Assessing the Research and Education Quality of the Top Brazilian Computer Science Graduate Programs

Alberto H. F. Laender[1] Carlos J. P. de Lucena[2] José Carlos Maldonado[3] Edmundo de Souza e Silva[4] Nivio Ziviani[1]

[1]Departarment of Computer Science
Federal University of Minas Gerais
Belo Horizonte, Brazil
{laender,nivio}@dcc.ufmg.br

[2]Departament of Informatics
Pontifical Catholic University of Rio de Janeiro
Rio de Janeiro, Brazil
lucena@inf.puc-rio.br

[3]Institute of Mathematical Sciences and Computing
University of Sao Paulo
São Carlos, Brazil
jcmaldon@icmc.usp.br

[4]System Engineering and Computer Science
Program - COPPE
Federal University of Rio de Janeiro
Rio de Janeiro, Brazil
edmundo@cos.ufrj.br

Abstract: This article reports about a study conducted to assess the quality of the top Brazilian Computer Science graduate programs. The study is based on data from DBLP and considers the scientific production of these programs in the triennial 2004-2006. A comparison of the scientific production of the Brazilian programs against that of reputable programs in North America and Europe indicates that the former compares well with these programs, both in terms of publication rate and number of graduates. The study also shows that the Brazilian programs follow international publication ratios of more than two conference papers per journal article. These results are a clear indication that the Computer Science field has reached maturity in Brazil.

Categories and Subject Descriptors: K.3.0 [Computers and Education]: General; K.3.2 [Computers and Education]: Computer and Information Science Education - *Computer science education*
General Terms: Human Factors, Economics
Keywords: Computer Science, Graduate Program, Evaluation Process, Brazil

1. HISTORICAL PERSPECTIVE

Since 1977, the Brazilian Ministry of Education's agency, CAPES[1], has been evaluating all graduate programs in the country in all fields of knowledge and, every three years, generates a quality ranking of these programs based on the recommendations of committees specially appointed for this task. The aim of this assessment is not only to produce for society and students, in particular, a measure of the quality of existing programs but also to establish a ranking to subsidize government spending in advanced research and education. This process is based on a very thorough peer review system and has been responsible for the good reputation enjoyed by the Brazilian graduate school system.

Although the evaluation process is well established, in the last few years the Computer Science (CS) field underwent a fantastic boom that resulted in the creation of several new programs (a growth of 200% in the last 10 years). This is certainly a unique characteristic of CS compared to other fields of study in the country. What used to be a straightforward comparison of 16 programs became a comparative analysis involving more than 50 programs. Moreover, the enormous growth of the number of programs was followed by an expressive increase in the number of MSc and PhD graduates as well as in the scientific production in the form of journal articles and conference papers. According to data from CAPES[2], over the last two evaluation triennium there has been an increase of 67% and 44% in the graduation rate of Master's and PhDs, respectively (see Table 1). In addition, the scientific production in this last triennial has reached a total of 911 journal articles and 3,669 conference papers. In this scenario, indexes based on bibliometric data and other indicators became necessary as ancillary tools to perform the job at hand.

[1] http://www.capes.gov.br
[2] http://www.capes.gov.br/sobre/estatisticas

In order to assess the maturity of the Brazilian CS graduate programs, it was also necessary to select those that could be considered as "reference programs." An additional issue faced by CAPES was the comparative analysis between CS and other fields, such as Biology and Physical Sciences, used as a tentative to make the evaluation process uniform, taking into consideration all fields. Certainly, the last task is inherently difficult and prone to errors.

Table 1: MSc and PhD graduates.

	1998-2000	2001-2003	2004-2006
MSc	1,350	2,354	2,254
PhD	174	238	312

One possible way to try to deal with the above issues is to promote a vertical comparative analysis with CS graduate programs, with unquestionable good reputations, from other countries. For instance, CS graduate programs in North America and Western Europe that would share similarities with those in Brazil in terms of organization and research goals. CS has become a key discipline for the social and economic development of developing countries. More so than in developing countries, it is strategic to establish a ranking of the local CS programs to guide government education and research policies for the field. As far as the authors know, this has not been done yet in developing countries other than Brazil.

Thus, the goal of this article is to report the Brazilian experience in assessing the quality of its CS graduate programs so that it can be used in other countries at the same stage of development of Brazil. This assessment effort started six years ago and has undergone two triennial evaluations: 2001-2003 and 2004-2006. We begin by presenting a brief historical overview of the CAPES CS evaluation in Brazil and comment on some of the major issues of such process. Then we present the methodology recently used to address such issues and discuss relevant results of the study we conducted to compare the scientific production of the top CS graduate programs in Brazil against that of reputable programs in North America and Europe. This comparison study is based on data from DBLP [6] and considers the scientific production of these programs in the triennial 2004-2006. We also discuss important open problems that must be addressed for improving the quality analyses of the field.

2. CAPES CS EVALUATION PROCESS

The process conducted by CAPES to evaluate all graduate programs in Brazil considers several different quantitative and qualitative metrics. Many of these metrics are identical for all fields, but there are key differences. We focus on those used for the CS field. Examples of these metrics include a subjective evaluation of the physical infrastructure of the programs (laboratory facilities, libraries, etc.); the coverage of the courses offered with respect to a CS core and their main subfields (based on, for instance, ACM publications on CS curricula); the curriculum vitae of the faculty members; the quality of theses and dissertations, measured, for instance, by the fraction of papers published in journals and conferences to the number of students graduated; and the national and international reputation of the program under evaluation.

One of the most relevant items in the evaluation process is clearly the number of publications per faculty member and some measure to assess the publication quality. Needless to say, this is the most difficult issue in the process. CAPES requires that the advisory committees for each field of knowledge propose a set of publication venues (journals and conferences) to be considered in the evaluation; that is, only the publication venues in this set, called the field Qualis, are counted for measuring the programs' publication efficiency. This set is subdivided into two for publications considered national or international, and each subset is further divided into categories A, B, or C, according to some measure of publication quality. For international publications, most scientific fields of knowledge consider only journals for their Qualis and use the ISI JCR impact factor[1] to classify the publications in categories A, B or C, after establishing thresholds for each subset.

However, considering only journals is unacceptable for the CS field, as in fact is observed by several existing works on the subject. First, it is well known that peer reviewed conference proceedings have a major impact on the development of CS [12, 13]. It is evident that conference papers must be taken into account in any evaluation of CS programs. Second, ISI journals alone do not accurately cover the broad CS field [9]. Young but prestigious journals exist and must be taken into account. Therefore, one must include other data sources in the process [1, 15]. A third issue arises when one compares CS with other fields of knowledge. This is true since not only the impact factor of journals differs significantly among fields but also the publications per faculty member ratio may be quite different.

[1] http://scientific.thomson.com/products/jcr/

To address these issues, our first approach, back in 2001, to build the CS Qualis for international publication venues was to consider not only ISI journals but also journals and conferences based on CiteSeer[1] and to use the respective impact indexes. Using CiteSeer was a fast and reasonably reliable way to establish a first ranking mechanism that included a rich set of CS international conferences. Unfortunately, the CiteSeer index stopped being maintained and, consequently, it was evident that one would have to use broad rules and some form of ranking to keep updating the CS Qualis. As a basic rule, it was included in the CS Qualis conferences sponsored by major scientific societies, such as ACM and IEEE, and ranked according to the data available for them (e.g., the number of years the conference has taken place, the fraction of accepted papers in the last many years, the international prestige of the technical program committee or editorial board, among others).

The CS CAPES advisory committee has also consulted the Brazilian Computer Society[2] (SBC). Through the SBC representatives for each CS subfield, several documents were produced providing information and evidence of the quality of the conferences and journals considered relevant for each subfield. Senior faculty members in Brazilian CS programs have also been asked to contribute information on international conferences and journals. A similar approach was taken for national conferences. Today all CS national conferences are providing SBC with statistics concerning acceptance rate, number of submissions, etc., in order to build a reliable database that can be used by CAPES.

The CS CAPES advisory committee has also taken into account the multidisciplinary aspects of CS field. In this way, the journal ranking of sister fields such as Mathematics, Applied Mathematics, Electronic Engineering, Statistics and Optimization have also been included in the CS Qualis. In order to ensure that graduate programs are built on a solid CS core, the committee enforces a limit on the fraction of publications from other fields with respect to the core CS publications that can be included in the evaluation process. Thus, the CS CAPES advisory committee has used the ranking and thresholds of the CS Qualis to obtain indexes that guide the evaluation of each program.

Today, although we have a good number of journals and conferences in the CS Qualis ranked according to the guidelines mentioned above, it is clear that the indexes alone cannot provide an accurate metrics to compare different fields of knowledge as argued above. Furthermore, although these indexes may serve as a first order indicator of the evolution of the CS field in Brazil, they are not appropriate for comparing the Brazilian programs with other CS programs in major universities abroad. But international comparison is of major importance to the CAPES evaluation process since the programs ranked in the top CAPES levels are considered to perform as well as major programs in North America or Europe.

3. COMBINING BIBLIOMETRICS WITH OTHER QUALITY AND PRODUCTIVITY INDICATORS

In order to properly assess the quality of the top Brazilian CS graduate programs and determine the level of internationalization of these programs with respect to their scientific production, we conducted a study to compare them with programs of unquestionable reputations in North America and Europe. In this study, we used DBLP - Digital Bibliography & Library Project[3] [6] as the main data source and an important productivity index expressed by the graduated PhDs per advisor ratio. Special consideration was given to the question of the importance of conferences as a venue for presenting research results in the CS field [12, 13] and the difficulty of establishing a conference ranking that would greatly improve the bibliometrics applicable to the field.

Table 2: List of requirements for a CS graduate program to be considered a reference for the field.

#	Description
1	Research activities within the program, including new fields of knowledge in CS.
2	Creative collaboration in innovative projects with the Brazilian industry and society.
3	Participation in a broad collaborative scientific network on the domestic and international scene, involving top and emerging institutions that would benefit from the experience and quality of the top program.
4	Production of Masters' and PhD degree holders in the CS field in the quantity and level of quality equal to the worlds' main institutions.
5	PhD degree holders from the program working in top tier universities and research institutes in Brazil and overseas.
6	International insertion evidenced by participation in qualified international conference program committees, editorial boards of first line periodicals, teaching and research activities in internationally renowned universities and research institutes, significant participation in domestic and international scientific societies and analogous activities exercised by a significant number of professors.
7	Awards and distinctions of national and international repercussion obtained by a variety of the professors in the program.
8	A quantity and quality of the publications in journals and proceedings of conferences of international reputation at a level equivalent to the top CS programs in the world.

.

[1] http://citeseer.ist.psu.edu/
[2] http://www.sbc.org.br/
[3] http://dblp.uni-trier.de

The CS CAPES advisory committee stated upfront a comprehensive list of requirements for a CS graduate program to be considered a reference in the field (i.e., those programs that correspond to the levels 6 and 7 in CAPES grading scale), which is presented in Table 2. The CAPES evaluation process attributes grades of 6 and 7 to doctoral and masters' programs with "international insertion." For the 2004-2006 triennial these requirements were used to assess those programs that, potentially, could receive these grades. Those that fully satisfied the requirements were considered to be level 7, whereas those that demonstrated to be closely complying with the requirements received a grade of 6. Two of these requirements, respectively, requirements 4 and 8 in Table 2, are the ones that were the target of our study.

3.1 The Choice of the Samples: Brazilian, North American and European Institutions

Our study addressed the top-8 Brazilian graduate programs according to the CAPES ranking, 16 North American graduate programs (from Canada and US) and 6 European graduate programs (from England, France and Switzerland), as follows:

- *Brazil*: Federal University of Minas Gerais, Federal University of Pernambuco, Federal University of Rio de Janeiro, Federal University of Rio Grande do Sul, Pontifical Catholic University of Rio de Janeiro, University of Campinas, University of São Paulo at São Paulo, and University of São Paulo at São Carlos.
- *North America*: Brown University, California Technology Institute, Carnegie Mellon University, Cornell University, Harvard University, University of Illinois at Urbana-Champaign, Massachusetts Institute of Technology, Princeton University, Stanford University, University of California at Berkeley, University of British Columbia, University of Texas at Austin, University of Toronto, University of Washington, University of Waterloo, and University of Wisconsin.
- *Europe*: Cambridge University, École Polytechnique, ETH Zürich, Imperial College, Oxford University, and Université Pierre et Marie Curie –Paris VI.

Eastern institutions were not considered in our study because they differ from the Brazilian cultural patterns. The sample was based on the ease of accessing information about the institutions, which are all included in the top positions of existing rankings of CS programs.

3.2 The Methodology

The data gathering process for our study involved four main steps. In the first step, we extracted from the home pages of the respective institutions the names of the faculty members of the 30 graduate programs considered. In the second step, using the list of extracted names, we collected the respective pages from DBLP and extracted the corresponding bibliographic citation data. In the third step, we assigned a Qualis category to each journal and conference found in DBLP. We notice that the CS Qualis includes a subset of the publication venues covered by DBLP and, therefore, not all journals and conferences got a category assigned to them. Finally, in the fourth step, we stored the resulting data in a relational database in order to make use of a more flexible querying environment.

As result of the data gathering process, 52,156 bibliographic citations from 2,007 faculty members from the 30 programs were stored in the database. The bibliographic citations were collected from the DBLP repository on June 27, 2007 and refer to articles published between 1954 and 2007 in 452 journals and in the proceedings of 1,611 conferences.

From the relational database it is possible to derive a variety of statistics about the 30 programs addressed in the study, such as: the total number of publications per faculty member per period, the average number of publications per faculty member per period, the average number of publications in journals per faculty member per period, the average number of publications in conference proceedings per faculty member per period, the most popular publication venues (journals and conferences) per program, the most prolific faculty members per program, among others. Moreover, because most publications venues are classified according to the CS Qualis, this allows for a qualitative evaluation of the scientific production of each program to take place.

A key step in our data gathering process was the assignment of a CS Qualis category to each publication venue (journal or conference). Although the existing version of the CS Qualis covered a reasonable number of CS journals, since it is mainly based on the CiteSeer and ISI JCR indexes, its coverage of conferences was quite limited since it basically contemplated those venues in which Brazilian researchers have published over the past years. To overcome this limitation and to avoid any bias towards the Brazilian programs, we carried out a subjective poll with the participation of members of the Brazilian CS research community to expand the CS Qualis conference set. To support this poll, we developed a Web-based platform to collect the opinion of experts from several distinct CS subfields. For this, we registered 916 researchers and faculty members from 51 CS graduate programs in Brazil. The poll was conducted in two phases:

- In the first phase, carried out from September 19 to October 1, 2007, we set a forum to discuss an initial list of 962 conferences distributed into 24 CS subfields, with a participation of 502 subjects from the 916 registered in the system. This initial list of conferences comprised conferences taken from the original CS Qualis, the DBLP

conference index page[1] and the CORE ranking[2]. However, only conferences with at least four editions and regularly held in the last three years were considered. In addition, regional conferences, such as those from the Pacific-Asian and Eastern-European regions, were also discarded because they are not usually a target venue to the Brazilian CS community. Based on the suggestions received from 226 messages posted during the discussion forum, we ended up with a final list of 1,052 conferences distributed into 27 subfields (90 conferences and 3 subfields were added to the initial proposed list), as shown in Table 2. These 27 subfields basically reflect the scope of the 24 special interest groups of the Brazilian Computer Society.

- In the second phase, carried out from October 6 until October 30, 2007, we conducted a subjective poll in which each participant subject classified the conferences belonging to each subfield into five categories: A, B, C, NC ("not considered"), and NE ("not evaluated"). The rules to classify each conference of a subfield into these categories followed the guidelines established by the CS CAPES advisory committee for classifying publication venues. Thus, to conform to these guidelines, a maximum of 40% of the conferences could be classified as A and the sum of votes in A or B could not exceed more than 80% of the total number of conferences on the ballot. In the NC category should be included those conferences that, according to the subject's opinion, could not be classified as A, B or C. Those conferences for which the subject did not have enough information to fit them in A, B, C or NC should be left without evaluation (NE). In this phase, we had 312 subjects and 875 fits (2.8 fits in distinct subfields per subject on average).

3.2.1 The Data Source

Similar studies, in different fields of knowledge, have been made based on the Thomson ISI[3] database. This database contains references to more than 8,700 journals in different fields. However, the Thomson ISI database is not adequate for bibliometric studies in the CS field, since it mainly focuses on the fields of Natural Sciences and Life Sciences, and covers only a small fraction of conferences. Mattern [9], in a study based on the 2003 scientific production of ETH, Zürich, shows that the Thomson ISI database covered only 14% of the CS publications of that institution, while it covered approximately 60% of the publications in Physics, Chemistry and Biology. Moreover, a look at the Thomson ISI list of the top 250 researchers with more citations in the CS field shows that this list includes none of the ACM Turing Award[4] winners in the last 10 years, clear evidence that this database does not cover the relevant publication venues in the field.

Instead, we used DBLP, a digital library that contains one of the largest repositories of CS bibliographic citations. Currently, DBLP indexes more than 1,000,000 publications, from more than 600 journals and 3,600 conferences [8]. The DBLP service evolved from a small, specialized bibliography repository to a digital library covering most subfields of CS. Although this coverage is not uniform for all CS fields, it is quite representative for important subfields, such as Algorithms and Theory, Artificial Intelligence, Computer Architecture, Computer Graphics and Image Processing, Computer Networks, Databases, Programming Languages, and Software Engineering, as we discuss next. Moreover, DBLP is maintained with massive human intervention in the acquisition and loading stages, with rigid data quality requirements [7]. Its content is highly normalized with respect to author and publication venue names, limiting the occurrence of ambiguity among author names and facilitating the correct identification and classification of journals and conferences [5]. Additionally, the DBLP content is totally open, with no commercial interest involved, which facilitates the access to its data.

As a consequence, DBLP has been frequently used as a primary data source for bibliometric studies and analysis of collaboration and citation networks in the CS field [2, 3, 4, 10, 11, 14].

3.2.2 DBLP Coverage

Considering that conference papers have a major impact on any bibliometric study in the CS field, we performed an analysis of the DBLP conference coverage. To determine this coverage, we considered the CS field divided into the 27 subfields shown in Table 3 and used as reference the list of 1,052 top conferences compiled for the poll we conducted to expand the Qualis conference set. Taking the complete list of conferences indexed by DBLP [8], we looked for each conference in the reference list. As we can see from Table 3, the average coverage considering all subfields achieves 67%. For 17 out of 27 subfields, DBLP covers more than 65% of the conferences in the reference list. For 11 subfields the coverage is greater than 75% and for only four subfields the coverage is smaller than 40%. These numbers show that the DBLP conference coverage is expressive and contemplates a wide range of CS subfields.

It is worth noting that, in addition to its good conference coverage, DBLP also covers a substantial part of the most important CS journals, particularly those published by two of the most prestigious CS scientific societies, ACM and IEEE.

[1] http://www.informatik.uni-trier.de/~ley/db/conf/index.a.html.
[2] http://www.core.edu.au
[3] Now Thomson Scientific: http://www.scientific.thomson.com/isi/.
[4] The most important award given annually by the Association for Computing Machinery to researchers with outstanding contributions to the advancement of CS.

Table 3: DBLP conference coverage.

CS Subfield	Poll	DBLP	Coverage (%)
Algorithms and Theory	45	37	82
Applied Computing	46	16	35
Artificial Intelligence	60	41	68
Computational Biology	12	10	83
Computer Architecture, High Performance Systems and Operating Systems	69	49	71
Computer Education	24	8	33
Computer Graphics, Image Processing and Computer Vision	69	46	67
Computer Networks, Distributed Systems and P2P Systems	85	50	59
Databases, Information Retrieval, Digital Libraries and Data Mining	45	43	96
Embedded, Real Time and Fault Tolerant Systems	24	19	79
Formalism, Logics and Computational Semantics	35	29	83
Games and Virtual Reality	27	15	56
Geoinformatics	13	7	54
Human Computer Interaction	27	22	81
Information Systems	14	10	71
Integrated Circuits Design	46	30	65
Machine Learning	40	30	75
Multi-thematic	17	13	76
Natural Language Processing	32	16	50
Operational Research and Combinatorics	28	6	21
Programming Languages	42	32	76
Robotics and Control & Automation	40	7	18
Security	39	31	79
Simulation and Modeling	16	10	63
Software Engineering and Formal Methods	83	60	72
Ubiquitous Computing	31	25	81
Web and Multimedia & Hypermedia Systems	43	35	81
Total	1,052	700	67

3.3 Comparison of the Scientific Production

Table 4 summarizes the main results of our study and shows the average number of publications per faculty member in the triennial 2004-2006 for all 30 programs addressed. In this table, columns All, A and B, under the heading Journal Articles, correspond, respectively, to all types of article (qualified[1] and not qualified), only articles of category A and only articles of category B (likewise for the columns under the heading Conference Papers). The last column All corresponds to all types of publication (qualified and not qualified). We notice that Table 4 does not show a column for category C publications because, for most programs, the average number of such publications per faculty number was not significant (overall mean of 0.02 and 0.18 for journal articles and conference papers, respectively) and, therefore, had a small impact in our comparative analysis. In the following, we analyze the relative position of the programs without identifying them and considering three distinct groups: NGP –North-American Graduate Programs, EGP –European Graduate Programs, and BGP –Brazilian Graduate Programs. For further reference, each program is numbered according to its relative position within its respective group.

Considering all publications in the period, the Brazilian programs present an average number of publications per faculty member between 3.84 and 8.10. These numbers show that the volume of publications of the Brazilian programs can be directly compared with those of at least seven of the overseas programs, three North Americans and four Europeans (bottom half of the table). Moreover, the top-2 Brazilian programs present an average number of publications per faculty member of 8.10 and 7.23, respectively, which places them close to the overall mean of the 30 programs evaluated (8.49 ± 1.04). The next three Brazilian programs, with an average number of publications per faculty member of 6.53, 6.41 and 6.34, respectively, are not far from the overall mean and above two North American programs and three Europeans.

[1] In our study, we refer to a publication as qualified if it appears in a publication venue (journal or conference) that has a Qualis category (A, B or C) assigned to it.

Table 4: Average number of publications per faculty member in the triennial 2004-2006.

	Programs	Journal Articles			Conference Papers			All
		All	A	B	All	A	B	
1	NGP-1	4.50	2.92	0.39	10.48	6.88	1.42	14.98
2	NGP-2	4.04	2.73	0.12	9.00	5.73	1.51	13.04
3	NGP-3	3.65	2.29	0.15	8.83	5.56	1.20	12.48
4	EGP-1	3.42	1.64	0.30	7.76	4.06	1.09	11.18
5	NGP-4	2.29	1.44	0.18	8,89	6.06	1.21	11.18
6	NGP-5	3.82	1.53	0.41	7.29	5.41	0.82	11.11
7	NGP-6	3.77	1.97	0.26	7.14	5.00	0.74	10.91
8	NGP-7	2.82	1.57	0.20	7.65	5.04	0.96	10.47
9	NGP-8	2.78	1.71	0.19	7.59	4.86	1.11	10.37
10	EGP-2	2.70	1.37	0.26	7.23	3.84	0.88	9.93
11	NGP-9	2.56	1.79	0.21	7.12	4.59	1.03	9.68
12	NGP-10	2.84	1.78	0.04	6.67	4.09	0.67	9.51
13	NGP-11	2.77	1.50	0.27	6.23	4.33	0.67	9.00
14	NGP-12	2.32	1.67	0.12	6.67	4.36	0.65	8.99
15	NGP-13	3.09	1.74	0.51	5.78	2.64	1.03	8.87
16	NGP-14	2.44	1.51	0.15	6.38	4.59	0.82	8.82
17	BGP-1	1.86	1.14	0.33	6.24	2.62	1.43	8.10
18	EGP-3	3.27	1.53	0.20	4.43	2.43	1.07	7.70
19	BGP-2	2.50	1.47	0.67	4.73	1.53	1.33	7.23
20	BGP-3	1.21	0.41	0.21	5.32	2.76	0.65	6.53
21	BGP-4	1.89	1.22	0.26	4.52	2.30	1.00	6.41
22	BGP-5	1.09	0.54	0.20	5.25	2.66	0.52	6.34
23	NGP-15	2.33	1.50	0.00	4.00	2.42	0.42	6.33
24	NGP-16	1.48	1.14	0.02	4.46	2.38	0.84	5.94
25	EGP-4	1.49	0.62	0.26	4.39	2.05	0.60	5.88
26	BGP-6	1.03	0.50	0.27	4.70	1.97	1.00	5.73
27	EGP-5	1.95	0.92	0.16	2.79	1.37	0.37	4.74
28	BGP-7	1.74	1.15	0.21	3.00	1.41	0.41	4.74
29	EGP-6	1.63	0.74	0.29	3.10	1.29	0.38	4.73
30	BGP-8	1.57	1.00	0.37	2.27	0.73	0.43	3.84

We also observe that the Brazilian programs' scientific production is strongly centered on conference papers in a ratio of 2.89 conference papers to each journal article. Although slightly higher than those of the North American (2.49:1) and European (2.10:1) programs, this ratio shows that the Brazilian CS community follows a similar pattern and also considers conferences as the preferable venue for presenting their research results. This is a very important finding because it shows that any kind of productivity assessment in the CS field, should it be institutional or seen from an individual viewpoint (e.g., in the case of grant or scholarship applications), cannot be carried out without considering conference publications. Moreover, in today's global scenario, the attendance of major international conferences opens a huge spectrum of opportunities, particularly for researchers from developing countries like Brazil, which makes these conferences key publication venues in the field.

A more detailed analysis of the data in Table 4 can be found in the graphics of Figures 1 to 3, which consider only highly qualified publications, i.e., those publications that have appeared in journals or proceedings of conferences classified as A or B according to the CS Qualis.

Figure 1 shows that one Brazilian graduate program (BGP-2) presents an average number of highly qualified journal publications per faculty member of 2.14, which places it among the top-6 programs in this specific comparison. Four other Brazilian programs (BGP-4, BGP-1, BGP-8 and BGP-7), with an average number of highly qualified journal publications per

faculty member between 1.37 and 1.48, appear just behind one North American program (NGP-15) and ahead of another North American program (NGP-16) and three Europeans (EGP-5, EGP-6 and EGP-4).

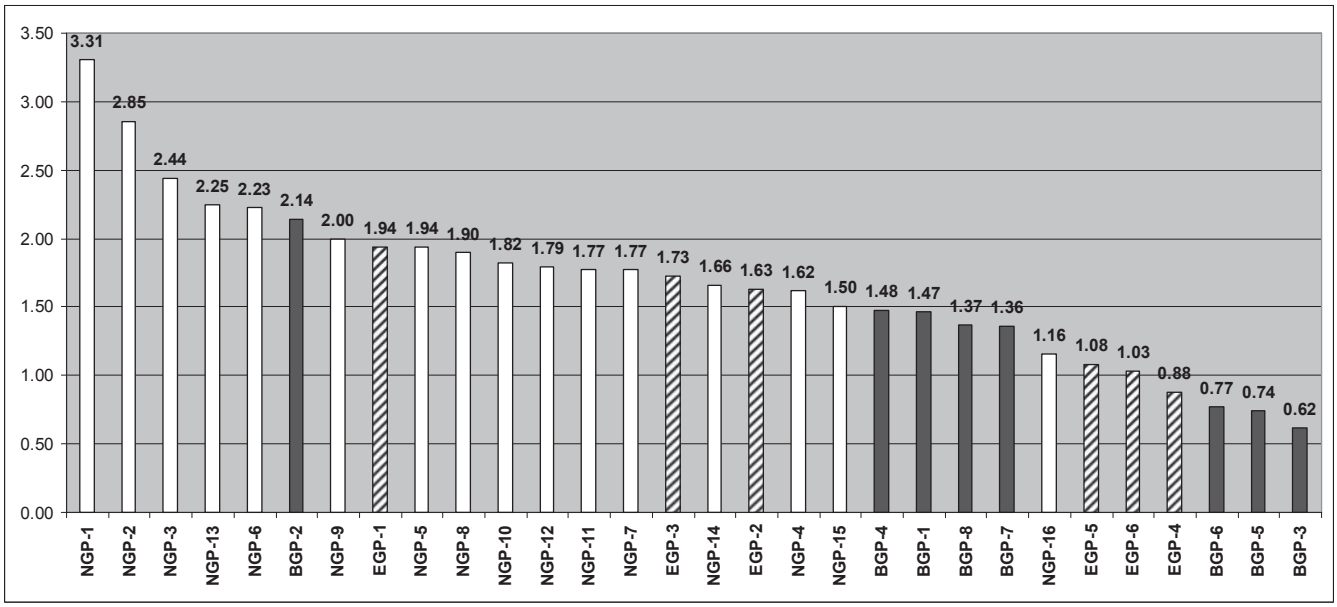

Figure 1: Highly qualified journal publications (Qualis A and B).

Figure 2 shows that one Brazilian graduate program (BGP-1) presents an average number of highly qualified conference publications per faculty member of 4.48, which places it among the top-16 programs in this specific comparison. Five other Brazilian programs (BGP-3, BGP-4, BGP-5, BGP-6 and BGP-2), with an average number of highly qualified conference publications per faculty member between 2.86 and 3.41, present a production equivalent to two North American programs (NGP-16 and NGP-15) and superior to three European programs (EGP-4, EGP-5 and EGP-6).

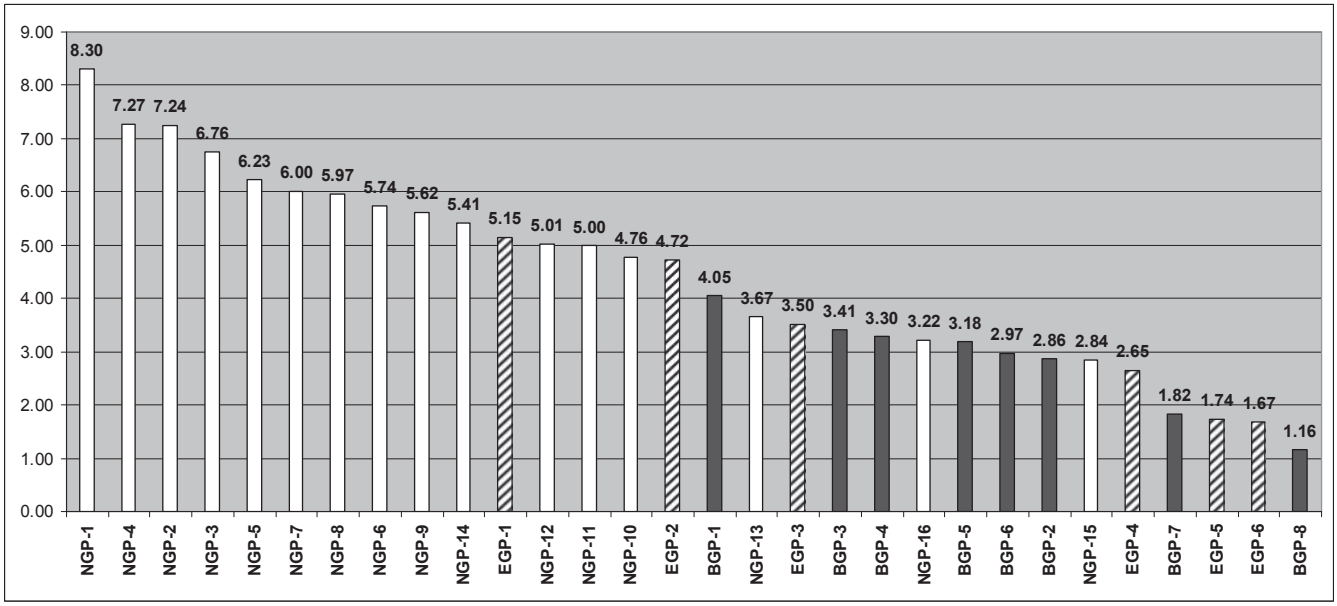

Figure 2: Highly qualified conference publications (Qualis A and B).

Finally, considering now only top qualified publications (journals and conferences Qualis A), as shown in Figure 3, two Brazilian programs (BGP-1 and BGP-4) present an average number of qualified publications per faculty member of 3.76 and 3.62, respectively, which places them in the same range of one North American program (NGP-15) and above another North American program (NGP-16) and three Europeans (EGP-4, EGP-5 and EGP-6). Five other Brazilian programs (BGP-5, BGP-3, BGP-2, BGP-7 and BGP-6), with an average number of qualified publications between 2.47 and 3.20, present a production equivalent to one European program (EGP-4) and superior to two other European programs (EGP-6 and EGP-5).

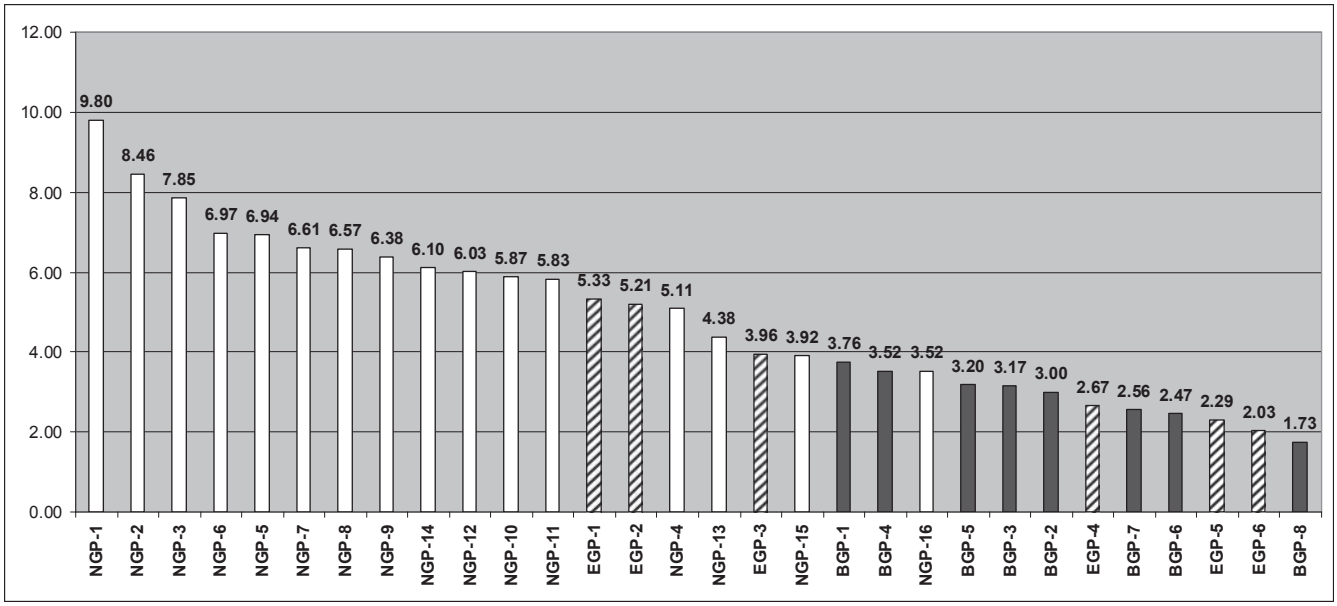

Figure 3: Top qualified publications (journals and conferences Qualis A).

Summarizing, we can say that the scientific production of the Brazilian graduate programs, both in quantity and quality, is comparable to the production of some North American programs and very similar to the production of most Europeans.

3.4 Production of PhD Graduates

Another important index for assessing the productivity of a graduate program is the graduated PhDs per advisor ratio. The graphic in Figure 4 shows this index for the eight Brazilian graduate programs, eight of the North Americans (NGP-1, NGP-2, NGP-9, NGP-10, NGP-11, NGP-13, NGP-14 and NGP-15), and two of the Europeans (EGP-1 and EGP-2), for the triennial 2004-2006. The corresponding data for the Brazilian programs were obtained from CAPES[1]. For the North American and European programs the data were obtained from the home pages of the respective institutions or directly from their officials.

As we can see from Figure 4, considering this index, the Brazilian programs perform very well when compared with the North Americans and Europeans. The top Brazilian program (BGP-1) presents a graduated PhDs per advisor ratio of 2.08, which is close to the one of the top European program (EGP-1) and greater than the ratio of the top North American program (NGP-2). Just after, appear five Brazilian programs (BGP-2, BGP-6, BGP-4, BGP-5 and BGP-3), with graduated PhDs per advisor ratios between 1.08 and 1.69, which place them ahead of six North American programs and one European.

It is important to notice that, although the overseas programs considered here correspond to a subset of the programs analyzed in the previous section, for the sake of this specific comparison they are quite representative since they include the top-2 North American and European programs according to the ranking expressed by Table 4. Thus, these results confirm that the Brazilian CS graduate programs compare well against programs of unquestionable research reputation in North America and Europe.

[1] http://conteudoweb.capes.gov.br/conteudoweb/CadernoAvaliacaoServlet

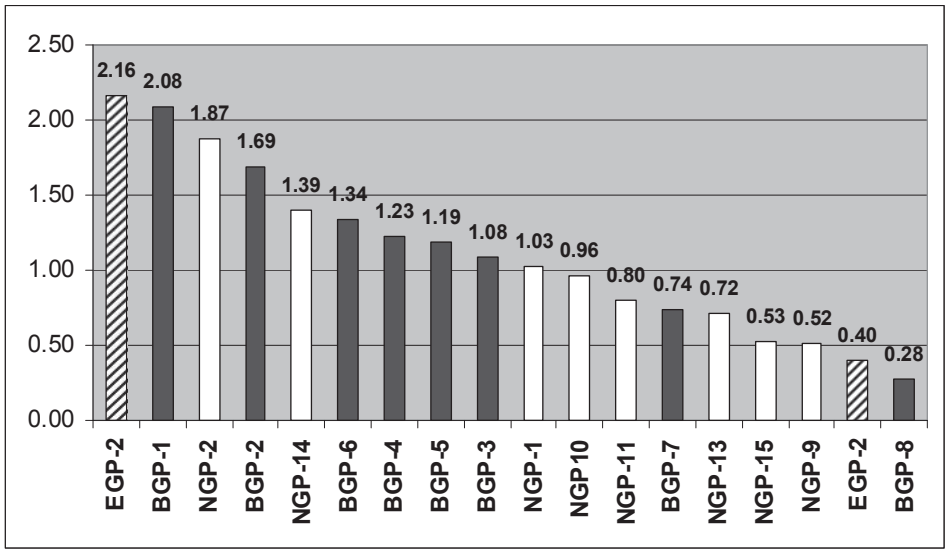

Figure 4: Graduated PhDs per advisor ratio.

4. CONCLUSIONS

Our main aim in this article was to present the results of a study we conducted to assess the quality of the top Brazilian CS graduate programs. Clearly, the number of publications per faculty member and metrics to assess the quality of the publications are important indicators to any evaluation process. However, simply importing metrics commonly used in the evaluation of other fields of knowledge, such as the ISI impact factor, is prone to errors due to key differences that exist in the publication pattern of CS researchers all over the world when compared with those from other fields. Therefore, we argue that only a careful vertical comparison of our CS graduate programs with reputable programs abroad can better assess the programs' excellence.

By using a large data source of CS publications that provides good coverage for most CS subfields, we have shown that the top Brazilian CS programs do well compared with the North American and European programs chosen for the study. We also show that the Brazilian CS programs follow international publication ratios of more than two conference papers per journal article. Therefore any evaluation process cannot ignore conference publications since these are important vehicles for disseminating CS research.

Establishing a ranking for conferences is a difficult and challenging enterprise due to the large diversity of events and the large number of quality indicators to be considered [16]. We have given a first and simple step towards obtaining a reasonable ranking by performing a subjective poll open to all CS graduate faculty members in the country. We have also calculated the ratio of graduated PhDs per advisor of a subset of the graduate programs evaluated. Similar conclusions to those obtained for publication production can be draw from the results. In particular, the Brazilian programs compare well with the European ones.

The evaluation of Brazilian CS programs is a continuing effort and different methodologies should be tried to reduce inaccuracies inherent to any evaluation process. For instance, one slightly different approach than that detailed here and that we would like to pursue in the future is to perform a careful evaluation of a chosen large set (say in the order of 50 to 70) of major programs abroad in terms of their publication records. We would then consider as the CS Qualis the set of all publication venues of the chosen set of these well established programs. A possible measure of goodness for our graduate programs would take into account the publication rate of the faculty abroad against that of the Brazilian counterparts. We plan to investigate this and other methodologies.

Although the increase in the number of CS programs in Brazil in the last decade was fantastic, it was also founded on a healthy growth of the top and well established programs. The conclusions presented here are clear. The CS field in Brazil has reached maturity.

ACKNOWLEDGEMENTS

The authors would like to thank their colleagues, members of the CS CAPES advisory committees, who participated in the last two triennial evaluation processes of the Brazilian CS graduate programs and supported the study presented in this

article, and the graduate and undergraduate students from the Federal University of Minas Gerais, specially Guilherme Vale Ferreira Menezes and Waister Silva Martins, who implemented the systems used to support the data gathering process and helped in many ways during the course of this study. This work was partially supported by the authors' individual research grants from CNPq.

REFERENCES

[1] Butler, D. Free journal-ranking tool enters citation market. *Nature*, 451:3, 2008
[2] Elmacioglu, E.; Lee, D. On Six Degrees of Separation in DBLP-DB and More. *SIGMOD Record*, 34(2): 33-40, 2005.
[3] Hassan, A. E.; Holt, R. C. The Small World of Software Reverse Engineering. In: *Proceedings of the 11th Working Conference on Reverse Engineering*, Delft, The Netherlands, 2004, pp. 278-283.
[4] Klink, S.; Reuther, P.; Weber, A.; Walter, B.; Ley, M. Analysing Social Networks within Bibliographical Data. In: Bressan, S.; Küng, J.; Wagner, R. (Eds.). Database and Expert Systems Applications, 17th International Conference, DEXA 2006, Kraków, Poland, September 4-8, 2006, Proceedings. *Lecture Notes in Computer Science 4080*, Springer, Berlin, 2006, pp. 234-243.
[5] Lee, D.; Kang, J.; Giles, C. L.; On, B-W. Are your citations clean? *Communications of the ACM*, 50(12): 33-38, 2007.
[6] Ley, M. The DBLP Computer Science Bibliography: Evolution, Research Issues, Perspectives. In: Laender, A. H. F.; Oliveira, A. L. (Eds.). String Processing and Information Retrieval, 9th International Symposium, SPIRE 2002, Lisbon, Portugal, September 11-13, 2002, Proceedings. *Lecture Notes in Computer Science 2476*, Springer, Berlin, 2002, pp. 1-10.
[7] Ley, M.; Reuther, P. Maintaining an Online Bibliographical Database: The Problem of Data Quality. In: Ritschard, G.; Djeraba, C. (Eds.). *Actes des Sixièmes Journées Extraction et Gestion des Connaissances*, Lille, France, 2006, pp. 5-10.
[8] Ley, M. Personal communication, 2008.
[9] Mattern, F. Bibliometric Evaluation of Computer Science –Problems and Pitfalls. Invited Talk, SARIT 06: Swiss IT Professors' Day, Swiss Association for Research in Information Technology, Basel, Switzerland, 2006.
[10] Nascimento, M. A.; Sander, J.; Pound, J. Analysis of SIGMOD's Co-Authorship Graph. *SIGMOD Record*, 32(3): 8-10, 2003.
[11] Nebel, B. Ranking? Publikationen, Zitate, Drittmittelprojekte und Promotionen an deustschen Informatikfalkultäten im Spiegel des WWW. *Informatik Spectrum*, 24:234-249, 2001.
[12] Patterson, D.; Snyder, L.; Ullman, J. Best Practices Memo: Evaluating Computer Scientists and Engineers for Promotion and Tenure. *Computer Research News*, Sept. 1999. http://www.cra.org/resports/tenure_review.pdf.
[13] Patterson, D. A. The Health of Research Conferences and the Dearth of Big Idea Papers. *Communications of the ACM*, 47(1): 23-24, 2004.
[14] Rahm, E.; Thor, A. Citation analysis of database publications. *SIGMOD Record*, 34(4):48-55, 2005.
[15] Rossner, M.; Van Epps, H.; Hill, E. Show me the data. *The Journal of Cell Biology*, 176(6): 1091-1092, 2007.
[16] Zhuang, Z.; Elmacioglu, E.; Lee, D.; Giles, C. L. Measuring conference quality by mining program committee characteristics. In: *Proceedings of the 7th ACM/IEEE Joint Conference on Digital Libraries*, Vancouver, Canada, 2007, pp. 225-234.