SENSYS08-RALEIGH	8	8	27	19	4	0	0	3.00
SENSYS09-BERKELEY	9	20	26	16	3	1	0	2.75
ISCA09-TEXAS	15	22	23	14	3	1	0	2.75
SENSYS07-SYDNEY	8	17	18	18	3	1	0	2.75
ASPLOS08-SEATTLE	7	16	17	16	2	2	0	2.50
MICRO06-FLORIDA	15	15	14	15	2	2	0	2.50
SENSYS06-COLORADO	20	41	31	22	2	2	0	2.50
ISCA08-CHINA	16	30	20	16	1	3	0	2.25
IMC07-SANDIEGO	26	35	20	19	1	3	0	2.25
MOBIHOC07-CANADA	20	34	17	20	1	3	0	2.25
IMC06-BRAZIL	27	34	17	20	1	3	0	2.25
ISCA07-SANDIEGO	22	32	20	15	0	4	0	2.00
ASPLOS09-WASHINGTON	19	28	26	12	1	2	1	2.00
MOBIHOC06-ITALY	21	40	19	17	0	4	0	2.00
ISCA06-BOSTON	25	32	17	16	0	4	0	2.00
ASPLOS06-	28	30	17	15	0	4	0	2.00

SANJOSE								
MICRO07-ILLINOIS	26	30	13	14	0	4	0	2.00
MICRO09-NEW YORK	26	30	13	13	0	4	0	2.00
MICRO08-ITALY	22	30	8	15	0	3	1	1.75
IMC09-CHICAGO	32	51	15	16	0	3	1	1..75
IMC08-GREECE	40	42	11	15	0	2	2	1.50
MOBIHOC08-HONGKONG	47	58	15	15	0	2	2	1.50
MOBIHOC09-NEWORLEANS	62	91	17	15	0	2	2	1.50
FCCM06-CALIFORNIA	34	52	6	10	0	1	3	1.25
FCCM08-CALIFORNIA	47	54	4	8	0	0	4	1.00
FCCM07-CALIFORNIA	49	54	4	11	0	0	4	1.00
FCCM09-CALIFORNIA	50	54	4	9	0	0	4	1.00

TABLE 13. Clustering for Networking and Electronics conferences – Second Centroid Initialization Method.

As we can see from Tables 12-13, which have very minor differences in the ranking order of individual venues, we derive the same number of Excellent conferences for both methods, and +/-1 Good/Medium conferences. This almost negligible difference shows that our approach of running the k-means algorithm for a very large number of different initial centroid choice leads to a minimization of the problem of possible convergence to a local minimum. It also shows, of course, that the use of the vector [μ-σ, μ, μ+σ] as initial centroids suffices; hence, the first method of choosing centroids is not necessary, as the second method can guarantee that the computational complexity of our classification approach will be small.

The number of conferences in each category (Excellent, Good, Medium) is selected as follows: by implementing, e.g., the second method, we find that seven conferences are Excellent in terms of their PRmean(1/4), six conferences are Excellent in terms of their PRmean(1/3), six conferences are Excellent in terms of their Efficiency and five conferences are Excellent in terms of their authors' average h-index; this leads us to an average of six ((7+6+6+5)/4 = 6) conferences in the overall Excellent category. However, there are three conferences which are ranked between fifth and seventh place, with Evaluation equal to 2.50.

This is where the location factor could be utilized. More specifically, as we saw earlier, both in the case of the ARWU ranking and of the LEIDEN ranking when larger conferences are considered, conferences that take place in America (and mostly in the USA) have, on average, much better PR results than those held outside of America. This means that we can use location as a tie-breaker: if two (or more) conferences are tied in their Evaluation values, those that have been held outside USA should be favored in the classification, as they start with a "disadvantage". Still, in this work we have focused only on top-tier conferences, for which the LEIDEN ranking gives on average tied results between conferences held in and outside America. Therefore, we don't use location as a tie-breaker and instead we consider all first seven conferences to be in the Excellent category, in the previous example. Finally, regarding the point system related to the first centroid initialization method, we need to point out the following: in the cases where, for an individual metric, the cluster of a category was computed to contain N conferences, but there was a tie between 2 or more conferences (hence there were conferences "hanging in the middle" between categories), we gave that conference half a point for each category.

Table 14 presents the average results of our point system for all venues of the conferences presented in the previous Tables. Note that the results are almost identical for both methods.

Method	**1st**		**2nd**	
Ranking	**Symposium**	**Points**	**Symposium**	**Points**
1	SenSys	2.75	SenSys	2.75
2	ISCA	2.25	ISCA	2.25
3	ASPLOS	2.17	ASPLOS	2.17
4	MICRO	2.00	MICRO	2.06
5	IMC	2.03	IMC	1.94
6	MobiHoc	1.81	MobiHoc	1.81
7	FCCM	1.06	FCCM	1.06

TABLE 14. Final cumulative ranking of conferences (all venues).

2.4.3 *Final Comparison with AR*

J. Chen and J. A. Konstan reported in [13] a rather strange phenomenon. They found that conferences with 10-15% AR have papers that are less cited than those of conferences with 15-20% AR. When considering only the average Efficiency of a conference (Section 2.1) our results agree with those in [13]: four conferences that have an AR in the range 10-15% have an average Efficiency equal to 17.75, while nine conferences with 15-20% AR have an average Efficiency equal to 19.78. Even more interestingly, 5 out of the 9 conferences with 15-20% are rated "Excellent", 3 out of 9 are rated "Good" and only 1 out of 9 belongs to the "Medium" category, while 3 out of 4 of the 10-15% AR conferences are rated "Good" and 1 out of 4 belongs to the "Medium" category. These results once again show the inadequacy of the AR in revealing substantial information about the quality of a venue, and, as a consequence, of an author's work.

3. Conclusions and Future Work

The use of the paper Acceptance Ratio of a conference is not a trustworthy metric in order to evaluate a researcher's contribution in his/her field. The PR metric, proposed in [1], is a rough coarse-grain metric proposed for making quick assessments of a researcher's contribution. Although PR is an idea towards the right direction in our view, it has some important shortcomings which we tried to point out, and then to overcome by proposing a new, finer-grain conference classification method which incorporates PR but alleviates its shortcomings.

Our future work will include a study on the possibility of incorporating the φ-index into our classification approach and the ranking of a much larger number of conferences, using the conference location as a tie-breaker in the way explained in Section 2.4.2.

We believe that our proposed conference classification approach can offer a better and more complete (thus fairer) method for judging the quality of the venue where an author publishes his/her work. Furthermore, with this new approach each conference can make its annual self-evaluation and draw useful conclusions from the gathered data over time. We hope that our approach can bring a fresh look into the current conference evaluation system.

References

[1] S. Nelakuditi, C. Gray and R. Roy Choudhury, "Snap Judgement of Publication Quality: How to Convince a Dean that you are a Good Researcher", ACM SIGMOBILE Mobile Computing and Communications Review, Volume 15, No. 2, 2011, pp. 20-23.

[2] http://grad-schools.usnews.rankingsandreviews.com/best-graduate-schools/top-science-schools/computer-science-rankings

[3] http://graduate-school.phds.org/rankings/

[4] http://www.arwu.org/FieldENG2010.jsp

[5]http://www.shanghairanking.com/SubjectCS2012.html

[6] Networking Conference Statistics, http://www.cs.ucsb.edu/~almeroth/conf/stats/

[7] http://vldb.org/archives/statistics.html

[8]http://students.engr.scu.edu/~sschaeck/acceptancerates/

[9]http://dl.acm.org/citation.cfm?id=1934902&picked=prox

[10]http://dl.acm.org/citation.cfm?id=1815961&picked=prox

[11]A. Schubert, “A Hirsch-type index of co-author partnership ability”, Scientometrics, Vol. 91, No. 1, 2012, pp. 303-308.

[12] S. Uddin, L. Hossain, A. Abbasi and K. Rasmussen, “Trend and efficiency analysis of co-authorship network”, Scientometrics, Vol. 90, No. 2, 2012, pp. 687-699.

[13] J. Chen and J. A. Konstan, “Conference Paper Selectivity and Impact”, Communications of the ACM, Vol. 53, No. 6, 2010, pp. 79-83.

[14] http://www.leidenranking.com/ranking.aspx

[15] J. E. Hirsch, "An Index to Quantify an Individual's Scientific Output", Proceedings of the National Academy of Sciences of the United States of America, Vol. 102, No. 46, 2005, pp. 16569-16572.

[16] A. K. Jain, M. N. Murty and P. J. Flynn, "Data Clustering: A Review", ACM Computing Surveys, Vol. 31, No. 3, 1999, pp. 264-323.